Genocide Studies and Prevention: An International Journal

Volume 18
Issue 1 *Evidence-Based Approaches to Preventing Genocide and Mass Atrocities*

Article 8

10-2-2024

Leveraging a Multi-Method Approach to Improve Mass Atrocity Forecasting

Hollie Nyseth Nzitatira
The Ohio State University

Trey Billing
The Ohio State University

Eric W. Schoon
The Ohio State University

Leveraging a Multi-Method Approach to Improve Mass Atrocity Forecasting

Acknowledgements

The authors wish to thank the editorial team and the other contributors to this special issue. This research was sponsored by the Political Instability Task Force (PITF), which is funded by the US Central Intelligence Agency. The views expressed herein are the authors' alone and do not represent the views of the US Government.

Leveraging a Multi-Method Approach to Improve Mass Atrocity Forecasting

Hollie Nyseth Nzitatira
Department of Sociology, The Ohio State University
Columbus, Ohio, U.S.A.

Trey Billing
Mershon Center for International Security Studies, The Ohio State University
Columbus, Ohio, U.S.A.

Eric W. Schoon
Department of Sociology, The Ohio State University
Columbus, Ohio, U.S.A.

Introduction

During the last three decades, researchers, policymakers, and practitioners have set their sights on forecasting mass atrocities. Guided by the notion that accurately predicting atrocities can aid in prevention efforts, these endeavors have sought to identify how to best forecast the onset of a mass atrocity like a genocide or a targeted mass killing. Government task forces, non-profit organizations, and academics have devoted numerous resources to such projects, and the United Nations has likewise developed their own method for predicting whether and when an atrocity might occur.

Broadly, such atrocity forecasting efforts can be categorized into two methodological camps. On the one hand, some efforts emphasize quantitative analyses that enable the assessment of numerous countries across time periods. These quantitative tools privilege the breadth and predictive power that come with large datasets, as well as the impact of isolated factors that increase the risk of mass atrocity. On the other hand, other forecasting efforts have employed case-based analyses that take concurrent risk factors into account when assessing the likelihood of mass atrocities in certain locations. These approaches typically privilege the identification of how certain factors intersect to produce situations in which the occurrence of a mass atrocity is possible.

In this article, we suggest that atrocity forecasts should capitalize on the strengths of each of these general methodological approaches.[1] Specifically, we propose a new forecasting approach that integrates case-oriented and quantitative toolkits to provide a more robust, holistic assessment of the risk of mass atrocities. Notably, we also provide a methodological innovation by using the case-oriented approach to first analyze the factors associated with the *absence* of mass atrocities prior to predicting their onset. To illustrate, we analyze mass atrocities between 1972 and 2022. We conceptualize mass atrocities as mass killings by government and/or non-governmental actors in which at least 500 people from a specific group are killed within a 12-month period—a conceptualization we further detail in our methods section.

In what follows, we begin by surveying a brief history of forecasting efforts, including their inception as well as some of the major efforts that exist today. After addressing the methodological differences across these efforts, we outline our methods, which involve configurational analysis of the *absence* of mass atrocity as well as two more quantitative models

---

[1] We are not the first to make this argument. See, for instance, Ernesto Verdeja, "Predicting Genocide and Mass Atrocities," *Genocide Studies and Prevention: An International Journal* 9, no. 3 (2016), 13–32, accessed July 26, 2024, http://dx.doi.org/10.5038/1911-9933.9.3.1314.

Hollie Nyseth Nzitatira, Trey Billing, and Eric W. Schoon. "Leveraging a Multi-Method Approach to Improve Mass Atrocity Forecasting." In "Evidence-Based Approaches to Preventing Genocide and Mass Atrocities." Edited by Jamie D. Wise and Kristina Hook. Special Issue, *Genocide Studies and Prevention* 18, no. 1, 54–83. https://doi.org/10.5038/1911-9933.18.1.1951.

that identify the countries that are most at risk of mass atrocity. We conclude by assessing our forecasts and addressing the utility of our approach.

Forecasting Mass Atrocities: A Brief Overview

As mass atrocities unfolded in Rwanda and Bosnia-Herzegovina in the early 1990s, people around the world questioned whether and how such widespread harm could have been prevented. Soon, government officials across continents began calling for efforts to predict mass atrocities, guided by the sound assumption that a crucial step in preventing atrocities is predicting their onset. Since then, a robust body of research, as well as associated efforts seeking to forecast mass atrocities, have emerged.

Apart from methodological divergences, forecasting efforts generally fall into two groups: (1) analyses of risk factors, or structural factors and related situations that influence whether mass atrocities are likely to occur; and (2) analyses of escalatory factors and triggers, or specific events and processes that more directly influence the onset or escalation of violence.[2] In line with the distinction between risk factors and triggering/escalatory factors, models of the onset of mass atrocities can be classified into two categories: risk assessment models and early warning models. Risk assessment models typically address a country's structural conditions that affect the possibility mass atrocities could occur, while early warning models focus on more proximate dynamics that escalate or trigger violence, as well as indicators that violence is imminent.[3]

In this article, we focus on the risk factors of mass atrocity and, as such, on risk assessment models. We do so because risk assessment models should logically be undertaken prior to early warning efforts. Put simply, risk assessment models can narrow the countries and communities that are at high risk of mass atrocities, while early warning models can aid in identifying whether such risk may materialize, as well as when. Nonetheless, it is also important to note that many risk assessment models (including ours) incorporate some escalatory and triggering factors as well, indicating that the boundaries between these categorizations are somewhat blurry.

That said, the earliest risk assessment models were quantitative, large-N[4] assessments that emerged in the late 1980s and early 1990s. These models drew from case-based research that sought to understand why mass atrocities—and, specifically, genocide—had occurred.[5] As such, it is vital to underscore that case studies of genocide and other mass atrocities laid the foundation for knowledge regarding the risk factors of future violence.[6] Collectively, this early quantitative work found that prior genocide, autocracies, exclusionary ideologies, political

---

[2] Scott Straus, *Fundamentals of Genocide and Mass Atrocity Prevention* (Washington, DC: United States Holocaust Memorial Museum, 2016). Risk factors are also termed priming factors. See also Alexander Laban Hinton, *Why Did They Kill?: Cambodia in the Shadow of Genocide* (Berkeley: University of California Press, 2005).

[3] Birger Heldt, "Risks, Early Warning and Management of Atrocities and Genocide: Lessons from Statistical Research," *Politorbis* 45 (2009), 65–70, accessed July 25, 2024, https://ssrn.com/abstract=2025007; Barbara Harff, "Detection: The History and Politics of Early Warning," in *Responding to Genocide: The Politics of International Action*, ed. Adam Lupel and Ernesto Verdeja (Boulder: Lynne Rienner, 2013), 85–110; Verdeja, *Predicting Genocide and Mass Atrocities*.

[4] "N" refers to the sample size, and a large-N study typically involves quantitative assessment of numerous countries across years.

[5] Helen Fein, "Accounting for Genocide after 1945: Theories and Some Findings," *International Journal on Group Rights* 1, no. 2 (1993), 79–106, accessed July 25, 2024, https://www.jstor.org/stable/24674446; Leo Kuper, *Genocide: Its Political Use in the Twentieth Century* (New Haven: Yale University Press, 1983). Specifically, Fein's analysis was a bivariate analysis, meaning she examined a series of relationships between sets of two variables.

[6] Nevertheless, these studies did not seek to forecast the future but rather endeavored to ascertain the factors associated with the onset of genocide and politicide.

upheaval, contention regarding the ethnicity of those in power, and low trade openness are tied to the occurrence of genocide and politicide.[7]

The subsequent decades saw a vast expansion of quantitative and case-based research regarding the risk factors of genocide and politicide, as well as related forms of violence such as mass killing.[8] While space does not allow for a full review of the risk factors identified in this work, they pertain broadly to political upheaval and threat (e.g., assassinations), state structure and capacity (e.g., infant mortality rate), ideology and social divisions (e.g., state-led discrimination), conflict and human rights history (e.g., recent history of atrocities), and international factors (e.g., membership in certain international organizations)—risk factors that we provide further detail about in our methods section.[9] More recent studies have also examined the predictive power of risk assessment models given continued methodological advancements. For instance, as quantitative studies can have statistically significant results that fare poorly with respect to forecasting the future,[10] newer studies employ different tools to assess the predictive capabilities of models.[11]

Atrocity Risk Assessment Efforts

Today, numerous initiatives draw upon the work we have reviewed thus far to develop annual atrocity risk assessment models. While we cannot consider all risk assessment endeavors, we briefly outline some of the most prominent and public efforts to provide a broad overview. In doing so, we underscore major methodological divergences with respect to quantitative and case-based approaches.

Several key forecasting efforts rely upon quantitative analysis for risk assessment. Specifically, one of the earliest atrocity forecasting bodies was the State Failure Task Force, a United States government forecasting project that the government created in the wake of the

[7] Matthew Krain, "State-Sponsored Mass Murder: The Onset and Severity of Genocides and Politicides," *Journal of Conflict Resolution* 41, no. 3 (1997), 331–360, accessed July 25, 2024, https://doi.org/10.1177/0022002797041003001; Barbara Harff, "No Lessons Learned from the Holocaust? Assessing Risks of Genocide and Political Mass Murder since 1955," *American Political Science Review* 97, no. 1 (2003), 57–73, accessed July 25, 2024, https://doi.org/10.1017/S0003055403000522; Barbara Harff and Ted Robert Gurr, "Systematic Early Warning of Humanitarian Emergencies," *Journal of Peace Research* 35, no. 5 (1998), 551–579. Barbara Harff and Ted Gurr continued to forecast genocide and politicide until 2018.

[8] Hollie Nyseth Brehm, "Re-examining Risk Factors of Genocide," *Journal of Genocide Research* 19, no. 1 (2017), 61–87, accessed July 25, 2024, https://doi.org/10.1080/14623528.2016.1213485; Benjamin Valentino et al., "'Draining the Sea': Mass Killing and Guerrilla Warfare," *International Organization* 58, no. 2 (2004), 375–407, accessed July 26, 2024, https://www.jstor.org/stable/3877862; Daniel Krcmaric, "Varieties of Civil War and Mass Killing: Reassessing the Relationship between Guerrilla Warfare and Civilian Victimization," *Journal of Peace Research* 55, no. 1 (2018), 18–31, accessed July 25, 2024, https://doi.org/10.1177/0022343317715060; Gary Uzonyi, "Civil War Victory and the Onset of Genocide and Politicide," *International Interactions* 41, no. 2 (2015), 365–391, accessed July 26, 2024, https://doi.org/10.1080/03050629.2014.932786; Nicholas Rost, "Will It Happen Again? On the Possibility of Forecasting the Risk of Genocide," *Journal of Genocide Research* 15, no. 1 (2013), 41–67, accessed July 25, 2024, https://doi.org/10.1080/14623528.2012.759398; Gary Uzonyi, "Domestic Unrest, Genocide and Politicide," *Political Studies* 64, no. 2 (2016), 315–334, accessed July 26, 2024, https://doi.org/10.1111/1467-9248.12181; Erica Chenoweth et al., "State Repression and Nonviolent Resistance," *Journal of Conflict Resolution* 61, no. 9 (2017), 1950–1969, accessed July 25, 2024, https://doi.org/10.1177/0022002717721390; Michael Colaresi and Sabine C. Carey, "To Kill or to Protect: Security Forces, Domestic Institutions, and Genocide," *Journal of Conflict Resolution* 52, no. 1 (2008), 39–67, accessed July 25, 2024, https://doi.org/10.1177/0022002707310427; Charles H. Anderton and John R. Carter, "A New Look at Weak State Conditions and Genocide Risk," *Peace Economics, Peace Science and Public Policy* 21, no. 1 (2015),1–36, accessed July 25, 2024, https://doi.org/10.1515/peps-2014-0008; Charles H. Anderton and Jurgen Brauer, "Mass Atrocities and Their Prevention," *Journal of Economic Literature* 58, no. 4 (2021), 1240–1292.

[9] Hollie Nyseth Nzitatira, "Predicting Genocide," in *Genocide: Key Themes*, ed. Donald Bloxham and A. Dirk Moses (New York: Oxford University Press, 2022), 45–74.

[10] Michael D. Ward et al., "The Perils of Policy by P-Value: Predicting Civil Conflicts," *Journal of Peace Research* 47, no. 4 (July 2010), 363–375, accessed July 26, 2024, https://www.jstor.org/stable/20752194.

[11] For instance, forecasters examine the Receiver Operator Characteristic (ROC) curve—which plots the relationship between the rate of false positives and the rate of true positives—to better assess predictive capacities. Many also rely upon out-of-sample forecasts.

1994 genocide in Rwanda.[12] This task force, which later became known as the Political Instability Task Force (PITF), used quantitative analysis of approximately 20 risk factors to develop annual forecasts for mass killings.[13] Also in the United States but stemming from nongovernmental bodies, the Early Warning Project forecasts mass killings committed by state or nonstate actors within a state's borders. Like the PITF, the Early Warning Project engages in quantitative forecasting via a model of roughly 20 risk factors.[14] Australia's Atrocity Forecasting Project likewise publishes annual risk assessments of the top 15 countries at risk of mass atrocity—though focuses specifically on genocide or politicide—based on quantitative models of 19 variables.[15]

While the PITF, the Early Warning Project, and the Atrocity Forecasting Project each rely upon quantitative forecasting methods, other major atrocity forecasting endeavors take a case-based approach. For instance, Genocide Watch (a non-profit organization based in the United States) uses a ten-stage model to predict genocide, politicide, and what they deem genocide-like crimes.[16] The United Nations Office of the Special Adviser on the Prevention of Genocide likewise engages in a case-based risk assessment. Specifically, this office draws upon their published *Framework of Analysis for Atrocity Crimes* to identify countries at risk of genocide (alongside other atrocity crimes) based on a series of 14 risk factors.[17] As one of several other examples of case-based approaches, Minority Rights Group International compiles an annual risk assessment called Peoples Under Threat.[18] This assessment aims to identify the risk of genocide, mass killing, or other systematic and violent repression, and it assesses the co-occurrence of 10 risk factors[19] (see also the U.S. State Department Atrocity Risk Assessment Framework).

To be certain, Western countries are heavily overrepresented in risk assessment efforts, indicating that there is a great need for capacity building for forecasting efforts in the Global South, as we certainly would have included more efforts from around the world if such efforts were public.[20] Yet, we also do not aim to review all forecasting projects in existence but rather have highlighted prominent risk assessment models to underscore the methodological diversity involved in predicting mass atrocity. Indeed, the PITF, the Early Warning Project, and the Australia Forecasting Project take a more quantitative approach, while Genocide Watch, the United Nations Office of the Special Adviser on the Prevention of Genocide, and Minority Rights Group International take a case-based approach.

---

[12] Much of the early forecasting work undertaken by this body drew upon the pioneering work of Barbara Harff.

[13] These models and their results are not publicly available.

[14] The Early Warning project was launched as a joint initiative of the U.S. Holocaust Memorial Museum and Dartmouth College in 2015: https://earlywarningproject.ushmm.org/methodology-statistical-model. It defines mass killing as the death of 1,000 non-combatants from a discrete group who were harmed in a period of sustained violence.

[15] See all variables and the current forecasts here: https://politicsir.cass.anu.edu.au/sites/default/files/AFP_2021-23_BriefReport.pdf; see also Benjamin E. Goldsmith et al., "Forecasting the Onset of Genocide and Politicide: Annual Out-of-Sample Forecasts on a Global Dataset, 1988–2003," *Journal of Peace Research* 50, no. 4 (2013), 437–452, accessed July 25, 2024, https://doi.org/10.1177/0022343313484167.

[16] Gregory H. Stanton, "The Ten Stages of Genocide," *Genocide Watch*, 2023, accessed July 25, 2024, https://www.genocidewatch.com/tenstages. The alerts identify what the organization calls genocide watches, genocide warnings, and genocide emergencies. Note also that the stages are not linear.

[17] United Nations, *Framework of Analysis for Atrocity Crimes: A Tool for Prevention*, 2014, accessed July 26, 2024, https://www.un.org/en/genocideprevention/documents/about-us/Doc.3_Framework%20of%20Analysis%20for%20Atrocity%20Crimes_EN.pdf. The United Nations does not make public lists of countries at risk available, but the Special Adviser on the Prevention of Genocide publishes public warning statements regarding countries at risk.

[18] See the website of Minority Rights Group and type in the search term "Peoples Under Threat"; the search results will yield the relevant reports.

[19] While these indicators are measured quantitatively, Minority Rights Group International considers their co-occurrence by creating a scale to incorporate each of the 10 risk factors, hence aligning with case-oriented methods that examine how combinations of factors coalesce to impact the onset of atrocity.

[20] Verdeja, *Predicting Genocide and Mass Atrocities*. See also Deborah Mayersen, "Predicting Genocide and Mass Killing," *Journal of Genocide Research* 23, no. 1 (2021), 81–104, accessed July 25, 2024, https://doi.org/10.1080/14623528.2020.1818478.

While there is overlap such that the quantitative/case-based binary we have drawn can be blurry, there are also important differences in these methods. Quantitative approaches have greater capacity for data such that the models regularly integrate 20 or more risk factors. This is an important strength as the models can accommodate for many risk factors that increase the predictive power of the forecasts. Put another way, in most instances, a model with 15 relevant risk factors fares better with respect to prediction than a model with four relevant risk factors, which is precisely why these models typically include between 15 and 20 variables. Additionally, quantitative approaches allow for numerical assessments of risk. While these models thus enable an assessment of high or low risk, they can also provide a specific numeric value, which aids in interpreting the likelihood of mass atrocities.[21]

Nevertheless, large-N quantitative analyses have also been criticized for assumptions built into modeling strategies.[22] Most notably, quantitative methodologies typically isolate the influence of each risk factor—such as an authoritarian government—net of other factors in the model. As such, the relative influence of each risk factor is considered separately. Moreover, such methods typically operate on the assumption of symmetric relationships; that is, the notion that if a certain factor increases the risk of mass atrocity, the absence of that factor decreases such risk.[23]

Case-oriented approaches to forecasting wield different strengths and weaknesses. Rather than examining the impact of risk factors net of other factors, case-based approaches typically enable assessment of whether and how risk factors intersect to impact the occurrence of an outcome. Indeed, risk factors of mass atrocity do not occur in isolation but rather coalesce in multifaceted social situations that influence violence, which is precisely why so many researchers have employed case-based analyses of mass atrocity. Case-oriented approaches also allow for the possibility of asymmetric relationships. In this sense, the presence or absence of a mass atrocity can be treated as causally distinct, which is why, as we explain shortly, these methods enable assessment of the presence *or* the absence of mass atrocity. Nonetheless, case-oriented approaches have limits with respect to the number of risk factors that can be examined, as a rigorous case study of 20 risk factors across all countries would be exceedingly difficult, and it is likewise difficult to quantify risk via these approaches.

As such, we suggest that integrating case-based and quantitative methods to atrocity forecasting will capitalize on the strengths—and better protect against the weaknesses—of these two broad approaches. In fact, as Ernesto Verdeja has illustrated, most early warning models already incorporate both approaches even though risk assessment efforts have tended to rely upon one or the other.[24] Here, we integrate the strengths of quantitative and case-based approaches to create risk assessment models. As such, we align ourselves with much social scientific work that underscores the value of triangulation, or the application of several research methods to study the same phenomena, in pursuit of more robust findings.[25]

Methods

To integrate quantitative and case-based risk assessment models, we first constructed a database of country-years. Specifically, we created a database of 174 countries ranging from 1972 to 2022.[26] Unfortunately, data on most key measures are not available for 2023, which is a key

---

[21] See Mayersen, *Predicting Genocide and Mass Killing*, for more on the importance of including relative risk in forecasts.

[22] Verdeja, *Predicting Genocide and Mass Atrocities*; Ward et al., *The Perils of Policy by P-Value*; Mayersen, *Predicting Genocide and Mass Killing*.

[23] Implicitly, researchers tend to interpret probabilistic models undertaken as part of quantitative analyses as explaining the presence of the outcome, but they are technically explaining the difference between presence and absence.

[24] Verdeja, *Predicting Genocide and Mass Atrocities*.

[25] Paulette M. Rothbauer, "Triangulation," in *The Sage Encyclopedia of Qualitative Research Methods*, ed. Lisa M. Given (Thousand Oaks: SAGE Publications, 2008), 892–894.

[26] We initially constructed a database beginning in 1955. However, some countries are missing much data on key measures in the 1950s and 1960s. Additionally, our analysis yielded that models were *more* predictive when using a dataset spanning from 1972 to 2022 than one spanning from 1955 to 2022.

drawback that we address in our discussion. Countries that are excluded from our analysis are mostly small island nation-states with high amounts of missing data as detailed in the footnote.[27] We also employ listwise deletion for missing values and, as a result, there are 168 countries included in our configurational analyses and 164 in the other analyses, as we further explain below. In what follows, we first address how we operationalize mass atrocities alongside the risk factors included in our analysis, followed by an overview of the three main methods we employ.

Dependent Variable: Group-Based Mass Atrocities

We conceptualize mass atrocities as violence in which 500 people from a specific group are killed by state or non-state actors within a 12-month period. This definition includes all genocides, and it also includes mass killings as typically defined with a threshold of 1,000 deaths. However, the definition follows newer work that lowers the threshold of mass killings to 500 given the relative rarity of mass killings and the sheer importance of predicting as many instances of mass atrocity as possible.[28]

While we would ideally include data on injuries, sexualized violence, and the many other forms of violence that unfold during a mass atrocity, the unfortunate reality is that these forms of violence are not well measured. What is more, existing datasets that do track civilian injuries (e.g., Worldwide Atrocities Dataset) or sexualized violence (e.g., Sexual Violence in Armed Conflict Dataset) do not have information regarding whether specific groups are targeted.[29]

We selected two sources of data to create our measure of group-based mass atrocities. These datasets include the: (1.) Political Instability Task Force Dataset on Mass Killings (1972–2022); and the (2.) Targeted Mass Killing Dataset (1972–2020, with data directly provided to the authors for 2021 and 2022).[30] The PITF mass killings variable was meticulously coded by the PITF each year and provides an important measure of mass atrocity, involving violence committed by governments that results in 1,000 or more deaths over a sustained period of time against a discrete group. Specifically, the sustained period lasts no more than two years; it begins in the first year in which more than 500 people are killed, and ends when two consecutive years of fewer than 500 fatalities are recorded.

The Targeted Mass Killings Dataset allowed us to include government and non-governmental violence. This dataset defines a targeted mass killing as the "direct killing of noncombatant members of a group by a formally organized armed forced that results in twenty-five or more deaths in an annual period, with the intent of destroying the group or intimidating the group by creating a perception of imminent threat to its survival."[31] We consider violence that rises above the threshold of 500 deaths in any given year to focus on large mass atrocities, as previously noted. Additionally, note that while the PITF mass killing variable includes any discrete group, the Targeted Mass Killing Dataset includes targeting of ethnic, political, or religious groups. These measures are correlated at .38, indicating a modest relationship yet also underscoring that they are covering distinct episodes of violence.

---

[27] These include Andorra, Antigua and Barbuda, Bahamas, Barbados, Belize, Brunei, Curacao, Dominica, French Guinea, Grenada, Guadalupe, Holy See, Hong Kong, Lichtenstein, Maldives, Malta, Monaco, New Caledonia, Palau, San Marino, Sao Tome-Principe, Seychelle, St. Vincent and the Grenadines, Taiwan, Vanuatu, and Western Samoa. Several other countries that are often missing data in international datasets are likewise not included in the analysis (e.g., South Vietnam).

[28] Charles Butcher et al., "Introducing the Targeted Mass Killing Data Set for the Study and Forecasting of Mass Atrocities," *Journal of Conflict Resolution* 64, no. 7–8 (2020), 1524–1547, accessed July 25, 2024, https://doi.org/10.1177/0022002179896405.

[29] Please contact the first author for information about models of mass atrocities that are not committed against a specific group, as we have also employed this method to forecast these types of atrocities as well.

[30] While data beyond 2020 are not public, the creators of the dataset provided us with episodes of targeted mass killing that are included in the 2021 and 2022 data.

[31] Butcher et al., *Introducing the Targeted Mass Killing Data Set*.

Thus, our dependent variable includes group-based mass killings by government and non-governmental actors in which at least 500 people are killed within a 12-month period. The actors engaging in the violence do not need to be within a single state's borders, though in most cases episodes of violence are confined within a single nation-state. In total, 784 country-years in the full database were coded as "1," involving 85 onsets (operationalized as a year with a group-based mass atrocity when the prior year did not involve a mass atrocity), between 1972 and 2022.[32]

Independent Variables

Given prior literature, we explored a host of risk factors that studies have linked to mass atrocities and/or that are included in existing risk assessment models. Specifically, we assessed variables included in recent studies of mass atrocities, the six models reviewed in our literature review, as well as a myriad of other theoretically meaningful variables, ultimately exploring hundreds of variables. To warrant inclusion, we examined the ROC curve to assess whether a possible risk factor impacted model predictiveness, as well as typical measures of model fit. The definitions, sources, and operationalizations for the independent variables included in all analyses are detailed in our online appendix, though we briefly explain them here as well.

Conflict and Political Upheaval

Much research has found that prior atrocity is predictive of future atrocity; put simply, countries that experienced atrocity are more likely to experience it again, for a variety of reasons. Accordingly, one of our core predictors is prior atrocity, and we measure prior atrocity in two ways. First, we include the percentage of years in our dataset that a country experienced an atrocity. Second, we also measure whether there was an atrocity in the prior year. An atrocity in the prior year is most predictive of future atrocity, while the percentage of years with prior atrocities more fully considers a country's history.

One of the other strongest risk factors of atrocity is the presence of current conflict. We thus incorporate a measure of whether armed conflict is present, including internal conflict, internationalized internal conflict, ethnic war, revolutionary war, or armed attack. Threat to those in power can impact the occurrence of atrocity as well, often because leaders respond to threat with violence. Consequently, even though we examine atrocities committed by state and non-state actors, we include several measures to indicate whether a regime is experiencing threat. These include coups (both whether there were coup attempts and successful coups), riots/demonstrations, and assassinations. Additionally, we include whether there were major elections (e.g., President, parliamentary, etc.). Again, descriptions, sources, and operationalization of all independent variables can be found in the online appendix.

State Structure, Capacity, and Society

We also include an indicator of whether the leader of the country has unlimited authority. Such leaders may be more likely to commit atrocities, and their regimes may also see more unrest from anti-state actors. We likewise incorporate an electoral democracy index to capture whether elections are free and fair, as well as the freedom of the press. Moreover, we include a measure of whether there is factionalism in the regime since prior work has tied factionalism to atrocities.

As unplanned leadership changes can facilitate the rise of repressive leaders and unrest, we also include indicators of extra-legal and unconventional leadership change. To capture potential struggles over power and how these struggles map onto identities within a country, we also incorporate an indicator of whether the ethic or religious identity of presidents, prime ministers, cabinet members, or other political elites is a recurring issue of contention.

Long-standing regimes are less likely to experience atrocities, so we include a measure of political instability as well as a measure of whether the regime is in transition. A country's infant mortality rate is one of the best indicators of state capacity that also has consistent data

[32] In the final analyses, some country-years drop out due to missing data. This is thus the count of what occurred, though if data are missing on any independent variable for a country-year, that country-year drops the entire year out of the analysis.

across many countries and years, and we consequently use this variable as well. To capture discrimination—which is often a precursor to atrocities that target groups—we incorporate a measure of state-led discrimination. Finally, we include population and population density.

International Factors

We include several core variables to assess international dynamics as well. Nearby conflict has been tied to the presence of conflict within one's own borders, and we accordingly include a measure of conflict within 200 kilometers. Trade openness may protect against atrocities such that we include whether a country is a member of the World Trade Organization (previously GATT). Finally, we include an indicator for region, for time, as well as for the Post-Cold War period given the difference in geopolitics following the end of the Cold War. Again, full descriptions and sources of each of these variables can be found in the online appendix.

Strategies of Analysis

In light of our goal of integrating case-based and quantitative methods for atrocity forecasting, we engage in three distinct strategies of analysis, including: (1.) a configurational analysis of the absence of mass atrocity; (2.) event history analysis of the presence of mass atrocity; and (3.) random forest analysis of the presence of mass atrocity, as we detail shortly. We conduct each analysis separately and then integrate the results to arrive at a set of forecasts. All models include the years 1972 to 2022 (or, for the configurational analysis, 2022 unless otherwise specified). We outline the three strategies before turning to our results examining the likelihood of group-based atrocities, including countries with ongoing atrocities.[33]

Configurational Analysis of the Absence of Mass Atrocity

We begin with configurational analysis. Configurational analysis—most commonly associated with Qualitative Comparative Analysis (QCA)—is a case-oriented approach that was developed to identify necessary and sufficient conditions for an outcome or, as we address shortly, the absence of an outcome.[34] This approach is particularly well suited to forecasting mass atrocity, as configurational analysis focuses on identifying multiple paths to a single outcome and privileges causal complexity.[35] Thus, rather than assessing the average effect of a particular variable when all other variables are held constant, configurational analysis aids in the identification of combinations of conditions that are associated with an outcome of interest. Configurational analysis also allows for the identification of asymmetric relationships between conditions; in other words, an association between the presence of a condition and an outcome does not mean that the absence of the condition will be associated with the outcome's absence.

Importantly, in our models, the outcome of interest is the *absence* of mass atrocity. Stephen McLoughlin has notably argued that little research on mass atrocity has actually examined the causes of peace and stability.[36] In line with this, newer work in forecasting rare outcomes has illustrated the value in analyzing the conditions in which the rare outcomes did *not* occur.[37] This value stems from the fact that the conditions leading to the onset of mass atrocity are indeed not the same as the conditions leading to their absence. As such, because mass atrocities may emerge for a variety of reasons, a productive approach to forecasting is to

---

[33] Contact the first author for information on models regarding new onsets.

[34] Despite the value of this method, little work has used configurational analysis to study mass atrocity. For an important exception, see Timothy Williams, "More Lessons Learned from the Holocaust - Towards a Complexity-Embracing Approach to Why Genocide Occurs," *Genocide Studies and Prevention* 9, no. 3 (2016), 137–153, accessed July 26, 2024, http://dx.doi.org/10.5038/1911-9933.9.3.1306.

[35] Benoît Rihoux and Charles C. Ragin, *Configurational Comparative Methods: Qualitative Comparative Analysis (QCA) and Related Techniques* (Thousand Oaks: SAGE Publications, 2008).

[36] Stephen McLoughlin, *The Structural Prevention of Mass Atrocities: Understanding Risk and Resilience* (New York: Routledge, 2014).

[37] Eric W. Schoon et al., "Precluding Rare Outcomes by Predicting Their Absence," *PloS One* 14, no. 10 (2019), 1–13, accessed July 25, 2024, https://doi.org/10.1371/journal.pone.0223239.

invert the problem and focus specifically on identifying conditions that are historically associated with the complete absence of mass atrocities. We follow this advice, suggesting that identifying conditions that are not associated with the onset of mass atrocities enables the reduction of the population of candidate cases for mass atrocities, thereby facilitating a more targeted approach to identifying cases at great risk of onset.

For our configurational analysis, the presence and absence of every condition and outcome is identified for each case and numerically coded as either a 1 or a 0.[38] We then follow the steps outlined by Eric Schoon and co-authors in a recent article on predicting rare outcomes, which involve running Monte Carlo experiments as further explained in their article.[39] Please see the online appendix for additional details on methodology.[40]

Event History Analysis

The second method we use is event history analysis. Specifically, we employ a discrete-time hazard model. In this model, the hazard is the instantaneous propensity that an event will occur. Unlike many models, hazard models analyze the effect of time and allow for time-varying predictors.[41] Please see the online appendix for more information about the model, as well as for the actual tables tied to the model and information regarding missing data and lags.

We also undertook two methods to assess the predictive capacity of the event history models. Following Michael D. Ward, Brian D. Greenhill, and Kristin M. Bakke, we assessed the models' predictive power using the area under the ROC curve, which is known as the AUC and has an optimal value of 1.[42] The ROC curve plots the relationship between the rate of false positives (the number of incorrectly predicted atrocities divided by the total number of cases where atrocities did not happen) and the rate of true positives (the number of correctly predicted atrocities divided by the total number of cases where atrocities did happen). A second, stricter option is to assess something known as out-of-sample forecasts. While the ROC curve allows a researcher to assess how a model predicts outcomes within the same data (called in-sample forecasting), out-of-sample forecasts predict outcomes in new data. Essentially, part of the dataset is used to predict other data that are excluded from the models. Accordingly, we ran K-fold cross validation tests to assess the average out-of-sample AUC. In line with Ward and colleague's suggestions, we assessed five folds across 10 iterations.

Random Forest Models

Finally, we fit random forest models using the same data and operationalizations as the event history analysis. In this sense, the variables are the same, though they are fit differently. Random forests are a standard machine learning algorithm that have several benefits for this task. Specifically, they automatically fit non-linear relationships[43] between predictors and the outcome, allow for more complex interactions among predictors, and incorporate automated variable selection.

However, random forests are harder to interpret than standard approaches as they do not produce simple coefficients and standard errors. Additionally, the flexibility of random

---

[38] Configurational analysis can be conducted on valued data as well, but for our purposes we focus on binary outcomes, commonly referred to as crisp-set analysis. Note also that a drawback of this method is that it does not accommodate large numbers of variables.

[39] Schoon et al., *Precluding Rare Outcomes by Predicting Their Absence*.

[40] As we explain in the appendix, for the purposes of this article, a coverage score of 20 percent is treated as a minimum threshold for inclusion. See Charles C. Ragin, *Redesigning Social Inquiry: Fuzzy Sets and Beyond* (Chicago: University of Chicago Press, 2008).

[41] Paul D. Allison, *Event History Analysis: Regression for Longitudinal Event Data* (Thousand Oaks: SAGE Publications, 1984).

[42] Ward et al., *The Perils of Policy by P-Value*.

[43] Unless otherwise specified, models like the discrete time hazard model assume linear relationships. Yet, complex relationships in real life are rarely so simple (despite the assumption of linear relationships across a host of modeling approaches).

forest models can lead to overfitting if the models are not properly cross-validated and tested. That said, relatively new approaches for interpretable machine learning and fairly simple validation techniques can address this drawback. As such, to guard against overfitting and performing a predictive test, we split the sample into training and test samples. The training set includes all country-years until 2017, while the test set includes observations from 2018-2022. The algorithm was therefore trained on the data until 2017,[44] while the most recent years were used to assess out-of-sample performance. If the model predicts well in the training sample but not the test sample, it is clearly overfit and is unlikely to predict future atrocities well. As further indicated below, this was not the case, and the models perform well.

Results

We first present results of the configurational analysis to identify the combinations of factors that are consistently associated with the absence of mass atrocities. This allows us to ascertain the countries that are at particularly low risk of mass atrocities and, in essence, remove them from the risk pool.[45] We then turn toward two forms of quantitative analysis of the presence of mass atrocity.

Configurational Analyses of the Absence of Mass Atrocity

We wrote a function that computed all possible configurations of the variables in our online appendix with the goal of finding all configurations that are associated with the complete absence of mass atrocities and account for at least 20 percent of the data (for a discussion of why we use 20 percent as a cutoff, please see the online appendix). The analysis identified four distinct configurations of conditions collectively account for 58.9 percent of all country-years. Individual countries can be represented by more than one configuration if they exhibit the characteristics of more than one configuration, leading to varying degrees of overlap in coverage. Again, the way the conditions are operationalizations is detailed in the online appendix.

Table 1. Configurational Coverage for the Absence of Group-Based Mass Atrocities

Configuration	Conditions Representing Configuration	Coverage
Configuration 1	No history of (or present-day) armed conflict, low infant mortality, constrained executive	31.59%
Configuration 2	No history of (or present-day) armed conflict, no history of atrocity, open trade, no major elections, no salient ethnicity	28.66%
Configuration 3	No history of (or present-day) armed conflict, no history of atrocity, low infant mortality, no salient ethnicity	24.86%
Configuration 4	No history of atrocity, low infant mortality, open trade, no major elections, constrained executive	23.86%

Table 1 lists the conditions associated with each configuration, as well as the coverage for each configuration. In essence, these configurations highlight distinct conjunctions of conditions that are each associated with the complete absence of mass atrocity. While each configuration is distinct, there is overlap in the cases they represent, which is why the total coverage for the analysis is not the sum of the coverage for each of the four solutions. Given that some countries are represented by multiple configurations and others are represented by only one, we use representation to distinguish levels of risk such that countries represented by more

[44] To tune the algorithm on the training data, we used three-fold cross-validation, repeated three times. This approach finds the best turning parameters while also guarding against overfitting.

[45] To be clear, we still include these countries in the quantitative analyses because they provide important data. However, as we detail below, if countries have the conditions associated with the lack of mass atrocity, we remove them from the final risk assessments.

configurations should be at lower risk than countries represented by fewer. Examining how many configurations of conditions are associated with the absence of mass atrocities is at best a rough heuristic for distinguishing among cases. However, it allows us to draw some preliminary distinctions between cases within the analysis. This approach is supported by a case-level analysis, which shows that countries with the most configurations tend to be countries with high levels of contemporary stability and peace.

These results paint a fairly coherent picture when viewed together. Enduring stability—measured by either by the absence of armed conflict or an absence of group-based atrocities since 1972, or both—is particularly meaningful and is present in each configuration. Low infant mortality appears in three of the four configurations, underscoring its importance. Finally, constrained executives, open trade, no major elections, and no salient elite ethnicity each appear in two configurations.

Table 2 lists the 98 countries accounted for by each of these configurations as of the year 2022. Based on this, a preliminary assessment is that the 49 countries that represent three or four configurations are at very low risk. Furthermore, the 49 countries within one or two configurations can be classified as relatively low risk, though we note that only appearing in one configuration (or perhaps even two configurations) is not enough to remove the country from the risk set. Finally, Table 3 lists the 68 countries that were in no configuration. These countries should thus be considered at higher risk than all other countries in the analysis, and we now move on to our analysis of countries at risk of mass atrocity before integrating the results.

Table 2. Countries Represented by Configurational Analysis of the Absence of Group-Based Atrocities

Very Low Risk	Total Configurations	Relatively Low Risk	Total Configurations
Albania	4	Belarus	2
Armenia	4	Qatar	2
Australia	4	Argentina	1
Canada	4	Belgium	1
Cuba	4	Benin	1
Denmark	4	Bosnia and Herzegovina	1
Estonia	4	Botswana	1
Finland	4	Cabo Verde	1
Germany	4	Chile	1
Greece	4	China	1
Honduras	4	Colombia	1
Ireland	4	Croatia	1
Jamaica	4	Cyprus	1
Japan	4	Djibouti	1
Lithuania	4	Ecuador	1
Luxembourg	4	El Salvador	1
Mauritius	4	Eswatini	1
Mexico	4	Gabon	1
Netherlands	4	Georgia	1

New Zealand	4	Ghana	1
North Macedonia	4	Guinea-Bissau	1
Norway	4	Guyana	1
Oman	4	Hungary	1
Panama	4	Iran	1
Poland	4	Israel	1
Slovakia	4	Jordan	1
Spain	4	Lebanon	1
Switzerland	4	Libya	1
Taiwan	4	Madagascar	1
Thailand	4	Malawi	1
United Arab Emirates	4	Malaysia	1
United Kingdom	4	Moldova	1
Uruguay	4	Mongolia	1
Venezuela	4		
Montenegro	1	Morocco	1
Austria	3	Namibia	1
Brazil	3	Niger	1
Bulgaria	3	Paraguay	1
Costa Rica	3	Peru	1
Czechia	3	Romania	1
Fiji	3	Russia	1
France	3	Serbia	1
Italy	3	Singapore	1
Korea, South	3	Solomon Islands	1
Kuwait	3	Sri Lanka	1
Latvia	3	Suriname	1
Portugal	3	Tanzania	1
Slovenia	3	Trinidad	1
Sweden	3	Vietnam	1
Tunisia	3		

Table 3. Countries in No Configurations (Higher Risk)

Afghanistan	Kyrgyzstan
Algeria	Laos

Angola	Lesotho
Azerbaijan	Liberia
Bahrain	Mali
Bangladesh	Mauritania
Bhutan	Mozambique
Bolivia	Nepal
Burkina Faso	Nicaragua
Burma	Nigeria
Burundi	Pakistan
Cambodia	Papua New Guinea
Cameroon	Philippines
Central African Republic	Republic of Congo (ROC)
Chad	Rwanda
Comoros	Saudi Arabia
Cote d'Ivoire	Senegal
Democratic Republic of Congo (DRC)	Sierra Leone
Dominican Republic	Somalia
Egypt	South Africa
Equatorial Guinea	South Sudan
Eritrea	Sudan
Ethiopia	Syria
Gambia	Tajikistan
Guatemala	Timor-Leste
Guinea	Togo
Haiti	Turkey
India	Turkmenistan
Indonesia	Uganda
Iraq	Ukraine
Kazakhstan	Uzbekistan
Kenya	Yemen
Korea, North	Zambia
Kosovo	Zimbabwe

Event History Analyses and Random Forests

The online appendix includes the models from the event history analyses, and here we present the predicted probabilities of a group-based mass atrocity derived from these models.

Importantly, we chose to measure a key independent variable in two distinct ways, and we present all forthcoming results with both measures. Specifically, prior atrocity is highly predictive of future atrocity, though there are several ways to measure prior atrocity. We thus ran all models with two different measures: (1.) the percentage of all country-years included in the analysis with a prior group-based atrocity,[46] as well as; (2.) whether there was a group-based atrocity in the prior year (in essence, a one-year lag). We chose the latter because it is the most predictive way to measure prior atrocity, and we chose the former because the presence of prior atrocities more generally aids in the identification of new onsets of atrocity. As such, these two measures yield similar yet distinct findings, and we thus show models with each key independent variable.

With respect to the predictive power of the models, the AUC (again, area under the curve with regard to the receiver operating characteristic curve) is .9734 when prior atrocity is measured as a percentage of country-years experiencing a mass atrocity since 1972, and it is .9923 when prior atrocity is lagged one year. A perfect score is 1.00; as such, the model performs very well.[47]

Out-of-sample forecasts perform similarly well. K-fold cross-validation tests revealed that the average out-of-sample AUC (for five folds across 10 iterations) was .9745 with the percentage of mass atrocities measure included in the models, and .9925 with the lagged atrocity measure included. For comparison, the out-of-sample AUC in Jay Ulfelder's forecasts of mass atrocity was approximately 0.8.[48]

Table 4 includes the countries with the highest predicted probabilities of group-based atrocity. We include 40 countries but note that the probability of mass atrocity declines significantly within the first 10 countries on each list. Before thoroughly assessing the countries listed in Table 4, however, we first turn toward an overview of the results of the random forest models.

Table 4: Event History Analysis Predictions of Group-Based Atrocity

Country	**Prior Atrocity (Lag)**	**Country**	**Prior Atrocity (Percentage)**
Ethiopia	0.9813178	**Burma**	0.952249
Burma	0.9810907	**Ukraine**	0.8484928
South Sudan	0.9807745	Iraq	0.7636069
Syria	0.9770141	**South Sudan**	0.5375273
Korea, North	0.8733243	**Ethiopia**	0.454165
Philippines	0.6177385	**Korea, North**	0.4026543
Sudan	0.4438525	Nigeria	0.3677929
Ukraine	0.3075983	Democratic Republic of Congo	0.3335294
Cameroon	0.0986332	**Syria**	0.3267213
Burkina Faso	0.068939	India	0.2339687
Somalia	0.0684514	Afghanistan	0.1809483
Iraq	0.0492098	Pakistan	0.1767996

[46] As we start analysis in 1972, we also explored the percentage of years since 1955 (when data become available on the dependent variable), and results were not substantially different.

[47] These values are significantly higher than those reported by common models of civil war onset, which often range between .75 and .86. See Ward et al., *The Perils of Policy by P-Value*. Note, however, that Rost's AUC for his study of genocide was .962. See Rost, *Will It Happen Again?*.

[48] Jay Ulfelder, "Forecasting Onsets of Mass Killing," accessed July 25, 2024, http://dx.doi.org/10.2139/ssrn.2056306.

Nigeria	0.0432139	Somalia	0.1507694
Pakistan	0.0367839	Sri Lanka	0.1351897
Democratic Republic of Congo	0.0316439	Burkina Faso	0.1350735
India	0.0294759	Cameroon	0.1322279
Mali	0.0189328	**Philippines**	0.1293552
Azerbaijan	0.018238	Turkey	0.1096285
Egypt	0.0179278	Indonesia	0.0675647
Yemen	0.0175253	Mozambique	0.0675348
Turkey	0.0172192	Angola	0.0472629
Mozambique	0.0159727	Azerbaijan	0.0469588
Central African Republic	0.0144276	**Mali**	0.0393489
Afghanistan	0.0133843	Central African Republic	0.0342379
China	0.0106548	Burundi	0.028075
Equatorial Guinea	0.0093824	Algeria	0.0269215
Saudi Arabia	0.0069819	Colombia	0.0250758
Sri Lanka	0.0062042	Egypt	0.0249789
Turkmenistan	0.0060248	Yemen	0.0224538
Dominican Republic	0.0055677	Republic of Congo (ROC)	0.0184638
Indonesia	0.0054003	Equatorial Guinea	0.018234
Algeria	0.0050383	Lebanon	0.0147935
Lebanon	0.0045131	Russia	0.0142623
Bahrain	0.003884	Zimbabwe	0.0141122
Uzbekistan	0.0037696	Rwanda	0.0131804
Republic of Congo (ROC)	0.0032953	Iran	0.011911
Iran	0.0032904	**Sudan**	0.0110853
Zimbabwe	0.0032125	Kosovo	0.0109777
Haiti	0.003115	Peru	0.0096966
Burundi	0.0030736	Turkmenistan	0.008609

The random forest models incorporate the same data and operationalizations as the event history analyses (see Table 1 in the online appendix). As described in our methods section, these models employ a significantly different modeling approach. To illustrate, we fit the models on a dataset covering 1972 to 2017, and we test them based on a dataset that runs from 2018 to 2022.

Like the event history analyses, these models fare incredibly well. The AUCs for the models are extremely high. Specifically, for the models measuring prior atrocity as a lagged variable, the AUC is .998 for the training data (1972–2017), and .999 for the testing data (2018–2022). The AUC drops somewhat for the models where prior atrocity is measured as a percentage, with a training AUC of .1 and a testing AUC of .982.

Table 5 presents the results from the random forest models. As with the event history analyses, we display the predicted probabilities with prior atrocity measured in two main ways. We likewise show the top 40 countries but again note that the level of risk decreases significantly after the first 10.

Table 5: Random Forest Model Predictions of Group-Based Atrocity

Country	Prior Atrocity (Lag)	Country	Prior Atrocity (Percentage)
South Sudan	0.930423	Korea, North	0.927391
Korea, North	0.899419	South Sudan	0.861344
Ethiopia	0.888884	Burma	0.851664
Burma	0.866478	Nigeria	0.716132
Syria	0.800855	Iraq	0.700184
Sudan	0.460374	Syria	0.615482
Philippines	0.286716	Ethiopia	0.604657
Afghanistan	0.235915	Afghanistan	0.595321
Ukraine	0.188385	Democratic Republic of Congo	0.289119
Nigeria	0.142789	Philippines	0.250928
Iraq	0.138671	Burkina Faso	0.223159
Burkina Faso	0.135674	Azerbaijan	0.184326
India	0.103479	Ukraine	0.180413
Azerbaijan	0.097341	Pakistan	0.163568
Cameroon	0.079252	India	0.154675
Uzbekistan	0.076706	Cameroon	0.136678
Egypt	0.07282	Egypt	0.131689
Pakistan	0.070409	Bulgaria	0.100686
Sri Lanka	0.061167	Central African Republic	0.09873
Bulgaria	0.051213	Uzbekistan	0.093258
Yemen	0.048469	Mozambique	0.091015
Democratic Republic of Congo	0.042819	Sri Lanka	0.09034
Mali	0.034424	Somalia	0.087268
Somalia	0.032617	Colombia	0.078055
Serbia	0.031618	Mali	0.077377
Turkey	0.028754	Yemen	0.065166
Lebanon	0.027047	Turkey	0.06071
China	0.0267	Indonesia	0.047426

Mozambique	0.026427	Angola	0.04704
Kazakhstan	0.01975	Lebanon	0.041009
Bahrain	0.018838	Russia	0.034902
Thailand	0.016654	Equatorial Guinea	0.034236
Chad	0.015774	Chad	0.0314
Poland	0.015639	Haiti	0.030645
Ecuador	0.014903	Kazakhstan	0.026872
Bosnia and Herzegovina	0.014698	Poland	0.024516
Saudi Arabia	0.01351	Thailand	0.02415
Equatorial Guinea	0.013153	Peru	0.021157
Russia	0.01189	Libya	0.020966
Central African Republic	0.011842	Serbia	0.020335

Assessing Risk Across the Models

As a last major step, we integrate results from all models presented thus far to ascertain the countries most at risk of a group-based mass atrocity. Specifically, Table 6 indicates the number of the four quantitative models that resulted in a country being listed in the top 40 countries at risk. This number can be found in the second column, and a higher number indicates greater risk. Finally, the last column in Table 6 indicates how many configurations each country appeared in for the configurational analysis outlined in Tables 1–3; and again, appearing in all four configurations would signify the lowest risk, while zero signifies the highest risk.

Table 6. Assessing Risk of Group-Based Mass Atrocity Across the Models

Country	Average Probability for Group-Based Atrocities	Total Configurations
Burma	0.9128705	0
South Sudan	0.8275172	0
Korea, North	0.7756971	0
Ethiopia	0.7322558	0
Syria	0.6800182	0
Iraq	0.4129178	0
Ukraine	0.3812221	0
Philippines	0.3211845	0
Nigeria	0.3174822	0
Afghanistan	0.2563921	0
Sudan	0.2324719	0
Democratic Republic of Congo	0.1742779	0
Burkina Faso	0.1407114	0

India	0.1303997	0
Pakistan	0.11189	0
Cameroon	0.1116978	0
Azerbaijan	0.086716	0
Somalia	0.0847764	0
Sri Lanka	0.0732253	1
Egypt	0.0618537	0
Turkey	0.054078	0
Mozambique	0.0502374	0
Uzbekistan	0.0446463	0
Mali	0.0425206	0
Central African Republic	0.0398092	0
Bulgaria	*0.0388742*	*3*
Yemen	0.0384036	0
Indonesia	0.0320452	0
Colombia	0.0261528	1
Angola	0.0258479	0
Lebanon	0.0218406	1
Equatorial Guinea	0.0187515	0
Russia	0.015891	1
China	0.0152261	1
Serbia	0.0145586	1
Chad	0.0124539	0
Kazakhstan	0.012213	0
Bahrain	0.0117598	0
Haiti	0.0115684	0
Algeria	0.011494	0
Thailand	*0.0113883*	*4*
Peru	0.0108084	1
Poland	*0.0103831*	*4*
Burundi	0.0093177	0
Ecuador	0.0082177	1
Bhutan	0.0080811	0
Bosnia and Herzegovina	0.0080024	1

Cambodia	0.0077459	0
Bangladesh	0.0070783	0
Turkmenistan	0.0069193	0

There are 50 countries in Table 6, though a number of these countries appear in several configurations. We thus would remove Bulgaria (three configurations), Thailand (four configurations) and Poland (four configurations) from the list and have accordingly italicized them in the table. Each of these countries had exceedingly low probabilities in the quantitative analyses as well. There are other countries that appear in two configurations and that have significantly higher probabilities, however, meaning that we would not remove them from the risk set. We would, nonetheless, weigh the results from all analyses, and situations like these are where the configurations are particularly useful.

Figure 1 depicts the results. Taken together, the top 15 countries with the highest predicted probabilities that are all above .1 *and* that appear in no configurations are what we would deem the countries most at risk. These countries can be found in the first 15 lines of Table 6 and are, in order of highest risk, Burma/Myanmar, South Sudan, North Korea, Ethiopia, Syria, Iraq, Ukraine, Philippines, Nigeria, Afghanistan, Sudan, the Democratic Republic of the Congo, Burkina Faso, India, and Pakistan. Notably, our own case-based knowledge of these cases supports these findings. Another way to calculate high risk would be to identify the countries with a predicted probability of .5 or greater in any of the four models that also do not appear in any configurations. These include Burma/Myanmar, South Sudan, North Korea, Ethiopia, and Syria. Finally, while we emphasize those with the highest predicted probabilities and no configurations here, we note that countries appearing in three or four of the quantitative model tables that do not appear in more than one configuration should be seen as relatively high risk as well.

Figure 1: Risk Assessments for Ongoing Group-Based Atrocities

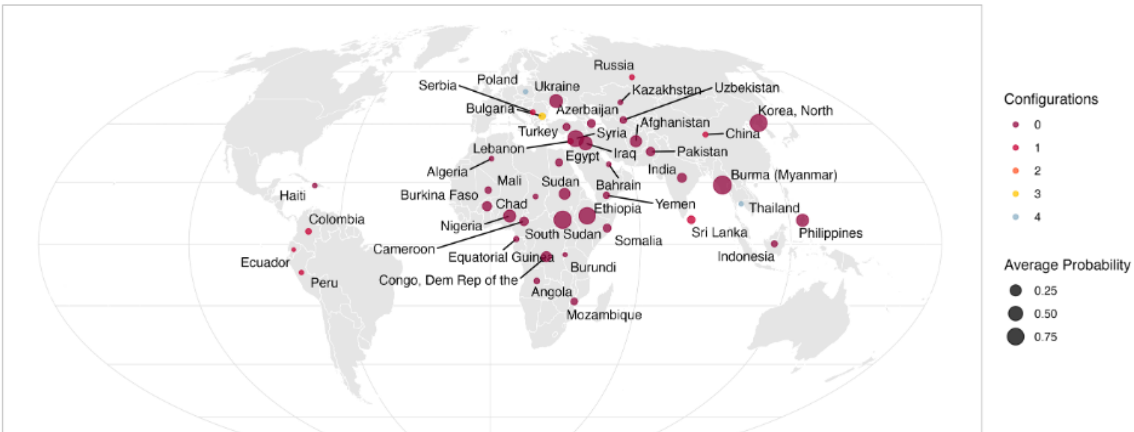

All countries with a risk above .1 should be viewed as at risk of group-based mass atrocities. We also repeated our efforts to forecast new onsets only (excluding ongoing mass atrocities). All models and information can be obtained from the first author, and we include in Table 7 a list of countries at risk of new onsets that can be used alongside Table 6 (and, again, were produced using the same methods outlined above).

Table 7. Predicted New Onsets of Group-Based Mass Atrocity

Country	Probability of New Onset of Group-Based Atrocity	Number of Configurations
Burkina Faso	0.0692616	0

Pakistan	0.0575881	0
Cameroon	0.0406754	0
Nigeria	0.0370434	0
Yemen	0.0275626	0
Turkey	0.0150126	0
Egypt	0.0121469	0
Somalia	0.0114449	0
Republic of Congo (ROC)	0.011212	0
Central African Republic	0.0105529	0
Azerbaijan	0.0096429	0
Kazakhstan	0.0083294	0
Democratic Republic of Congo (DRC)	0.0066166	0
Mozambique	0.0041144	0
Iraq	0.0040466	0
Angola	0.0036097	0
Zimbabwe	0.0033072	0
Gabon	0.0029618	1

Discussion and Conclusion

This article has outlined an approach to forecasting mass atrocities that integrates case-based and quantitative methods. In doing so, we present what we believe is a more holistic view of risk than methods employing either quantitative or case-based methods in isolation. Specifically, we examine the risk of group-based mass atrocities, which again we operationalize as mass killings targeting specific groups. Our assessments are well above 97% accurate with respect to forecasting, and Table 6 presents forecasts with various ways to measure risk.

While we believe these forecasts are robust, it is important to qualify our risk assessments in several ways. First, we have only examined group-based mass atrocities. Forecasts of mass atrocities that do not target civilians (not shown) result in some similarities but also core differences with respect to the countries at greatest risk (e.g., the DRC, Nigeria, Pakistan, Iraq).

Second, as other forecasting approaches that we are aware of, the methods we have employed rely upon dated data. Indeed, data collection efforts always lag several years behind real time, which presents a significant problem for forecasting. One option to address this issue would be to impute future data to be able to carry the forecasts forward. Many of the variables do not vary meaningfully over time, however, such that the imputations may only provide marginal improvements, if at all. As such, we suggest that a main way to remedy this issue would be to supplement the forecasts with qualitative data. Such data would include information on the countries that experienced mass atrocities after data collection efforts ended (e.g., in this case, in 2023), as well as countries that experienced changes in significant risk factors.

These challenges are acutely demonstrated by the fact that our models do not identify Israel as being at risk for group-based atrocities. Hamas's October 7, 2023 attack on Israel and Israel's subsequent violence in Gaza both meet our operational definition for a group-based mass atrocity. However, Israel is only identified as being "relatively low risk" by the configurational analysis and is not flagged by the others, ultimately not making it onto our list of countries most likely to experience a group-based mass atrocity. A separate analysis of *non*-group-based mass atrocities does flag Israel as being among the top 20 countries at risk for atrocities. The fact that our data ends in the preceding year, yet fails to flag this case, highlights the methodological implications of different approaches to operationalizing mass atrocities. This difference highlights an additional challenge related to data, which is that the data used for forecasting are not politically neutral. For example, we contend that decisions surrounding who to count as a civilian versus non-civilian are necessarily political decisions. Moreover, we believe the fact that Israel, but not Palestine, is included in available data likely contributes to our failure to forecast the events of 2023. Thus, the fact that our models did not identify Israel as a candidate for group-based mass atrocities pointedly illustrates some of the many persistent challenges of forecasting these events.

Keeping these issues in mind, we nonetheless believe there is value in forecasting. Moreover, we believe that integrating forecasting methods offers measurable improvements over any one method in isolation. For those seeking to use these forecasts, we would consider the countries in Table 6 as the primary risk set and again encourage taking both the predicted probabilities and the number of configurations countries appear in into account. As with other forecasting methods, these models should be coupled with early warning endeavors. In fact, the models may also inform such efforts. For instance, if the risk of atrocity appears higher in a country that is not in Table 6, an analyst could examine the configurational analysis to get a general sense of the risk for that country. We likewise underscore that local efforts to forecast violence are also incredibly important and must be integrated with global efforts like this one. We thus hope that our models will contribute to global efforts to better predict and, as such, prevent mass atrocity.

Bibliography

Allison, Paul D. *Event History Analysis: Regression for Longitudinal Event Data*. Thousand Oaks: SAGE Publications, 1984.

Anderton, Charles H., and Jurgen Brauer. "Mass Atrocities and Their Prevention." *Journal of Economic Literature* 59, no. 4 (2021), 1240–1292.

Anderton, Charles H., and John R. Carter. "A New Look at Weak State Conditions and Genocide Risk." *Peace Economics, Peace Science and Public Policy* 21, no. 1 (2015), 1–36. Accessed July 25, 2024. https://doi.org/10.1515/peps-2014-0008.

Atrocity Forecasting Project. "Brief Report on AFP 2021-23 Forecasts of Genocide/Politicide." n.d. Accessed January 29, 2023. https://politicsir.cass.anu.edu.au/sites/default/files/AFP_2021-23_BriefReport.pdf.

Butcher, Charles, Benjamin E. Goldsmith, Sascha Nanlohy, Arcot Sowmya, and David Muchlinski. "Introducing the Targeted Mass Killing Data Set for the Study and Forecasting of Mass Atrocities." *Journal of Conflict Resolution* 64, no. 7–8 (2020), 1524–1547. Accessed July 25, 2024. https://doi.org/10.1177/0022002179896405.

Chenoweth, Erica, Evan Perkoski, and Sooyeon Kang. "State Repression and Nonviolent Resistance." *Journal of Conflict Resolution* 61, no. 9 (2017), 1950–1969. Accessed July 25, 2024. https://doi.org/10.1177/0022002717721390.

Colaresi, Michael, and Sabine C. Carey. "To Kill or to Protect: Security Forces, Domestic Institutions, and Genocide." *Journal of Conflict Resolution* 52, no. 1 (2008), 39–67. Accessed July 25, 2024. https://doi.org/10.1177/0022002707310427.

Early Warning Project. "Methodology: Statistical Model." *Simon-Skjodt Center for the Prevention of Genocide of the United States Holocaust Memorial Museum/Dickey Center for International Understanding at Dartmouth College*. Accessed January 29, 2023. https://earlywarningproject.ushmm.org/methodology-statistical-model.

Fein, Helen. "Accounting for Genocide after 1945: Theories and Some Findings." *International Journal on Group Rights* 1, no. 2 (1993), 79–106. Accessed July 25, 2024. https://www.jstor.org/stable/24674446.

Goldsmith, Benjamin E., Charles R. Butcher, Dimitri Semenovich, and Arcot Sowmya. "Forecasting the Onset of Genocide and Politicide: Annual Out-of-Sample Forecasts on a Global Dataset, 1988–2003." *Journal of Peace Research* 50, no. 4 (2013), 437–452. Accessed July 25, 2024. https://doi.org/10.1177/0022343313484167.

Harff, Barbara. "Detection: The History and Politics of Early Warning." In *Responding to Genocide: The Politics of International Action*, edited by Adam Lupel and Ernesto Verdeja, 85–110. Boulder: Lynne Rienner, 2013.

----------. "No Lessons Learned from the Holocaust? Assessing Risks of Genocide and Political Mass Murder since 1955." *American Political Science Review* 97, no. 1 (2003), 57–73. Accessed July 25, 2024. https://doi.org/10.1017/S0003055403000522.

Harff, Barbara, and Ted Robert Gurr. "Systematic Early Warning of Humanitarian Emergencies." *Journal of Peace Research* 35, no. 5 (1998), 551–579. Accessed July 25, 2024. https://www.jstor.org/stable/425699.

Heldt, Birger. "Risks, Early Warning and Management of Atrocities and Genocide: Lessons from Research." *Politorbis* 45 (2009), 65–70. Accessed July 25, 2024. https://ssrn.com/abstract=2025007.

Krain, Matthew. "State-Sponsored Mass Murder: The Onset and Severity of Genocides and Politicides." *Journal of Conflict Resolution* 41, no. 3 (1997), 331–360. Accessed July 25, 2024. https://doi.org/10.1177/0022002797041003001.

Krcmaric, Daniel. "Varieties of Civil War and Mass Killing: Reassessing the Relationship between Guerrilla Warfare and Civilian Victimization." *Journal of Peace Research* 55, no. 1 (2018), 18–31. Accessed July 25, 2024. https://doi.org/10.1177/0022343317715060.

Kuper, Leo. *Genocide: Its Political Use in the Twentieth Century*. New Haven: Yale University Press, 1983.

Laban Hinton, Alexander. *Why Did They Kill?: Cambodia in the Shadow of Genocide*. Berkeley, CA: University of California Press, 2005.

Mayersen, Deborah. "Predicting Genocide and Mass Killing." *Journal of Genocide Research* 23, no. 1 (2021), 81–104. Accessed July 25, 2024. https://doi.org/10.1080/14623528.2020.1818478.

McLoughlin, Stephen. *The Structural Prevention of Mass Atrocities: Understanding Risk and Resilience*. New York: Routledge, 2014.

Nyseth Brehm, Hollie. "Re-examining Risk Factors of Genocide." *Journal of Genocide Research* 19, no. 1 (2017), 61–87. Accessed July 25, 2024. https://doi.org/10.1080/14623528.2016.1213485.

Nyseth Nzitatira, Hollie. "Predicting Genocide." In *Genocide: Key Themes*, edited by Donald Bloxham and A. Dirk Moses, 45–74. New York: Oxford University Press, 2022.

Political Instability Task Force, "Political Instability Task Force Mass Killing Measure (1955–2021)," (dataset). Accessed January 29, 2023.

----------. "Political Instability Task Force Worldwide Atrocities Dataset." Accessed January 29, 2023. https://eventdata.parusanalytics.com/data.dir/atrocities.html.

Ragin, Charles C. *Redesigning Social Inquiry: Fuzzy Sets and Beyond*. Chicago: University of Chicago Press, 2008.

Rihoux, Benoît, and Charles C. Ragin, eds. *Configurational Comparative Methods: Qualitative Comparative Analysis (QCA) and Related Techniques*. Thousand Oaks: SAGE Publications, 2008.

Rost, Nicholas. "Will It Happen Again? On the Possibility of Forecasting the Risk of Genocide." *Journal of Genocide Research* 15, no. 1 (2013), 41–67. Accessed July 25, 2024. https://doi.org/10.1080/14623528.2012.759398.

Rothbauer, Paulette M. "Triangulation." In *The Sage Encyclopedia of Qualitative Research Methods*, edited by Lisa M. Given. 892–894. Thousand Oaks: SAGE Publications, 2008.

Schoon, Eric W., David Melamed, Ronald L. Breiger, Eunsung Yoon, and Christopher Kleps. "Precluding Rare Outcomes by Predicting Their Absence." *PloS One* 14, no. 10 (2019), 1–13. Accessed July 25, 2024. https://doi.org/10.1371/journal.pone.0223239.

Stanton, Gregory H. "The Ten Stages of Genocide." *Genocide Watch*, 2023. Accessed July 25, 2024. https://www.genocidewatch.com/tenstages.

Straus, Scott. *Fundamentals of Genocide and Mass Atrocity Prevention*. Washington, DC: United States Holocaust Memorial Museum, 2016.

Ulfelder, Jay. "A Multimodel Ensemble for Forecasting Onsets of State-Sponsored Mass Killing." *American Political Science Association 2013 Annual Meeting* (September 2013), 1–23. Accessed July 25, 2024. https://ssrn.com/abstract=2303048.

----------. "Forecasting Onsets of Mass Killing." Accessed July 25, 2024. http://dx.doi.org/10.2139/ssrn.2056306.

United Nations. *Framework for Analysis for Atrocity Crimes: A Tool for Prevention*. 2014. Accessed on July 26, 2024. https://www.un.org/en/genocideprevention/documents/about-us/Doc.3_Framework%20of%20Analysis%20for%20Atrocity%20Crimes_EN.pdf.

Uzonyi, Gary. "Civil War Victory and the Onset of Genocide and Politicide." *International Interactions* 41, no. 2 (2015), 365–391. Accessed July 26, 2024. https://doi.org/10.1080/03050629.2014.932786.

----------. "Domestic Unrest, Genocide and Politicide." *Political Studies* 64, no. 2 (2016), 315–334. Accessed July 26, 2024. https://doi.org/10.1111/1467-9248.12181.

Valentino, Benjamin, Paul Huth, and Dylan Balch-Lindsay. "'Draining the Sea': Mass Killing and Guerrilla Warfare." *International Organization* 58, no. 2 (2004), 375–407. Accessed July 26, 2024. https://www.jstor.org/stable/3877862.

Verdeja, Ernesto. "Predicting Genocide and Mass Atrocities." *Genocide Studies and Prevention* 9, no. 3 (2016), 13–32. Accessed July 26, 2024. http://dx.doi.org/10.5038/1911-9933.9.3.1314.

Ward, Michael D., Brian D. Greenhill, and Kristin M. Bakke. "The Perils of Policy by P-Value: Predicting Civil Conflicts." *Journal of Peace Research* 47, no. 4 (2010), 363–375. Accessed July 26, 2024. https://www.jstor.org/stable/20752194.

Williams, Timothy. "More Lessons Learned from the Holocaust—Towards a Complexity-Embracing Approach to Why Genocide Occurs." *Genocide Studies and Prevention: An International Journal* 9, no. 3 (2016), 137–153. Accessed July 26, 2024. http://dx.doi.org/10.5038/1911-9933.9.3.1306.

Online Appendix

Table 1. Independent Variables and Conditions (Risk Factors)

Variable	Definition	Source	Event History Analysis and Random Forest Models	Configurational Analysis
Armed Conflict	Armed conflict is defined as having any internal conflict, internationalized internal conflict, ethnic war, revolutionary war or armed attack.	Created with data from PRIO, PITF, and the Cross-National Time-Series Data Archive	0 or 1 for each country-year	*Absence of conflict* is coded with 1 as consistently no armed conflict across a countries' history and 0 if the country ever experienced an armed conflict (since 1972)
Conflict Nearby	Conflicts within countries within 200 kilometers	PITF	0 or 1 for each country-year	
Riots or Demonstrations	Riots or anti-government demonstrations present in country.	Cross-National Time-Series Data Archive	0 or 1 for each country-year	
Assassinations	Any successful assassination attempt within a country	Cross-National Time-Series Data Archive	0 or 1 for each country-year	
Any coup	Whether or not there was a coup (successful or failed attempt)	Center for Systemic Peace	Dummy variable	

Salient elite ethnicity	Whether the ethnic or religious identity of presidents, prime ministers, cabinet, or military junta is a recurring issue of contention	Coded by the PITF until 2000 and has since been updated by Monty Marshall	0=elite ethnicity is not salient; 1=elite ethnicity is salient, the political leadership is representative of the largest communal groups or a coalition of several groups that together constitute a majority; 2 = elite; ethnicity is salient, the political leadership is representative of a minority communal group or a coalition of small groups that together constitute less than a majority.	0=elite ethnicity is salient; 1=elite ethnicity is n not salient
Partial democracy	Scores 1-6 on the Polity V full regime type scale	Coded in Polity V and updated by PITF	0 or 1 for each country-year	
Partial autocracy	Scores -1 to -6 on the Polity V full regime type scale	Coded in Polity V and updated by PITF	0 or 1 for each country-year	

Factionalism indicator	Factionalism indicator based on competitiveness of participation component from Polity data source. A code of 3 for POLPACMP indicates factional or factional / restricted patterns of competition. For this variable: 0=not factional; 1=factional.	PITF extension of Polity IV data on factionalis m	0 or 1 for each country-year; note that countries in transition are coded as 2 so as not to lose these country-years	
Government in transition	Indicator based on whether a government is in transition; this is largely because these governments are not coded on PolityV variables and would drop out of the analysis otherwise (so it is -66, -77, or -88)_	Polity V	0 or 1 for each country-year	
Unlimited authority	Executive constraints variable (polxcons) recoded as binary such that it measures unlimited authority versus other types of executive constraints.	PolityV	1 = Unlimited authority; 0 = Not unlimited authority	*Constrained Executive* is coded as 0 = Unlimited authority; 1 = Constrained executive
Unconventional leadership entry	Manner in which leader entered power (essentially unconventional or not).	Archigos	0=regular means; 1=irregular means or directly imposed by another state	
Extralegal leadership change indicator	Extralegal leadership change indicator (1=yes,0=no)	PITF	Extralegal leadership change indicator (1=yes,0=no)	

State-led discrimination	Coded as 1 = at least one group subject to state-led discrimination; 0 = no groups subject to state-led discrimination.	Center for Systemic Peace and PITF	Percentage	
Infant mortality rate	Number of deaths of infants under one year of age per 1,000 live births in a given year	U.S. Bureau of the Census via PITF	Rate	*Low Infant Mortality* coded as 0 or 1 for infant mortality in the bottom quartile worldwide (1) versus not (0)
Population	Total population based on the de facto definition of population	World Bank Developme nt Indicators	Logged to best fit the data	
Population growth	Annual population growth rate. Population is based on the de facto definition of population, which counts all residents regardless of legal status or citizenship--except for refugees not permanently settled in the country of asylum, who are generally considered part of the population of the country of origin.	World Bank Developme nt Indicators	Annual percentage of population growth	
Trade openness and engagement (WTO/GATT Membership)	Membership in WTO/GATT	World Trade Organizatio n	0 or 1 for each country-year. Note we also assessed the Penn World Tables' measure of trade openness (exports plus imports divided by GDP)	*Open Trade* is coded as 0 or 1 for each country-year with 1 signifying membership

Regime durability	Number of years since a 3-point change in the polity score over a short period of time	PolityV (and updated by the PITF)	Count of years	
Region	Region of the world	PITF	Africa, East Asia, Europe, Latin America, Near East, South and Central Asia; Africa is set as the reference category	
Prior group-based mass atrocity	Measure of whether country has experienced group-based mass atrocity to date, since 1972	Based on the dependent variable	The percentage of prior-country-years that experienced mass atrocity. Also measured as a dummy variable (0 or 1) as well	*No Prior Mass Atrocity* is essentially the reverse such that 1 = no atrocities since 1972
Post-Cold War	Measure to signify the shift in the international system following the Cold War	Country-years after 1990 operationalized as post-Cold-War	0 or 1 for each country-year	
Time	Indicator to capture the passage of time	Year	Numeric year; note models also assessed year-squared and year fixed-effects, though numeric year was best	

Table 2. Event History Analysis Models (N=164 countries)

	Model 1	**Model 2**
Risk Factors	Group-Based Atrocities	Group-Based Atrocities
Percent Group-Based Atrocities	1.050***	
	(1.040 - 1.061)	
Group-Based Atrocity (lagged)		412.293***
		(230.686 - 736.869)
Armed Conflict	14.755***	13.316***
	(8.600 - 25.315)	(8.027 - 22.091)
Riots or Demonstrations	0.837	0.876
	(0.580 - 1.206)	(0.536 - 1.432)
Any Coup, lagged	0.996	0.421*
	(0.583 - 1.701)	(0.172 - 1.029)
Extralegal Leadership Change	2.805***	5.712***
	(1.601 - 4.916)	(2.402 - 13.583)
Regime Durability	0.983	1.001
	(0.955 - 1.012)	(0.979 - 1.023)
Salient Elite Ethnicity (Maj.)	2.014*	1.998**
	(0.922 - 4.401)	(1.056 - 3.781)
Salient Elite Ethnicity (Min.)	1.735	2.079*
	(0.617 - 4.883)	(0.931 - 4.643)
State-Led Discrimination	2.325***	1.756**
	(1.283 - 4.211)	(1.069 - 2.884)
Partial Democracy	1.069	1.141
	(0.443 - 2.581)	(0.591 - 2.202)
Partial Autocracy	0.995	0.957
	(0.469 - 2.114)	(0.503 - 1.823)
Factionalism	0.963	0.930
	(0.483 - 1.919)	(0.548 - 1.578)
Government in Transition	1.149	1.005
	(0.358 - 3.693)	(0.390 - 2.587)
Assassinations	2.136***	2.208***
	(1.316 - 3.466)	(1.259 - 3.873)
Unconventional Leader Entry	0.785	0.926

	(0.368 - 1.674)	(0.508 - 1.687)
Unlimited Authority	3.251**	3.614***
	(1.267 - 8.347)	(1.608 - 8.124)
Population (logged)	1.222	1.248**
	(0.923 - 1.619)	(1.034 - 1.505)
Population Growth	0.697***	0.830**
	(0.582 - 0.836)	(0.706 - 0.977)
GATT/WTO Member	0.380***	0.425***
	(0.193 - 0.749)	(0.263 - 0.687)
Infant Mortality Rate (logged)	1.443	1.511
	(0.708 - 2.942)	(0.905 - 2.524)
Conflict Nearby	0.860	0.933
	(0.478 - 1.549)	(0.557 - 1.565)
East Asia	0.485	1.600
	(0.101 - 2.318)	(0.590 - 4.344)
Europe	0.260*	0.271**
	(0.061 - 1.102)	(0.083 - 0.887)
Latin America	0.604	1.610
	(0.152 - 2.405)	(0.635 - 4.086)
Near East	0.278**	0.595
	(0.088 - 0.873)	(0.253 - 1.402)
South & Central Asia	0.305**	0.413**
	(0.093 - 1.000)	(0.195 - 0.877)
Post-Cold War	1.360	0.809
	(0.515 - 3.593)	(0.365 - 1.795)
Year	0.982	1.004
	(0.942 - 1.023)	(0.971 - 1.039)
Observations	7,357	7,356
*** p<0.01, ** p<0.05, * p<0.1		