

Violence and Ethnic Segregation: A Computational Model Applied to Baghdad*

NILS B. WEIDMANN
Peace Research Institute Oslo

IDEAN SALEHYAN
University of North Texas

December 8, 2011

Article accepted for publication in *International Studies Quarterly*.

Abstract

The implementation of the US military surge in Iraq coincided with a significant reduction in ethnic violence. Two explanations have been proposed for this result: The first is that the troop surge worked by increasing counterinsurgent capacity, whereas the second argument is that ethnic unmixing and the establishment of relatively homogenous enclaves were responsible for declining violence in Baghdad through reducing contact. We address this question using an agent-based model that is built on GIS-coded data on violence and ethnic composition in Baghdad. While we cannot fully resolve the debate about the effectiveness of the surge, our model shows that patterns of violence and segregation in Baghdad are consistent with a simple mechanism of ethnically motivated attacks and subsequent migration. Our modeling exercise also informs current debates about the effectiveness of counterinsurgency operations. We implement a simple policing mechanism in our model and show that even small levels of policing can dramatically mitigate subsequent levels of violence. However, our results also show that the timing of these efforts is crucial; early responses to ethnic violence are highly effective, but quickly lose impact as their implementation is delayed.

* *Authors' notes:* Previous versions were presented at the 2010 Annual Meeting of the American Political Science Association, Washington DC, the 2009 Meeting of the Peace Science Society, and the New Faces in Political Methodology III conference, Penn State University, May 2010. We are grateful to the *Empirical Studies of Conflict* project at Princeton University (especially Jacob N. Shapiro and Joseph H. Felter) for access to the SIGACTS data, and to the International Conflict Research group at ETH Zurich for providing the computational infrastructure. Nils Weidmann gratefully acknowledges support from the Air Force Office of Scientific Research (grant #FA9550-09-1-0314). This and prior versions of this paper benefited significantly from comments by Navin Bapat, Andrew Enterline, Cullen Hendrix, Jacob Shapiro and Richard Stoll, as well as three anonymous reviewers at ISQ. An online appendix is available at <http://nils.weidmann.ws/>.

Ethnic conflicts often baffle observers as groups that had once lived together in relative peace commit horrendous acts of violence against one another. In places such as the former Yugoslavia, Rwanda, and Burundi, people co-existed in mixed neighborhoods for years before mass violence led to a hardening of ethnic identities and the intensification of ethnic hatred (see Fearon and Laitin 2003). Ethnic violence also escalated after the 2003 United States-led invasion of Iraq, as the dominant mode of violence shifted from attacks against an occupying force to sectarian killings by Sunnis, Shias, and Kurds. The bombing of the Shia al-Askari Mosque in Samarra in 2006 precipitated a spiral of violence that, by 2007, looked to be out of control, leading many in the US to call for a complete withdrawal of troops. Sunni–Shia violence in Baghdad was especially pronounced as the feeble government and inadequate US forces were unable to prevent frequent attacks by the Mahdi Army, Al-Qaeda in Iraq, and several other militias. At the same time, the nature of Baghdad changed from a city where the two main ethno-religious groups resided in mixed neighborhoods to one with well-defined ethnic neighborhoods, separated in some areas by a security wall.¹

Beginning in 2007, the United States adopted a new strategy (often called the “surge”) in which it increased the number of troops in Iraq, hoping to contain the level of violence through a significant security presence. In addition, the US sought accommodation with certain militias, changed counterinsurgency tactics, and placed greater emphasis on institution building. The training of Iraqi government forces also increased significantly. These measures preceded a significant decline in violence in 2008 and 2009, particularly in Baghdad, although the exact causes of this decline are debated (Matthews and Klein 2008; Ricks 2009). These patterns of ethnic violence and residential segregation lead to important questions for the academic and policy communities. What explains the spike of violence in Baghdad and the subsequent decline in the number of attacks? How do these events relate to inter-ethnic violence and forced migration elsewhere? What are the most important measures needed to prevent mass killings and ethnic segregation? While we focus on Baghdad in this paper, our analysis will hopefully shed light on the dynamics of ethnic conflict more generally. For scholars and policymakers, a proper understanding of the Iraqi case—particularly the effectiveness of different interventions—can help guide future action and avoid potential pitfalls.

Given a single case, it is difficult to provide a definitive answer as to why violence escalated and declined as it did in Baghdad. However, we can draw some conclusions about plausible combinations of factors that may have produced the observed dynamics of conflict. Did improved counterinsurgency work? Or, did ethnic partition help reduce levels of violence? To address these questions, we rely on a computational model, or computer representation, where we experiment with different hypothetical scenarios informed by current theories of ethnic violence.² In particular, we are interested in

¹Bright, Arthur. 2007. Baghdad’s Sunni/Shiite Security Wall. *Christian Science Monitor*, April 20.

²For other applications of computational modeling to ethnic conflict and war, see Hammond and Axelrod (2006); Bennett (2008); Epstein (2002); Geller and Alam (2010); Cederman (2008) and Lustick, Miodownik and Eidelson (2004).

how combinations of parameters interact to produce insurgent violence as well as ethnic settlement decisions. Rather than a completely artificial world, we use geo-referenced data on ethnic settlement patterns and violent attacks in Baghdad to inform our model. Although we cannot perfectly re-run history, we can construct a model based on actual data and compare the results of our artificial simulations with general empirical patterns. While it is difficult to draw firm conclusions about a single case, our models allow us to assess theoretically interesting combinations of parameters and we present a novel approach to validating the output of our simulations. Therefore, our aim is to shed light on theoretical debates in the literature on insurgency and counterinsurgency, and provide some insights—albeit tentative—into the dynamics of violence in Baghdad.

We create a computational model with very simple assumptions about intergroup behavior and local rules of interaction. We populate our artificial world with “Shia” and “Sunni” agents, along with a small share of violent individuals in each ethnic community who attack members of the other group. Agents are assigned to neighborhoods according to actual ethnic maps of Baghdad. When attacks against one’s group occur, agents in that neighborhood can respond by moving to other, safer neighborhoods, or by retaliating with their own violence. Thus, we can represent processes of violence as well as migration and settlement patterns. We can then determine how close our computational model comes to resembling actual data on violence and ethnic settlement patterns in Baghdad in order to draw conclusions about the parameter combinations at work. Even though we cannot say for certain what “actually” occurred in Baghdad and which factors were most important in turning the tide of the insurgency, our models can help to inform the debate over the efficacy of counterinsurgency operations.

In addition to our Sunni and Shia agents, we allow for a state policing mechanism to punish insurgents. According to many scholars, a strong central state is needed to enforce order, punish violent individuals, and maintain peace among ethnic groups (Fearon and Laitin 2003; Posen 1993; Muller and Seligson 1987; Tilly 1978). Thus, we introduce a state actor—a counterinsurgent force—who can punish insurgents for their transgressions and remove them from the simulation, thereby reducing discontent and fear among our agents. The ability to detect and capture insurgents is a parameter that we manipulate in our simulations. By introducing counterinsurgency efforts as a variable, we are able to shed light on debates about the efficacy of the “surge” and of increased state capacity in limiting violence.

In the following sections, we discuss current theories of interethnic peace and violence. Then, we introduce our agent-based model and its application to the Baghdad case. Rather than a completely artificial world, we ground our model in the current theoretical literature and use actual data from Baghdad to inform our modeling exercise. After examining the parameter combinations that could plausibly reproduce patterns of violence in Baghdad, we turn to a counterfactual analysis. In particular, we look at how the level of policing affects violence and ethnic segregation, altering both the effectiveness of counterinsurgency as well as the timing of its implementation. The simulations suggest three main findings. First, even without assuming that people have a preference for living with co-ethnics, we show that violence can dramatically increase ethnic segre-

gation as civilians search for safety. Second, we are able to show that ethnic segregation places natural limits on the extent of violence; that is, as communities become more perfectly segregated, violence declines even in the absence of effective counterinsurgency operations. Third, we show that small increases in state policing can significantly reduce the level of violence, but that it is far more effective if implemented early on in an insurgency. Our final section offers some concluding observations.

Theories of Interethnic Peace and Conflict

In multiethnic societies, most of the time, the vast majority of ethnic groups do not come into conflict with others. While ethnic divisions are ubiquitous, organized violence between communal groups is relatively rare. Ethnic cleavages are clearly necessary for ethnic conflict, but they are not a sufficient cause, even when historical grievances run deep (Collier and Hoeffler 2004; Fearon and Laitin 2003; Mueller 2000; Wilkinson 2004). For this reason, scholars in the ‘Hobbesian’ tradition have argued that diminished state capacity to control violence among communal groups is an essential ingredient in fostering conflict (Fearon and Laitin 2003; Herbst 2000; Posen 1993; Rose 2000). Even when latent tensions between groups exist, the vast majority of people would never contemplate murdering their neighbors and co-workers; rather, a small set of extremists are predisposed to violence and a strong state is needed to punish these offenders (Lake 2002; Mueller 2000). However, once state capacity erodes and extremists begin to attack members of the out-group, violence can quickly spiral out of control, leading ethnic identities to harden (Fearon and Laitin 2003). As the ethnic cleavage becomes dominant, even moderates will seek safety from attack by turning to their co-ethnics for support. Some—particularly young males—may be inclined to join militant groups for safety as well as to seek revenge (Urdal 2006). Those who are not willing or able to fight may seek safety in numbers by residing among their ethnic brethren.

This is precisely the pattern that was seen in Iraq. After the 2003 invasion, there were too few US troops on the ground to prevent interethnic conflict (Enterline and Greig 2007), and the nascent Iraqi government was too feeble to provide order on its own. Sunni and Shia extremists began to compete for influence, and their voices became louder than that of moderates. Many observers pointed out that prior to the invasion, Baghdad’s neighborhoods were quite mixed and intermarriage between Sunnis and Shias was common. Some saw intermarriage and ethnic mixing as a stabilizing force that could prevent an all-out civil war.³ But with rising violence one observer noted that: “mixed couples who symbolize Iraq’s once famous tolerance are increasingly entangled by hate.”⁴ As extremists attacked members of the other sectarian group, moderates were forced to choose sides if only for safety. Interethnic warfare reached new heights after the bombing of a Shia Mosque in Samarra in 2006, which caused widespread fear and anger among

³Telhami, Shibley. 2005. “Rush To Stabilize May Backfire in Polarized Iraq.” *San Jose Mercury News*, October 30, 2005. Opinion.

⁴Raghavan, Sudarsan. 2007. “Marriages Between Sects Come Under Siege in Iraq,” *Washington Post*, March 4.

the Shia community. Many in Baghdad joined the various militias operating there, while others moved to neighborhoods populated by their co-ethnics or fled the city entirely. Many of Baghdad's neighborhoods were "ethnically cleansed" as one or the other group was forced out of the area (Agnew et al. 2008).

Although normatively questionable, ethnic unmixing may provide one explanation for the decline of interethnic violence. Some scholars argue the best way to secure peace is to create a partition, or physical separation, between ethnic groups (Chapman and Roeder 2007; Kaufmann 1996; Sambanis 2000). According to this argument, once identities have hardened through violence and hatred runs deep, ethnic groups engaged in bloody conflicts will not be able to trust one another enough to co-exist peacefully. The best way to secure peace, then, is to separate groups through some form of partition. Partition may entail the complete separation of the state into two new entities. One could also argue that ethnic segregation within a state, or even a city, serves many similar functions by reducing contact between ethnic groups and establishing autonomous enclaves. Along these lines, studies of race riots in the United States have shown that decreasing levels of urban racial segregation increases contact, competition, and the frequency of riots (Olzak, Shanahan and McEneaney 1996). While ethnic segregation may not be desirable, or even stable in the long-term, it may serve as an immediate fix to large-scale bloodshed.

These claims have been given renewed attention in the context of the Iraq war. Edward Joseph and Michael O'Hanlon (2007) argue for a "soft partition" of Iraq and claim that "physical separation boosts security, but keeping the communities cheek-by-jowl makes residents angry and resentful (p. 5)." Even US military planners agreed, arguing that the Baghdad security wall, which separates ethnic communities, "is one of the centerpieces of a new strategy by coalition and Iraqi forces to break the cycle of violence."⁵ Researchers at the University of California, Los Angeles, using satellite imagery of nighttime lights in Baghdad to determine settlement patterns, concluded that the military surge was ineffective, but rather, "the diminished level of violence in Iraq since the onset of the surge owes much to a vicious process of interethnic cleansing (Agnew et al. 2008:2295)." In other words, Agnew et al. make a strong case that ethnic segregation in Baghdad, as discerned through satellite imagery, was in place prior to the troop surge and may explain declining levels of violence.

Therefore, one hypothesis, which we turn to later, is that ethnic segregation reduces ethnic violence. However, as an alternative to the ethnic cleansing of neighborhoods and the creation of mutually-exclusive enclaves, steps may be taken to strengthen the state and reestablish control. If the lack of a robust deterrent capability and inability to punish defectors led to ethnic violence in the first place, then measures to bolster the central government may help to reverse violence. Many would agree that—assuming that it is feasible—an impartial, competent state to resolve disputes and establish a monopoly of legitimate force is normatively preferable to ethnic cleansing. State strength clearly means many things, including the ability to provide public goods and execute administrative functions, but one important component of state strength is the ability

⁵Wong, Edward and David Cloud. 2007. "US Erects Baghdad Wall to Keep Sects Apart." *New York Times*, April 21.

to punish criminals and violent individuals (Hendrix 2010). If the state can detect and defeat insurgents as well as provide civilians adequate levels of protection (Kalyvas 2006), conflict levels may come down, internal displacement may subside, and over the long-run, ethnic trust may be established. In his focus on security institutions, Herbst (2004:367) argues that in order to reduce civil violence it is, "...essential to promote institutions—notably the police and the military—that are immediately responsible for order." Fearon and Laitin (2003) make a similar argument, and advocate increasing the state's capacity to govern its territory and root out insurgents.

The troop surge in Iraq promised to accomplish many of these functions. The United States sent more than 20,000 troops to Baghdad in 2007 and, along with Iraqi forces, implemented Operation *Fardh al-Qanoon* (Operation Imposing Law). Baghdad was divided into several zones and the US military worked to clear each area of hostile Sunni and Shia militias. In particular, military presence in residential areas was improved by setting up "combat outposts" in different neighborhoods of Baghdad.⁶ The combat outposts were used to ensure a quick response in cases of insurgent attacks, but also to work more closely with the local population. US forces may be seen as a "neutral" third party enforcement mechanism, not tied to any ethnic or religious faction, but with a mandate to enforce order over the entire city. Simultaneously, the US engaged in efforts to strengthen the Iraqi state by training troops and police officers while encouraging broad participation in the government by all groups. The stated goal was to secure Baghdad and provide the space needed for ethno-sectarian groups to find an acceptable accommodation and integrate into a strong, unified force that could provide security for the country as a whole.

Many have argued that the troop surge in Iraq was a success. US and Iraqi casualties declined significantly after thousands of troops were deployed to Baghdad and elsewhere in 2007. In his report to Congress, the US Commander in Iraq, General David Petraeus highlighted improved counterinsurgent capacity, stating, "One reason for the decline in incidents is that Coalition and Iraqi forces have dealt significant blows to Al-Qaeda-Iraq... We have also disrupted Shia militia extremists..."⁷ In the *Wall Street Journal*, US Senators John McCain and Joe Lieberman argued, "As Al-Qaeda has been beaten back, violence across the country has dropped dramatically... As the surge should have taught us by now, troop numbers matter in Iraq."⁸ Finally, in a speech in 2008, Lt. General Ray Odierno attributed the success of the surge to, "attacking the enemy," "establishing and maintaining a presence in places that had long been sanctuaries of Al-Qaeda," and "going after Shia extremists." He places special emphasis on the training of Iraqi forces, stating that, "the surge of Coalition forces also helped bring about a surge in Iraqi Security Force capacity."⁹ As these statements indicate, increasing the number of US

⁶Londono, Ernesto. 2007. "In Baghdad, a Flimsy Outpost." *Washington Post*, March 22.

⁷"Report to Congress on the Situation in Iraq", General David H. Petraeus, Commander, Multi-National Force-Iraq. 10–11 September 2007.

⁸McCain, John and Joe Lieberman. 2008. "The Surge Worked," *The Wall Street Journal*, January 10. Opinion.

⁹Lt. General Raymond T. Odierno. 2008. "The Surge in Iraq: One Year Later." Lecture on National Security and Defense, The Heritage Foundation, March 13. Accessed online at <http://www.heritage>.

troops, while simultaneously improving the efficacy of the Iraqi government, were seen as important components of the counterinsurgency effort.

The troop surge in Iraq was clearly a multi-pronged strategy that went beyond simply increasing troop levels. However, one of the most important objectives of the surge was to increase US and Iraqi ability to identify and eliminate insurgent forces. Part of this strategy included a change in tactics that de-emphasized heavy firepower and focused on establishing ties to local communities (see Lyall and Wilson 2009). This change in thinking on counterinsurgency operations and empowering locals to help with policing efforts were part of an overall strategy to better identify and defeat violent extremists. Building state capacity, moreover, is widely viewed by scholars (discussed above) as a key component of deterring and ending an insurgent campaign.

Therefore, there are two distinct arguments that have been made with respect to declining violence in Baghdad. The first is that the troop surge worked by increasing counterinsurgent capacity. Increasing the number of security forces in Baghdad, strengthening the Iraqi state, and changing tactics to better identify insurgents all helped to remove extremists from the population and reduce levels of violence. The second argument is that the troop surge was irrelevant. Instead, according to this argument, ethnic unmixing and the establishment of relatively homogenous enclaves were responsible for declining violence in Baghdad through reducing contact. Moreover, ethnically homogenous neighborhoods make it more difficult for insurgents from the other group to operate, since insurgents rely on networks of local support and it is relatively easier to identify individuals that do not “belong” to the community. Unfortunately, the Baghdad case is overdetermined as both of these processes occurred, and it is difficult to disentangle the relative strength of these claims with a single case. While we cannot fully resolve this debate, we develop an agent-based model that takes into account levels of violence, ethnic settlement patterns, and counterinsurgency effectiveness, to assess theoretically interesting combinations of parameters and we validate our modeling exercise with actual data taken from Baghdad.

In addition to looking at how ethnic settlement patterns may have influenced violence in Baghdad, we are also able to assess counterfactual claims regarding state capacity. In particular, we assess how effective augmenting counterinsurgent ability is in reducing violence through experimental “treatments” in our simulation. For any given level of counterinsurgent capacity, how much violence will we see? We are also able to introduce enhanced policing at various periods in the simulation to address important questions regarding the timing of counterinsurgent campaigns. Some, most notably US Army Chief of Staff Erik Shinseki, argued that a strong troop presence should have been put in Iraq immediately after the 2003 invasion in order to prevent widespread violence. Along these lines, Enterline and Greig (2007) find that a significant troop presence early in a counterinsurgency campaign, the “Shinseki plan,” is far more effective than troops added at a later date. By varying the level and timing of state policing, we can see how the dynamics of violence change in our models.

org/Research/Lecture/The-Surge-in-Iraq-One-Year-Later (April 1, 2010).

Model

A multitude of influences affected the dynamics of violence and segregation in Baghdad, and any attempt to understand the micro-mechanisms behind them seems difficult at best. Recognizing that we will not be able to perfectly “re-run” Baghdad history, we aim to find out whether the general interrelationships between violence and migration we discussed above bear some resemblance to the observed patterns. Similar to many other applications in the social sciences, this attempt needs to reflect the micro/macro distinction made in Coleman’s “bathtub” (Coleman 1990): the mechanisms we postulate are located at the micro-level, with insurgents deciding to attack in certain neighborhoods, triggering migration of civilians as a result. However, we can only observe the result of this process at the macro-level—the uneven distribution of violence across the city, and the stepwise segregation of ethnic groups.

One way to fill the gap between micro-mechanisms and macro-outcomes is by means of computational (or more precisely, agent-based) modeling (ABM). Cederman (2001) defines ABM as “a computational methodology that allows the analyst to create, analyze, and experiment with, artificial worlds populated by agents that interact in non-trivial ways and that constitute their own environment” (p. 16). Still, the implementation of this methodology far from standardized. First, ABMs differ widely as to the complexity of the computational agents and their behavior. At the lower end of this complexity scale is Schelling’s (1971) famous model of neighborhood segregation, where the only feature of an agent is its “color” and one simple rule—affinity for living with agents of the same color—drives their behavior. Axelrod’s (1984) computer simulations are also quite simple, with agents playing the familiar Prisoners Dilemma game in multi-player tournaments. These models are easy to understand, but may be too basic to represent reality. More recent examples such as Geller and Alam (2010) or Cioffi-Revilla and Rouleau (2010) have grown far more complex, featuring representations of entire countries with political and environmental sub-systems. The intention of this approach is to create a more realistic representation of the world, but this comes at a cost. With a model guided by dozens if not hundreds of parameters, one quickly encounters what de Marchi (2005) calls the “curse of dimensionality”: the complexity of the model itself makes it difficult to understand the interplay of its elements.

Along a second dimension, ABMs differ with respect to the level to which they incorporate empirical data, if at all. Computational models can serve as purely heuristic devices, to illustrate how a theorized micro-process leads to the emergence of macro-outcomes. Again, Schelling’s (1971) model is an example here, but more recent ones include Cederman’s (1997) model of war and state formation, Bhavnani’s (2006) norms model of mass participation in interethnic violence, or Siegel’s (2009) model of collective action in social networks. Models that incorporate empirical data can do so by comparing a generated output to an observed pattern. For example, Cederman (2003) shows how a simple model of interstate war generates a power-law distribution of war sizes, similar to what has been observed in reality. A similar approach was taken by the authors of the “Artificial Anasazi” model (Dean et al. 2006), who managed to reproduce

the settlement pattern of an ancient culture in an agent-based model. Geller, Rizi and Latek (2011) use actual data from Afghanistan to validate a computer simulation of the interaction between drug production and corruption. Despite some promising attempts, the empirical evaluation of computational models is the exception rather than the norm. As we will outline in the next paragraphs, our computational model adopts an intermediate approach between simplicity and complexity by keeping the number of moving parts within manageable ranges. At the same time, we conduct a model fitting exercise by letting empirical data guide the selection of plausible parameters. In the following paragraphs, we describe our implementation of this model, and the next section discusses the incorporation of empirical data.

The fundamental building blocks of our model is a space with a set of locations that represent the neighborhoods of Baghdad, and the agents populating that space. Agents can belong to one of two ethnic groups, Sunni or Shia. In addition, Sunni and Shia agents are designated to be either insurgents or civilians. The basic dynamics of the model are as follows. Insurgents of one group attempt attacks against civilians of the other group. The success of an attack depends on the local ethnic makeup of the location at which it is conducted. This generates fear in the targeted population, which leads civilian agents to consider migration to safer places in the city. This in turn alters the ethnic configuration, which can either increase or constrain the susceptibility for violence in these locations. In the following, we first describe the model structure (the agents and the model space) and then turn to the dynamics of the model.

Model Space and Agent Population The model space consists of N locations i that are populated by agents of two groups j , $j \in \{\text{Sunni}, \text{Shia}\}$. For a given group j , we refer to the other group as $\neg j$. The initial total number of agents at each location is assigned at the beginning of the simulation, and consists of both insurgents I and civilians C . The proportion of insurgents π is a model parameter that determines the size of the insurgent population. Locations differ with respect to their Sunni/Shia balance. As we will describe below, the initial ethnic balance is taken from an ethnic map of Baghdad. A certain percentage of agents at a location is randomly chosen to be insurgents. This percentage is equal across neighborhoods, and is a model parameter that is fixed at the beginning of a model run. Throughout the simulation, we measure the ethnic composition of a location at a given time simply as the proportion $p_{i,j}$ of agents of group j at location i . Once the space has been populated with insurgents and civilians of different groups, the model proceeds in a series of time steps t as described in the following paragraphs.

Violence At the beginning of a time step t , each insurgent $I_{i,j}$ attempts to stage an attack at its current location i against the members of the other group $\neg j$. Even though all insurgents aim to attack, only few of these will be successfully carried out. We assume that the probability of success of an attack depends on the local ethnic configuration of the location, as measured by the proportion of the insurgent’s co-ethnics in the respective location $p_{i,j,t}$. Theoretically, this effect could take one of two forms. On the one hand,

interethnic violence could be carried out as ethnic cleansing of small minorities of the other group, for example by violent raids. If that is the case, we expect the impact of the proportion of co-ethnics to be positive, implying that as neighborhoods become increasingly homogenous, there is more pressure for the minority group to leave. On the other hand, attacks against a group could increase with its populations share as attackers try to cause maximum damage to as many of these individuals as possible. A Shia bombing of a marketplace in a Sunni neighborhood would be an example for this type of attack, as it attempts to maximize casualties in the other group. If most attacks take this form, the effect of the proportion of co-ethnics would be negative. These are competing claims that we can evaluate in our model. In addition to this local impact on attack success, we want to allow for the possibility that the insurgent attacks are unaffected by local conditions, and are simply occurring with a constant rate of success across all neighborhoods. In order to incorporate these competing alternatives, we let a logit equation determine the probability of success:

$$pAttack_{i,j,t} = \text{logit}^{-1}(\alpha_0 + \alpha_1 p_{i,j,t}) \quad (1)$$

We then simply take a Bernoulli draw with the computed probability to determine whether an attack is successful. The intercept α_0 and the coefficient α_1 are model parameters selected at the beginning of a simulation run, which then drive all attack behavior during the subsequent time steps. Once all agents have attempted to attack, they relocate to a randomly selected location.

The logit model allows for a flexible specification of the attack mechanisms we are interested in. First, as described above, it makes it possible to determine whether a constant attack probability is sufficient to explain the patterns we observe in Baghdad, or whether the local ethnic mix does indeed have an effect as we hypothesized. If the latter, we should estimate α_1 to be different from 0. Depending on the sign of α_1 , the effect can go in one of two directions. First, if $\alpha_1 > 0$, a higher proportion of co-ethnics makes attacks more likely to be successful. This is consistent with the first mechanisms of ethnic violence we described above, where violence is carried out to cleanse neighborhoods dominated by one group from members of the other group. On the contrary, if $\alpha_1 < 0$, lower proportions of co-ethnics will lead to more attacks. This is essentially the car bomb violence mechanism, where the insurgent attempts to maximize damage to the other group, but at the same time avoids harm to his co-ethnics.

The result of the attack sub-procedure is for each neighborhood, a number of attacks $a_{i,j,t}$ against a particular group. This in turn affects the civilian population as we elaborate in the next paragraph.

Migration Once insurgents have carried out their attacks, civilians react to the threat and decide whether to relocate to safer areas. We consider two potential influences of insurgent violence. First, attacks in a given neighborhood directly influence the level of fear experienced by members of the victimized group. Here, it is the experienced violence $a_{i,j,t}$ that affects an agent's decision. Second, we argue that violence can have effects beyond the borders of a single neighborhood. Violent acts against one group

could instill fear in agents in other neighborhoods nearby. We model this spillover of fear by including the spatially weighted number of attacks as a determinant of a civilian’s migration decision. Following common practice, we compute the spatial lag $a_{i,j,t}^W$ as the average number of attacks against group j across all other neighborhoods, weighted by the normalized inverse distance to neighborhood i . This spatial lag will take high values if violence in proximate locations against group j is high, but low values if this violence occurred in distant neighborhoods. Again, we use a logit model to compute the probability of migration for a civilian of group j at location i , as given by the following equation:

$$pMigration_{i,j,t} = \text{logit}^{-1}(\beta_0 + \beta_1 a_{i,j,t} + \beta_2 a_{i,j,t}^W) \quad (2)$$

As for the attack model above, the migration model allows us to separate the baseline migration from the one that is (directly or indirectly) induced by ethnic violence. The constant β_0 and the coefficients β_1 and β_2 are model parameters fixed throughout a simulation run. After comparing our model runs to actual data, if we observe β_1 and β_2 to be positive in empirically plausible runs, we interpret this as evidence in support of our proposed violence-induced migration mechanism.

If an individual chooses to migrate, she does so by selecting a location that appears to be safer than the one she is currently residing at. Having observed the insurgent-generated violence in the current time step, she moves to a randomly selected location z that experienced lower levels of violence against her group, i.e. with $a_{z,j,t} < a_{i,j,t}$. Importantly, our models make no assumptions about ethnophilia, or an inherent desire to live with one’s co-ethnics, in contrast to Schelling’s (1971) models. Migration decisions are driven by safety concerns, not ethnic attachments.

In sum, our model relies on a set of 6 parameters. First, we have the proportion of insurgents π in the model. Next, the attack model is guided by the constant α_0 and the coefficient α_1 , and the migration model by β_0 , β_1 and β_2 . The purpose of our empirical analysis will be to find out which parameter settings can explain the variation in violence and segregation we observe in Baghdad, and in particular, if segregation alone is sufficient to explain the decrease in violence. Later, in our counterfactual experiments, we will add a simple punishment mechanism to the model to evaluate a different way of violence reduction and its effectiveness.

Data

Similar to many existing approaches (Schelling 1971; Cederman 1997; Epstein and Axtell 1996), our computational model as described above could be implemented in an entirely artificial space that uses cells as spatial units. However, since our intention is to apply the model to a particular case (Baghdad), we create a model that shares certain characteristic features of our case: first, we create the model based on the real geography of the city, using neighborhoods as unit of observation. Second, we use observed incidents of sectarian violence to get the distribution of violence across neighborhoods. Third, the ethnic composition of neighborhoods in the model is based on high-resolution ethnic

maps of Baghdad. This section explains our data sources and how they were integrated in the model.

Unit of Observation The model operates on the set of 85 Baghdad neighborhoods, which were geo-referenced from a map provided by the UNs Humanitarian Information Center for Iraq (HIC).¹⁰ Using these neighborhoods as our basic units of observation, we incorporate real data on the level of violence and the ethnic composition of each neighborhood.

Violence Data on incidents of sectarian violence was obtained from the Iraq SIGACTS database.¹¹ From all the events we selected those that correspond to sectarian violence by excluding (i) attacks involving US/Iraqi security forces, and (ii) minor/non-deadly event categories such as “looting”. The SIGACTS database in its present version covers the period from April 2004 through February 2009, but the information on violence initiator and target is available only in a comprehensive version of the database that covers April 2006 through September 2007. SIGACTS contains data at the event level. Each line in the main table list one event along with its spatial and temporal coordinates. Using the spatial coordinates of an event, we matched events to our units of observation to compute event counts of violence at the neighborhood level.

Ethnicity We obtained estimates about the ethnic distribution at the neighborhood level from M. Izady at the *Gulf 2000* project at Columbia University.¹² The project provides maps of the predominant ethnic group in a neighborhood for different periods (2003, 2006, early 2007, late 2007, 2009). The map shows areas of Baghdad either as mixed, Sunni- or Shia-dominated.¹³ We geo-referenced the maps and created a similar variable (Sunni/Shia/mixed) at the neighborhood level. Unfortunately, the data do not provide precise estimates of group proportions—that is, we cannot say if a “dominated” neighborhood is 70% or 80% populated by one group—but we can use the maps to ascertain overall settlement patterns.

Period of Observation The bombing of the Askari Mosque in Samarra in February 2006 is commonly taken as the most significant trigger of violence between Sunni and Shia groups in Iraq. During that year and especially in the autumn, violence levels in Baghdad increased steadily. This was part of the reason for the US to adopt a new

¹⁰Map available online at <http://www.humanitarianinfo.org/iraq/maps/280a\%20A4\%20Baghdad\%20districts\%20neighbourh\%20300dpi.pdf> (accessed February 4, 2010). Of the original 89 units, 4 primarily industrial areas with few residential homes were excluded.

¹¹Joseph H. Felter and Jacob N. Shapiro, Princeton University, 2008; see Berman, Shapiro and Felter (2011) for an application.

¹²M. Izady, *Ethnic-Religious Neighborhoods in Metropolitan Baghdad*. <http://gulf2000.columbia.edu/maps.shtml> (accessed February 4, 2010).

¹³The maps also show Christian populations in Baghdad. However, there is no homogenous Christian neighborhood in Baghdad due to their small population share (< 5%), so we do not treat them as a separate group in our model.

military strategy, the surge. Part of the implementation of the surge was the additional deployment of US military personnel, primarily to Baghdad, where most of the violence in Iraq occurred at that time. This deployment started in late January 2007, increased rapidly in February and March, and reached its peak in June 2007. Thus, the Samarra attacks in February 2006 can mark the onset of inter-ethnic violence in Iraq, until the implementation of the surge directly aimed to stop the hostilities between groups. Thus, we use data from the period between the Samarra bombings and the onset of the surge for our calibration of the model. Figure 1(a) illustrates the ethnicity and violence data for our observation period.

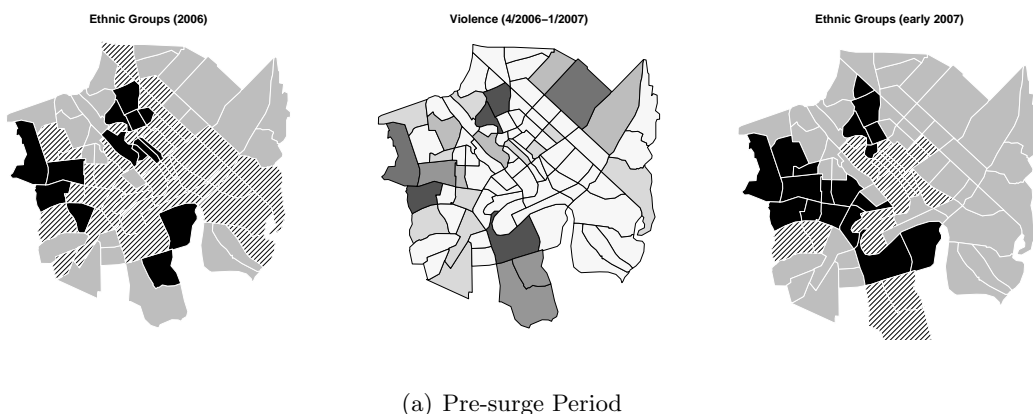


Figure 1: Empirical data used for seeding and validation of the model. Ethnic maps show Shia (grey), Sunni (black) and mixed neighborhoods (striped). The level of violence by neighborhood is displayed in different grey shades (center map).

Results

Our approach of incorporating empirical data into a computational model shares the essential features of the Bayesian approach to statistics, in which we start with a theoretical prior and update based on the incorporation of empirical data. Essentially, our approach is as follows:

1. We develop a simple model of violence and migration, with a set of parameters determining its behavior. This model was described above.
2. The model is run with parameter combinations sampled randomly from a large parameter space Θ , assuming no prior knowledge of the distributions of the individual parameters (our “noninformative priors”).
3. We select those parameter combinations $\Theta' \subset \Theta$ that produce empirically plausible runs, i.e. those that come closest to the empirical patterns in Baghdad.

4. We examine the distribution of the parameter values in Θ' (our “posterior distributions”). The examination of this posterior distribution of the model parameters reveals if a parameter is necessary to generate empirically plausible model outcomes, and if so, whether it has the expected sign.

Model Initialization and Dynamics

As described above, we conduct a many different runs of the model with different parameter vectors θ , and evaluate which parameter settings come closest to the observed patterns of violence and segregation in Baghdad, i.e. produce the best fit. Using the data we have available for our observation period (ethnic distribution at the beginning and the end of the period, and distribution of violence), we initialize the model based on these data, and compute its goodness of fit as the simulation unfolds. More precisely, we initialize the model run with the ethnic distribution at the beginning of this period. This is done by populating each neighborhood with 100 agents, assigning a 50/50 proportion of ethnicity in mixed neighborhoods (as given by our ethnic maps), and an 70/30 proportion in neighborhoods dominated by one group.¹⁴ Since we do not have data on absolute population at the neighborhood level, this procedure omits differences in population totals across neighborhoods. Figure 2 shows a screenshot of the model after initialization.




Figure 2: Screenshot of the computational model, initialized using the 2006 ethnic map. The shading indicates the group distribution (bright: Shia, dark: Sunni).

Once it has been initialized, the model is run for a certain number of time steps. As described in detail above, each step consists of the insurgents carrying out their attacks, and the civilians reacting to violence. All of these actions are determined by the model parameters θ set for the current run. Empirically plausible runs are those that approximate both the ethnic distribution at the end of the observation period, and the spatial pattern of violence during this period. In other words, we want those runs that

¹⁴Results are robust to using alternative proportions for dominated neighborhoods, see online appendix.

minimize the deviation of the model from the observed ethnic distribution at the end of the observation period, and second, those that at the same time produce a distribution of violence that is as close as possible to the observed one. Thus, we compute two error measures for the model. $\epsilon_e(t)$ is the proportion of incorrectly predicted neighborhood types (Sunni, Shia, or mixed). We count a neighborhood as a correct prediction if at time t it is dominated ($\geq 70\%$) by the same group as given on the ethnic map, or alternatively, if no group reaches a 70% share in the simulation and the neighborhood is “mixed” according to the map. The error $\epsilon_v(t)$ between the observed violence map and the simulated one at time t is simply the mean squared difference across neighborhoods between the observed normalized violence score and the normalized simulated number of attacks per neighborhood that occurred until time t (remember that we would like to approximate the cumulative distribution of attacks). As described below, the error measures $\epsilon_e(t)$ and $\epsilon_v(t)$ help us select those parameter combinations of the model that come closest to the observed patterns in Baghdad.

Parameter Estimates

First, we fit our model to the empirical data described above. This exercise shows us whether the model parameter estimates we obtain are in line with our theoretical reasoning. These parameter estimates help us conduct simple counterfactual experiments with our model, described in the next section.

Following our strategy described above, we run the model on the empirical data, using 20,000 randomly drawn parameter vectors. In principle, the coefficients in the attack and migration logit models could be unbounded. Practically, however, we only need to consider parameter ranges where the coefficients in the logit models have a significant impact on the computed probability; for this reason, we restrict the sampling range for the coefficients α_1 , β_1 and β_2 to $[-10, 10]$ and for the intercepts α_0 and β_0 to $[-20, 0]$. For the experiments reported below we use a fixed proportion of insurgents π of 0.02, indicating that only a very small share of the population is disposed to violence.¹⁵ Each run is limited to 500 time steps.¹⁶ We now need to select those parameter vectors θ that maximize the fit with respect to both the distribution of groups and violence across the city. More precisely, a “plausible” parameter vector eventually (i) produces a simulated ethnic map that approximates the observed one, (ii) produces a spatial distribution of violence similar to the observed one, and (iii) does so at the same time during the model run. Note that since we do not calibrate the model with respect to time, the model can reach a good fit at any time during the simulation.

In order to select parameter configurations based on these three criteria, we record for each model run (i) the time t_{opt} when $\epsilon_e(t)$ is minimal: $t_{opt} = \arg \min_{t \in [0, 500]} \epsilon_e(t)$, (ii) the value of $\epsilon_e(t_{opt})$, and (iii) the value of $\epsilon_v(t_{opt})$. Remember that model runs start with an ethnic configuration as given by the ethnic map at the beginning of the respective

¹⁵Our tests with alternative values of this parameter (0.01, 0.03) produce similar results, see online appendix.

¹⁶We conducted initial experiments showing that most of the model dynamics occur with the first 100-200 time steps.

period, so $\epsilon_e(0)$ is simply the prediction error between this map and the one at the end of the period.¹⁷ Many parameter combinations yield $t_{opt} = 0$, in other words, the model never improves upon the ethnic map at the beginning of the period. Moreover, a small t_{opt} could indicate that the improvement in ϵ_e is due to random fluctuation at the beginning of the model run. Therefore, when selecting our plausible parameter vectors, we first discard all runs with $t_{opt} \leq 5$, and from this sample, select those where both $\epsilon_e(t_{opt})$ and $\epsilon_v(t_{opt})$ are smaller than their respective mean.

This results in a set Θ' of 533 plausible parameter vectors. This low number emphasizes the narrow margin within which the model captures the dynamics similar to Baghdad: only about 3% of all parameter vectors are likely to have generated our data. In other words, the vast majority of parameter combinations can be ruled out, leaving us with a more limited range of theoretically interesting models. The fit of the model is fair but not overwhelming. In the pre-surge period, the model reduces the prediction error ϵ_e for ethnicity from an $\epsilon_e(0)$ of 0.45 to a minimum of 0.38. The minimal error ϵ_v between the simulated and the observed violence maps is 0.05. Still, given the lack of data for many essential features in Baghdad, a perfect replication of migration and violence for each of the 85 neighborhoods is near impossible. Rather, our models are able to capture the most important dynamics of the empirical case: increasing segregation along with decreasing violence. Most importantly, we are able to *exclude* many implausible parameter combinations and can focus attention on the most theoretically interesting estimates.

Proportion of Co-ethnics and Insurgent Attacks We first consider the coefficient α_1 in the attack model, which controls the effect of the proportion of co-ethnics on the likelihood of attacks. Figure 3 plots the density of α_1 in Θ' (solid line) against the full sampling range (grey line), which we drew from a uniform distribution.

Recall that a negative estimate for α_1 would be consistent with the “maximum impact” mechanism of violence: attacks become more likely the higher the proportion of the other group in order to maximize damage to the out-group. However, the plot shows that across all plausible parameter vectors, α_1 is almost consistently positive. This leads to two conclusions: First, a violence generating mechanism that depends on the local ethnic configuration is a necessary part of the model. In other words, we can show that a random allocation of insurgent attacks, independently of the local ethnic configuration, would not account for the patterns we observe in Baghdad. Moreover, the positive coefficient provides support for an “ethnic cleansing” type of violence, where higher proportions of co-ethnics encourage insurgents to attack small minorities of the other group. This is consistent with the argument that insurgents need a local support base in order to successfully carry out attacks and that attacks are motivated by the desire to create ethnically homogenous enclaves.

Experienced Violence and Migration Next, we consider the effect of violence on migration, as captured by the coefficients β_1 and β_2 in the migration model. β_1 deter-

¹⁷The model dynamics start at time step 1.




Figure 3: Density of α_1 in Θ' (solid line). The full sampling range in Θ is shown as a grey line.

mines the effect of violence in civilian's neighborhood on her decision to migrate, whereas β_2 captures the indirect effect of violence occurring in the proximity of the civilian (the spatial lag of violence). The densities of the two coefficients are shown in Figure 4.

For β_1 , we see a result similar to the previous one: the density of the parameter is clearly positive as we would expect. We can conclude that fear created by insurgent violence has a strong impact on the settlement pattern changes we see in Baghdad during the course of our study period. The right panel shows that, counter to our expectations, there does not seem to be a signaling effect of violence across neighborhoods. The spatially lagged violence does not seem to affect an individual's migration decision in a clearly discernible way, since the density of the β_2 in Θ' is virtually indistinguishable from the sampling density and can take both negative and positive values. In sum, our models produce a better fit when attacks against one's ethnic group are positively related to migration, but only in the locality where attacks occur. In short, as we would expect, violence in a locality causes members of the affected ethnic group to move to other locations.

Violence and Segregation over Time One of our principal claims is that there is a reciprocal relationship between violence and settlement patterns. Violence should induce migration, but ethnic homogeneity should reduce levels of violence, even in the absence of policing. In order to show how our key outputs of interest—the level of violence and segregation—evolve over time, we recorded these outputs at every time step in the 533 model runs with parameters in Θ' . Averages over these runs are plotted in Figure 5.

Clearly, increased segregation is generally related to lower levels of violence, which

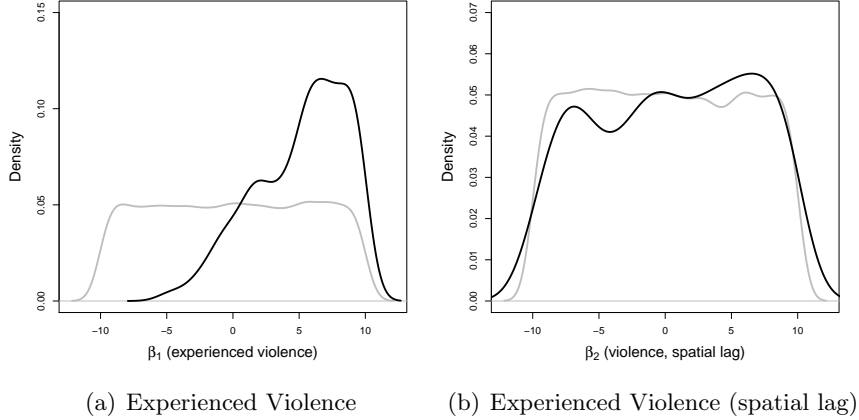


Figure 4: Density of β_1 (left panel) and β_2 (right panel) in Θ' (solid lines).

is also what we observed in Baghdad. However, what is the effect of *local* ethnic mixing on violence, i.e. how does the composition of neighborhood affect its risk of ethnic violence? Our results above suggest the following relationship: as mixed ethnic neighborhoods move towards homogeneity, attacks against the minority increase (positive effect of proportion co-ethnics). At the same time, however, the outflow of minority members reduces the number of potential targets, removing the risk of violence in homogenous neighborhoods. Consequently, there should be a curvilinear relationship between segregation and violence, with intermediate levels of segregation giving rise to many attacks, but a decline at high levels of segregation. A quick additional test confirms this. Table 1 shows the results of a regression of the (logged) number of attacks on the proportion of homogenous neighborhoods and its squared term (including dummies for each model run), which confirms the inverted U-shaped relationship.

Table 1: Notes: Regression of violence on segregation. Number of observations: 266,500. Adjusted R-squared: 0.83. Coefficients for the model run dummies not shown. Standard errors in parentheses, * indicates significance at $p < 0.05$.

	<i>Model 1</i>
Prop. homogeneous	3.73*
	(0.03)
Prop. homogeneous (squared)	-3.21*
	(0.03)
(Intercept)	0.45*
	(0.02)

What can we conclude from these results? First, we are able to show that violent attacks are consistent with an ethnic-cleansing mechanism: attacks are more likely to

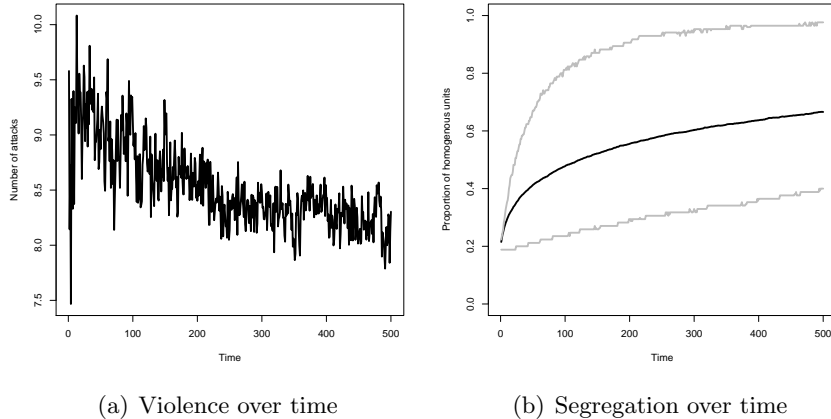


Figure 5: Evolution of violence (left) and segregation (right) over time, averaged over 533 models runs with parameters from Θ' .

occur in areas where there are small but significant minorities. Second, we show that violence increases the level of ethnic segregation over time. Significantly, we do not need to assume ethnophilia to generate this result. Ethnic segregation need not be driven by antipathy toward out-groups; the search for safety is all that is required to produce unmixing. Finally, we show that ethnic segregation—by itself—can reduce the level of violence in a system, even if it is not normatively desirable. Under conditions of anarchy, violence declines as neighborhoods become more homogenous—a finding in support of arguments that the troop surge was irrelevant for the drop in observed violence in Baghdad (Agnew et al. 2008). Thus, there is a reciprocal relationship between settlement patterns and violence: initial levels of ethnic mixing influence where attacks are likely to occur, these attacks produce migration and segregation, in turn, segregation reduces the level of violence. In the next section, we turn to policing and its ability to limit insurgent attacks.

Policing to Reduce Ethnic Violence

In this section, we provide a more nuanced perspective on alternative solutions to ethnic violence, in particular those that involve policing by outside actors (e.g. the state). Based on our above results, we cannot dismiss the effect of policing altogether; all we can say is that segregation is one way to reduce violence, and *could have been* sufficient to bring down violence in Baghdad. To what extent can policing be successful in limiting violence? Can enhanced policing efforts be effective relatively late in an insurgency, or is it better to have a robust force in place early on, as argued by General Shinseki? We address these questions in two steps. First, we analyze an “early on” implementation of policing, with policing starting at time step 1. In these simulations, we systematically

vary the level of policing success to assess its effect on the reduction of violence and ethnic segregation. Second, we perform a set of simulation runs to examine more closely the effect of timing, i.e. the time when counterinsurgent efforts start. This helps us to disentangle whether the *size* of the policing force, or the *timing* of these efforts matter more.

Implementing Policing As argued above, an alternative way of violence reduction is by means of policing. By adding a simple policing mechanism to our model, we can “re-run” history and assess different counterfactual worlds. Policing is modeled using a “success rate” of counterinsurgency, which accounts for the probability that an insurgent is removed from the scene by capture or killing, after having carried out an attack. Note that in this implementation, agents can only be captured once they have attacked successfully and thus have visibly identified themselves as insurgents. In the present version of the model, we remain agnostic about possible variation in this success rate across locations, and assume a constant probability of capture. After an insurgent has carried out an attack, we take a random draw from a Bernoulli distribution with success rate τ . If successful, the insurgent will be removed from the model. In addition, we can “switch on” policing at a particular time during the simulation. Furthermore, we assume that punished attacks do not generate fear in the targeted population in the same way as regular attacks do. Therefore, only the number of unpunished attacks enters the agents’ migration decision.¹⁸

We apply our counterfactual experiments—the variations in policing success and timing—to the parameter estimates obtained above (i.e. the parameter vectors in Θ' , with a 0.02 proportion of insurgents). We initialize the model using the ethnic map of 2003, which reflects the distribution of groups in Baghdad roughly at the beginning of the invasion. This map shows that Baghdad is much less segregated than in 2006, with only about 30% of the neighborhoods coded as homogenous. Based on this setup, we systematically vary the success rate of policing (experiment 1) and the onset of policing (experiment 2). For each value of these experimental variations, we run the model on all 533 parameter vectors in Θ' , again limiting each run to 500 time steps. We collect information about the level of violence (measured as the total number of attacks in a run) and segregation (measured as the proportion of homogenous units). The results shown below are averages over the 100 parameters combinations tested for each experimental variation.

Our first counterfactual experiment tests the effect of increased policing success on violence and segregation, and Figure 6 shows the results. We vary the policing success rate τ from 0 to 1 in increments of 0.1; policing is implemented from the beginning of the simulation. The left panel shows that even small levels of policing are highly effective in reducing violence: even a low success rate of 0.1 reduces violence by more than 75% compared to no policing. The effects of policing on segregation are not as strong. There is a continuous decline in segregation as policing increases, but with limited success. Even implausibly high policing success rates of 0.5 and above cannot prevent a segregation

¹⁸Our results do not depend crucially on this assumption, see online appendix.

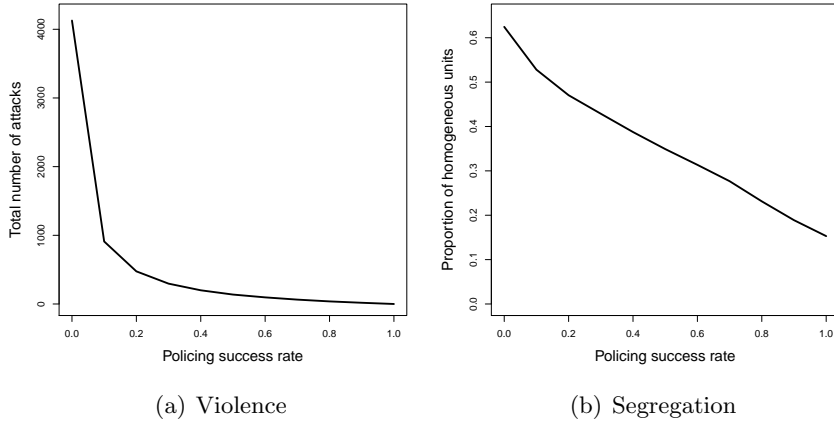


Figure 6: Effect of policing on the reduction of violence (left) and segregation (right).

up to a certain limit (recall that the initial proportion of homogenous units is 0.3). The reason is that for policing to be effective and insurgents to be captured, these insurgents have to attack first. This unavoidable level of violence may be sufficient to trigger the ethnic unmixing we see in this experiment. Nonetheless, we can conclude from this experiment that policing can have a large impact on reducing violence levels, even at relatively low levels of counterinsurgent effectiveness.

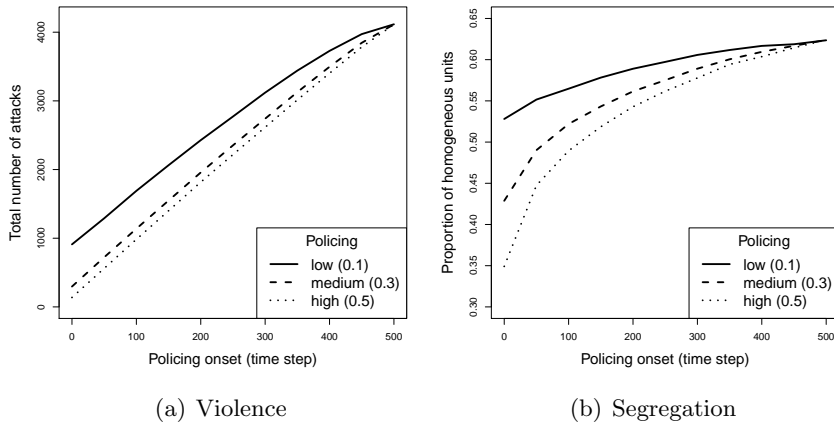


Figure 7: Effect of the onset of policing on the reduction of violence (left) and segregation (right).

In our first experiment described above, the policing mechanism is active from the first time step. In the next experiment, we examine the timing of policing, varying the

onset of policing from time step 0 to 500 in steps of 50. We do this for two different policing success rates: low ($\tau = 0.1$), medium ($\tau = 0.3$) and high ($\tau = 0.5$). Figure 7 shows the results for low (solid lines), medium (dashed) and high (dotted lines) levels of success. As we expected, the later the policing efforts start, the smaller the reduction in violence and segregation. Again, similar to the previous experiment, the effect on segregation is limited due to the migration occurring even with low levels of violence. Both plots, however, draw our attention to the importance of the timing, especially the first 100 time steps. In order to compensate for a delay of 100 time steps in the onset of our policing efforts, we would have to roughly triple policing efforts to achieve a comparable reduction in violence: we achieve the same reduction in violence if we implement policing with a low (0.1) success rate right from the beginning, or with a medium (0.3) success rate after 100 time steps. The effect on segregation reduction is similar. Together with our results from the previous experiments, which showed that even small success rates can mitigate violence considerably, these results emphasize even more the importance of a timely response. Thus, these results may be seen as a vindication of the Shinseki plan: early implementation of effective counterinsurgency can have a very large effect, while late implementation is much less successful.

Conclusion

In this paper, we examined the relationship between violence and ethnic segregation using an agent based model informed by actual data from Baghdad. While we are not able to re-run history and draw definitive conclusions about what actually happened in Iraq, our results do contribute to our understanding of ethnic conflict more generally. First, we show that ethnic settlement patterns influence where violent attacks are likely to occur. Violence is most likely when there are small but significant minorities in a locality as insurgents attempt to ethnically cleanse neighborhoods. Second, the desire for safety is enough to generate a considerable degree of ethnic segregation. Although people may desire to live in neighborhoods dominated by their co-ethnics, this is not necessary to produce ethnic unmixing in the context of escalating violence. All that is needed is for individuals to move to areas that are relatively safer. Third, we show that as neighborhoods become more segregated, the level of violence eventually declines. While partition into mutually exclusive ethnic enclaves may not be normatively desirable, it can help to reduce bloodshed in the short run. Thus, there is a reciprocal relationship between ethnic segregation patterns and violence. While we cannot provide the final answer as to what happened in Baghdad, our results indicate that one plausible explanation for why violence declined as it did was due to the fact that by the time of the troop surge, ethnic segregation was nearly complete.

Our modeling exercise also informs current debates about the effectiveness of counterinsurgency operations. We show that even small increases in policing effectiveness can dramatically mitigate the level of violence, but only if policing is implemented very early on. Delayed implementation of a robust counterinsurgency strategy may help to reduce violence somewhat, but it is not nearly as effective as efforts early on. This finding has

important implications for policy discussions. While increasing the number of troops in Iraq and Afghanistan may have some utility in the context of a raging insurgency, resources would have been better used if sufficient forces were in place from the start. In the future, policy makers must exercise extreme caution before engaging in military occupations, and do so only if they are willing to devote overwhelming resources to contain an insurgency.

Finally, in this paper we present an approach to agent-based modeling that is grounded in empirical data. We use actual data from Baghdad to initialize our model, and we attempt to select those parameters that best fit the empirical record. While abstract modeling exercises can be extremely useful in developing theories and contributing to scholarly debate, validating these models with empirical data can help to demonstrate the utility of a particular tool. No computational model can credibly claim to perfectly recreate history, but such tools can significantly aid in our understanding of complex social phenomena.

References

- AGNEW, JOHN, THOMAS W. GILLESPIE, JORGE GONZALEZ AND BRIAN MIN. (2008) Baghdad Nights: Evaluating the Us Military 'Surge' Using Nighttime Light Signatures. *Environment and Planning A* 40: 2285–2295.
- AXELROD, ROBERT. (1984) *The Evolution of Cooperation*. New York: Basic Books.
- BENNETT, D. SCOTT. (2008) Governments, Civilians, and the Evolution of Insurgency: Modeling the Early Dynamics of Insurgencies. *Journal of Artificial Societies and Social Simulation* 11(4).
- BERMAN, ELI, JACOB N. SHAPIRO AND JOSEPH H. FELTER. (2011) Can Hearts and Minds Be Bought? The Economics of Counterinsurgency in Iraq. *Journal of Political Economy* 119(4): 766–819.
- BHAVNANI, RAVI. (2006) Ethnic Norms and Interethnic Violence: Accounting for Mass Participation in the Rwandan Genocide. *Journal of Peace Research* 43(6): 651–669.
- CEDERMAN, LARS-ERIK. (1997) *Emergent Actors in World Politics: How States and Nations Develop and Dissolve*. Princeton, NJ: Princeton University Press.
- CEDERMAN, LARS-ERIK. (2001) Agent-Based Modeling in Political Science. *The Political Methodologist* 10: 16–22.
- CEDERMAN, LARS-ERIK. (2003) Modeling the Size of Wars: From Billiard Balls to Sandpiles. *American Political Science Review* 97(1): 135–150.
- CEDERMAN, LARS-ERIK. (2008) Articulating the Geo-Cultural Logic of Nationalist Insurgency. In *Order, Conflict, and Violence*, ed. Stathis Kalyvas, Ian Shapiro and Tarek Masoud. Cambridge: Cambridge University Press.
- CHAPMAN, THOMAS AND PHILIP G. ROEDER. (2007) Partition as a Solution to Wars of Nationalism: The Importance of Institutions. *American Political Science Review* 101(4): 677–691.
- CIOFFI-REVILLA, CLAUDIO AND MARK ROULEAU. (2010) MASON RebeLand: An Agent-based Model of Politics, Environment, and Insurgency. *International Studies Review* 12(1): 31–52.
- COLEMAN, JAMES S. (1990) *Foundations of Social Theory*. Cambridge, MA: The Belknap Press of Harvard University Press.
- COLLIER, PAUL AND ANKE HOFFLER. (2004) Greed and Grievance in Civil War. *Oxford Economic Papers* 56(4): 563–595.
- DEAN, JEFFREY S., GEORGE J. GUMERMAN, JOSHUA M. EPSTEIN, ROBERT AXTELL, ALAN C. SWEDLUND, MILES T. PARKER AND STEPHEN MCCARROLL. (2006)

- Understanding Anasazi Culture Change through Agent-based Modeling. In *Generative Social Science: Studies in Agent-based Computational Modeling*, ed. Joshua M. Epstein. Princeton: Princeton University Press.
- DE MARCHI, SCOTT. 2005. *Computational and Mathematical Modeling in the Social Sciences*. New York: Cambridge University Press.
- ENTERLINE, ANDREW J. AND MICHAEL GREIG. (2007) Surge, Escalate, Withdraw, and Shinseki: Forecasting and Retro-Casting American Force Strategies and Insurgency in Iraq. *International Studies Perspectives* 8(3): 245–252.
- EPSTEIN, JOSHUA M. (2002) Modeling Civil Violence: An Agent-based Computational Approach. *Proceedings of the National Academy of Sciences USA* 99(3): 7243–7250.
- EPSTEIN, JOSHUA M. AND ROBERT AXTELL. (1996) *Growing Artificial Societies*. MIT Press.
- FEARON, JAMES D. AND DAVID D. LAITIN. 2003. Ethnicity, Insurgency and Civil War. *American Political Science Review* 97(1): 75–90.
- GELLER, ARMANDO, SEYED M. M. RIZI AND MACIEJ M. LATEK. (2011) How Corruption Blunts Counternarcotic Policies in Afghanistan: A Multiagent Investigation. In *Social Computing, Behavioral-Cultural Modeling and Prediction, 2011 Conference. Lecture Notes in Computer Science 6489*, ed. J. Salerno, S.J. Yang, D. Nau and S.K. Chai. Springer pp. 121–128.
- GELLER, ARMANDO AND SHAH JAMAL ALAM. (2010) A Socio-Political and -Cultural Model of the War in Afghanistan. *International Studies Review* 12(1): 8–30.
- HAMMOND, ROSS AND ROBERT AXELROD. (2006) The Evolution of Ethnocentrism. *Journal of Conflict Resolution* 50(6): 926–936.
- HENDRIX, CULLEN. (2010) Measuring State Capacity: Theoretical and Empirical Implications for the Study of Civil Conflict. *Journal of Peace Research* 47(3): 273–285.
- HERBST, JEFFREY. (2000) *States and Power in Africa: Comparative Lessons in Authority and Control*. Princeton, NJ: Princeton University Press.
- HERBST, JEFFREY. (2004) African Militaries and Rebellion: The Political Economy of Threat and Combat Effectiveness. *Journal of Peace Research* 41(3): 357–369.
- JOSEPH, EDWARD AND MICHAEL E. O’HANLON. (2007) The Case for Soft Partition of Iraq. *Saban Center for Middle East Policy Analysis Papers* 12(June 2007).
- KALYVAS, STATHIS N. (2006) *The Logic of Violence in Civil War*. New York: Cambridge University Press.
- KAUFMANN, CHAIM. (1996) Possible and Impossible Solutions to Ethnic Civil Wars. *International Security* 20(4): 136–175.

- LAKE, DAVID. (2002) Rational Extremism: Understanding Terrorism in the Twenty-First Century. *Dialogue-IO* 1(15): 29.
- LUSTICK, IAN S., DAN MIODOWNIK AND ROY J. EIDELSON. (2004) Secessionism in Multicultural States: Does Sharing Power Prevent or Encourage it? *American Political Science Review* 98(2): 209–229.
- LYALL, JASON AND ISAIAH WILSON. (2009) Rage Against the Machines: Explaining Outcomes in Counterinsurgency Wars. *International Organization* 63(1): 67–106.
- MATTHEWS, DYLAN AND EZRA KLEIN. (2008) How Important Was the Surge? *The American Prospect*. July 28, 2008. Online edition.
- MUELLER, JOHN. (2000) The Banality of "Ethnic War". *International Security* 25(1): 42–70.
- MULLER, EDWARD AND MITCHELL SELIGSON. (1987) Inequality and Insurgency. *American Political Science Review* 81(2): 425–451.
- OLZAK, SUSAN, SUSANNE SHANAHAN AND ELIZABETH MCEANEY. (1996) Poverty, Segregation, and Race Riots: 1960-1993. *American Sociological Review* 61(4): 590–613.
- POSEN, BARRY R. (1993) The Security Dilemma and Ethnic Conflict. In *Ethnic Conflict and International Security*, ed. Michael E. Brown. Princeton: Princeton University Press.
- RICKS, THOMAS. (2009) *The Gamble: General David Petraeus and the American Military Adventure in Iraq, 2006-2008*. New York: Penguin.
- ROSE, WILLIAM. (2000) The Security Dilemma and Ethnic Conflict: Some New Hypotheses. *Security Studies* 9(4).
- SAMBANIS, NICHOLAS. (2000) Partition as a Solution to Ethnic War. *World Politics* 52: 437–483.
- SCHELLING, THOMAS C. (1971) Dynamic Models of Segregation. *Journal of Mathematical Sociology* 1: 143–186.
- SIEGEL, DAVID A. (2009) Social Networks and Collective Action. *American Journal of Political Science* 53(1): 122–138.
- TILLY, CHARLES. (1978) *From Mobilization to Revolution*. Reading: Addison-Wesley.
- URDAL, HENRIK. (2006) A Clash of Generations? Youth Bulges and Political Violence. *International Studies Quarterly* 50(3): 607–629.
- WILKINSON, STEVEN. (2004) *Votes and Violence*. Cambridge University Press.