Taking Stock of Early Warning for Atrocity Prevention

Report from the 2021 Sudikoff Interdisciplinary Seminar on Genocide Prevention

February 2022

UNITED STATES
HOLOCAUST
MEMORIAL
MUSEUM

SIMON-SKJODT CENTER
FOR THE PREVENTION OF GENOCIDE

CONTENTS

Editor's Note

In 2015 the United States Holocaust Memorial Museum, in partnership with Dartmouth College, launched the Early Warning Project (EWP), a first-of-its-kind publicly available early warning system for mass atrocities. The ultimate goal of the project is to contribute to the prevention of genocide and mass atrocities. By providing governments, civil society groups, and other influential actors with early and reliable warnings of mass atrocities, the project seeks to expand awareness of and opportunities for preventive action.

We know from studying the Holocaust and other genocides that such events are never spontaneous. They are always preceded by a range of early warning signs. If warning signs are detected and their causes addressed, it may be possible to prevent catastrophic loss of life.

The United States Holocaust Memorial Museum's founding charter, written by Holocaust survivor Elie Wiesel, mandates that our institution strive to make preventive action a routine response when warning signs appear. Wiesel wrote, "Only a conscious, concerted attempt to learn from past errors can prevent recurrence to any racial, religious, ethnic or national group. A memorial unresponsive to the future would also violate the memory of the past."

The Museum's Simon-Skjodt Center for the Prevention of Genocide was established to fulfill that vision by transmitting the lessons and legacy of the Holocaust, and "to alert the national conscience, influence policy makers, and stimulate worldwide action to confront and prevent genocide." The Simon-Skjodt Center's Early Warning Project works to fulfill this aspect of the Museum's mandate by using innovative research to identify early warning signs. In doing so, we seek to do for today's potential victims what was not done for the Jews of Europe.

Since the launch of EWP, the professional practice of early warning has grown and become more methodologically diverse. Several other risk assessment and early warning projects have emerged in recent years, both within and outside governments and international organizations. A few of these projects also assess risk of atrocities (e.g, the Atrocity Forecasting Project), while others forecast related crimes, armed conflicts, and other types of instability.[1] In addition to statistical forecasting, some projects—including EWP—have experimented with crowd forecasting methods and qualitative risk assessments.

The 2021 Sudikoff Interdisciplinary Seminar on Genocide Prevention explored methodological and policy-related issues around early warning for atrocity prevention. We sought to share lessons among various forecasting projects, contribute to improved early warning and response mechanisms within the US government, and generate new ideas for future research to improve the field of early warning.

This investment in better understanding of early warning for mass atrocities reflects the US Holocaust Memorial Museum's status as a "living memorial" that learns and adapts over time. It also reflects the commitment of the Museum's Simon-Skjodt Center for the Prevention of Genocide to equip decision

1

makers with the knowledge, tools, and support required to prevent and halt genocide and related crimes against humanity — and EWP's goals to improve the quality and depth of discussions about risks of genocide and mass atrocities and advance the science of early warning for mass atrocities.

Adapting to COVID-19 restrictions, the 2021 Sudikoff Seminar took the form of three webinars, each focused on a particular aspect of early warning: statistical methods for risk assessment and early warning, qualitative early warning assessments, and communicating about risk. In each session, expert discussants provided introductory remarks and then participants considered key questions selected to ensure sharing of learning across projects and between policy makers and scholars.

Experts reflect on early warning

This report compiles essays from several leading experts who participated in the seminar series. We asked them to reflect on one or more key themes and offer recommendations for research and/or policy initiatives that could help current and future leaders address mass atrocities. The collection of essays comprises diverse perspectives and spans reflections on everything from technical approaches to forecasting to practical recommendations on communicating risk to the overall efficacy of forecasting projects.

- **Quscondy-Mohamed Abdulshafi** provides recommendations for how mass atrocity monitoring and preventive interventions can make use of social media-generated big data and artificial intelligence (AI), and suggests the need for a strategic collaboration between social media companies and mass atrocity research institutions.
- **Chiara De Franco** discusses warning as persuasion, including the importance of clarity and specificity in the content of the warnings, the relationship between the warners and those expected to act on the warning, and the importance of developing a warning strategy that fits into the "swirl of intelligence" continuously confronting policy makers.
- **Håvard Hegre** outlines and weighs the benefits of different prediction targets. He argues that the field should move towards "continuous multi-period incidence" (CMPI) targets to avoid arbitrary coding thresholds, enable models to better use available information, and allow for a more nuanced approach to uncertainty in evaluating and describing results.
- **Victor Okorie** addresses the role of researchers in situations where atrocity risk is present, and describes the ethical and logistical challenges of those seeking to conduct research where gatekeepers and conditions demand loyalty to the state. In such circumstances, Okorie argues that researchers are obligated to take on an activist role, portray both victims and perpetrators, and strive for objectivity and reflexivity.
- **Clionadh Raleigh** questions the assumption of many early warning and prediction models that more accurate, timely warnings will result in the prevention of conflict. She argues that global models attempt to create consistency across cases where none actually exists, and that projects' focus on catastrophic events misses the need for response to low-level threats. She recommends a revised approach to early warning focused on studying each unique environment.
- **Paul Slovic** provides an overview of three psychological mechanisms that impede the impact and effectiveness of early warnings: psychic numbing, pseudoinefficacy, and the prominence effect. He concludes with recommendations for those who have been (or may become) victims of atrocity crimes, government officials, non-governmental organizations (NGOs), and the public.

State of the field: Growth, evolution, and recurring challenges

The breadth of perspectives captured in these essays, and in the Sudikoff Seminar discussions themselves, make it impossible to meaningfully distill key takeaways or overarching conclusions. However, here I offer some reflections on the tone of the discussions, attendees, and the future of the field.

The number and diversity of professionals from the US government, foreign governments, academic, and research NGO communities that were interested in participating in the discussions suggest that **the field continues to garner support and attention**. While participants acknowledged the limitations of early warning systems, researchers continue to strive to improve systems and methods and are interested in collaborating across projects for the benefit of the field.

A recurring theme throughout both the discussions and essays in this report is **the need to improve relationships, communications, and life-saving responses--in other words, early warning's human dimensions**--in addition to data, analytic methods, and bureaucratic designs. Whatever the specific approach, early warning initiatives share a core motivation of saving lives. Participants argued for keeping people at the center of discussions of various specific topics, including data availability, bridging the research-policy gap, improving understanding of local contexts, and addressing responsibility and objectivity of researchers who work on early warning systems.

The focus on qualitative research methods, their potential impact, and nuanced reflection on how the US government's Atrocity Assessment Framework can and should be adapted to each unique country-context reflects a general—and important—**trend of localization in the field**. Projects have acknowledged for years that cross-national assessments can only provide a starting point for analysis and that country-specific assessments should be conducted by local experts. The shift we are seeing now is increased recognition that analytical frameworks themselves should be adapted to the local context to ensure maximum impact.

Unsurprisingly, **the challenges of communicating about risk and achieving policy impact persist**. For researchers, communication is too often an afterthought, despite the fact that it is arguably the most important component of a risk assessment. Notably, the session on communicating risk garnered the most interest, suggesting a desire for more information on evidence-based approaches to achieving impact.

Finally, **no one method, model, or system is a silver bullet for early warning and prevention**. Temporal and geographic scope varies across projects and methods. "Early warning" encompasses everything from annual (or even five-year or ten-year) cross-national statistical analyses created in capitals far-removed from the localities at risk, to field-based warnings issued via SMS systems in real-time by individuals who themselves are at risk. Perspectives vary widely on the "best" approaches to risk assessment and warning, and while lessons can be gleaned from the risk communications literature, much depends on the individual receiver of warnings and the structure in which they operate. These observations have led me to conclude that a multi-method approach to risk assessment and early warning for the field as a whole is (a) likely to continue, as each project is developed through different organizational and funding streams, and (b) arguably preferable. Though the field would benefit from improved coordination, that coordination should not attempt to limit the diversity and varied approaches of warning systems.

Mollie Zapata, February 2022

Utilizing Social Media for Atrocity Monitoring and Prevention

Quscondy Abdulshafi, Freedom House

Social media has transformed social mobilization and movement building. Likewise, it is almost impossible to understand the dynamics of any atrocities without considering discourse on social media platforms. This essay argues that the future of atrocity monitoring and prevention must make greater use of social media-generated big data and artificial intelligence (AI) monitoring and preventive interventions. One way to promote this kind of shift would be a strategic collaboration between a mass atrocity research institution and a social media corporation. I illustrate this idea by discussing potential cooperation between the Museum's Simon-Skjodt Center and Facebook.

The power of social media and the digital organizing space

The growing access to mobile internet and social media in fragile states with corrupt ethno-religious extremism and structural social inequalities made social media a powerful tool for organizing the masses toward a shared cause. Despite the fact that social media has helped promote good causes, such as movements to promote local innovations and businesses, these platforms have also been tools for inciting and incubating mass atrocities and genocide.[2] As access to the internet increases and the number of social media users grows exponentially, the power of these platforms to both incite and prevent atrocities increases every day. A few examples below provide evidence of the power and efficacy of social media and the internet in directing and mobilizing the masses for a cause.

- On December 17, 2010, a young Tunisian man named Mohamed Bouazizi set himself on fire on the street in front of a government building in desperation, anger, and protest of government negligence. A video of Mohamed blazing himself was shared on social media and immediately ignited a revolution that would later be known as the Arab Spring.[3]
- In December 2018, a Sudanese professional association used Facebook to plan nationwide protests against the government and, in a few months, successfully toppled one of the world's most repressive regimes despite the government mocking its efforts.[4]
- In July 2021, a Facebook group in San Antonio sparked widespread demonstrations in Cuba after decades of repression and totalitarianism.[5]

Each of these discrete incidents of protest was given momentum by the groups and associations that formed afterward in the digital space and mobilized for real-life action. The heroes behind those social media movements did not necessarily know or meet each other before.

Unfortunately, it is easy to find examples of organized criminal and political groups with violent intentions, including state-sponsored actors and non-state armed actors, using digital space to mobilize and advance their causes.

- ISIS has used social media to mobilize tens of thousands of militants in over one hundred countries since its emergence.[6] Furthermore, ISIS used social media platforms to share its atrocities publicly, including executions.
- In December 2017, a video of a Hindu man killing a Nepali Muslim man in India went viral and ended in ethnic violence fueled by India's existing religious and ethnic divisions.[7]
- In Europe, rising nationalist groups connect and coordinate their anti-immigrant and xenophobic campaigns on social media platforms such as Facebook and Twitter.[8]
- The Myanmar military's genocidal campaigns against minority Rohingya were incited on Facebook platforms by dedicated officers misinforming the public to set a stage to justify the genocide.[9]

The above examples of the mobilization of masses for good and violent causes were made possible by social media platforms and leaders who knew how to effectively use them. Digital groups can raise funds through virtual networks and conduct financial transactions to support their causes. An individual or a small group can incite and incubate violent extremism, racism, anti-Semitism, and xenophobic ideas within a country and across international borders easily using social media platforms. Furthermore, social media platforms allow organized groups and individuals to have influence beyond their borders, making it difficult for any single actor to monitor or track when the actors have a violent and criminal intent.

Social media as a tool for atrocities monitoring and prevention

Understanding the nature and landscape of the digital and social media ecosystem, as well as the main actors and discourse on social media, is becoming more critical for the future of effective atrocities prevention. The massive amount of content uploaded to social media platforms such as Twitter, Facebook, TikTok, WhatsApp, Instagram, Telegram, YouTube, and others provides information essential to understanding the key issues, influential actors, motives, interests, and spoilers around any given atrocity-related context. As social media dynamics and users change, global online atrocities monitoring and prevention must become more complex. For example, traditional methods of analyzing speech struggle with the sheer magnitude of online speech. In addition, online monitoring for atrocity warning signs in multilingual contexts is challenged by the need to understand local dialects and interpretations. Prevention also becomes more difficult when authorities are implicated in encouraging atrocities on social media.

The need for strategic collaboration

The above trends confirm the importance of social media corporations and their advanced machine learning technologies in the future of atrocity monitoring and prevention efforts. Thus, there is a need for collaborative strategic partnerships between research institutions and private entities such as Facebook to engage AI and big data for timely, accurate early warning and preventive actions.

In recent years, private corporations—particularly Facebook—have faced increasing criticism from governments, inter-governmental institutions, and the public for being complicit in or having inadequate

policies to prevent violence and atrocities. This criticism influences users and negatively affects the profitability of these platforms, forcing the corporations to take a serious interest in preventing their platforms from becoming tools of violence incitement. Facebook updated its online community policies, which among others, include key atrocity monitoring components such as designating dangerous individuals and organizations and removing language that incites serious violence.[10] Facebook utilizes machine learning and a large number of local experts to moderate content in different languages and countries.

Despite these efforts, Facebook's mechanisms continue to fall short of effectively monitoring, identifying, and preventing violence and atrocity incitement. In 2018, the UN Independent Fact-Finding Mission on Myanmar found Facebook was used as a tool for organizing and inciting hate and violence against Rohingya and blamed its ineffectiveness and slow response for leading to real world violence.[11] Facebook's efforts, including its human content moderators, lack the necessary expertise and tools to analyze atrocities-related dynamics and design relevant preventive policies.

A social media company like Facebook could benefit greatly from a partnership with an institution with deep expertise on mass atrocities and genocide, such as the Museum's Simon-Skjodt Center. At the same time, a strategic collaboration could advance the Museum's global atrocities prevention efforts by increasing its ability to monitor atrocity warning signs on social media and understand the dynamics of online atrocities incitement, and by expanding the reach of its prevention messaging.

Several challenges might stand in the way of this type of partnership. For example, privacy concerns could prevent Facebook from sharing data with an institution like the Museum. Competitive interests could also prevent them from sharing technologies that determine how different types of content is spread across the platform. For its part, the Museum, as an independent establishment of the US government, might hesitate to partner with a corporation, let alone one that has come under significant scrutiny from the US Congress.

Mindful of these challenges, initially, a partnership should be narrowly defined and prioritize coordination on specific cases of urgency identified by the Early Warning Project. The partnership should seek to analyze atrocities-related dynamics and provide context-specific recommendations, including social media-specific policies. This type of initiative should be scaled up gradually to create a formal task force for mass atrocities and genocide prevention in social media.

Warning as Persuasion

Chiara De Franco, University of Southern Denmark

When discussing the role of warning in the prevention of mass atrocities, it is important to emphasize that warning is not a straightforward exercise in information transfer and that its outcome should not be just a list of warning indicators. Unfortunately, these are widespread misconceptions that have led to ineffective warning practices, which can be observed in many different institutional settings. Instead, warning about mass atrocities should be understood as *purposeful persuasion attempts—*possibly challenging long-standing beliefs about regions, countries or actors, age-old understandings of "national interest," and deeply rooted attitudes towards the use of diplomatic or military tools.

This means that warning "products" should be crafted and shared in ways that maximize their persuasiveness with specific recipients and in relation to specific goals. These goals may include adjusting beliefs about the probability of mass atrocities to occur, changing attitudes concerning the impact of such events on the recipients' values and interests, and fostering preventive or mitigating measures. Crucially, conceiving of warning as persuasion implies that a warning's success should be assessed in relation to the specific goal of the warning, not on the basis of ever-changing political ambitions.

Based on research I conducted between 2008 and 2018 on seven different cases of conflict escalation involving mass atrocities, I identified a few key factors as influencing the decision of which warnings will be listened to and acted upon and which will instead fall flat.[12]

First, persuasive warnings are characterized by *clarity* and *specificity,* not only in the way they identify and label the forecasted risk, but also in how they use evidence to support their claims and make the case for the risk at stake to have potential harmful consequences for the interests of the intended recipients. In the field of mass atrocity prevention, the choice of warning about genocide, ethnic cleansing, or just violence escalation is a critical one, which should be led by considerations of what kind of message the intended recipients need to receive to pay attention to the warning, accept it, and act upon it. Nevertheless, such a choice is often dictated by the working environment where warners operate. In certain organizational settings, for example, a legalistic culture might prevail and discourage the kind of bold and risky statements that a genocide warning requires even if and when it would be the most appropriate warning. Of course, this also means that to craft persuasive warnings, analysts need to develop confidence in their knowledge and motivation to take risks that might affect their career. This, in turn, might be more complicated at the beginning of a crisis when analysts can be less confident about their warnings, in cases of violence escalation exhibiting historical discontinuities, or where the conflict parties manage to operate in secret or be deceptive.

Second, even if the specific content of a warning certainly matters, the act of warning very much relies on a good *relationship* between warners and warning recipients. In many ways, warning is a people business and therefore it is affected by aspects of the relationship between warners and warning recipients such as closeness and trust. These factors are critical for the warners' ability to make key decisions about whom to warn, when, and how, as well as the warning recipients' capacity to evaluate the relevance of a warning. Recognizing the importance of relationships means that warners should invest in their network, track record, and reputation and that governments and international organizations (IOs) should equip themselves with instruments that allow warners and decision makers to meet regularly, talk openly, and develop a common language that can foster trust, mutual respect, and understanding. None of the organizations I studied, which include states such as the United Kingdom and the United States, and IOs such as the United Nations, the European Union, and the Organization for Security and Cooperation in Europe, have adequate mechanisms of this sort. On the contrary, they seem set up to reinforce a disconnect between warners and warning recipients. Take for example how many governments and IOs tend to emphasize "early" warnings as a sort of gold standard to strive for, thus forgetting that for a recipient, whose agenda may already be overloaded with current crises, early warnings can be too early to be valuable.

Third, a good *communication* strategy is critical to a warning's success. This does not necessarily mean that analysts should be skilled communicators, although some basic training in communication techniques would certainly not do them any harm. To be persuasive, warners need to develop awareness of the communication challenges arising from the evidence at their disposal and/or the prediction and recommendations they intend to put forward. These challenges are contextual in the sense that they should be assessed not in abstract terms but vis-à-vis the specific preferences and beliefs of the intended warning recipients, and in relation to specific organizational settings and political agendas. This often means that the most effective warners are seasoned officials with considerable material and social resources in an organization and significant practical knowledge about whom to get in touch with, when, and how to get a message successfully across.

Finally, it is worth stressing that warnings about armed conflict or mass atrocities are part and parcel of broader foreign policy debates and often just one piece of that incredibly complex puzzle that can be called the "swirl of intelligence." This means that their persuasiveness will also greatly depend on how congruent they are with other bits and pieces of information that policy makers receive from other sources. Thus, our understanding of the warning process should include both warnings that travel "inside-up" within a bureaucracy as well as warnings that percolate from the "outside-in." Journalists and NGOs officials, in particular, are as much in the warning business as analysts working for governments or IOs and their influence on how far a warning manages to travel can be significant. This is especially true in the early stages of a crisis when journalists and NGO officials might have stronger incentive to ring the alarm bell and feel less constrained by reputational risks, but also in later stages when government officials might resort to leaking reports or making public statements to finally achieve what they previously attempted through internal channels.

Early-warning Systems for Political Violence: How Should We Define our Prediction Targets?

Håvard Hegre, Uppsala University and the Peace Research Institute Oslo (PRIO)

As noted in the UN-World Bank *Pathways for Peace* report, early warning of armed conflict is necessary to facilitate early action. Stakeholders seeking to prevent a humanitarian disaster or preparing to assist will be better able to do so if they have an early indication of where major violence is likely to occur, as well as the likely intensity and duration of fighting if it occurs.

A number of quantitative early warning systems have recently been developed and are running with regular updates of their risk assessments. All of these systems are limited to binary (yes/no) forecasts of whether violence will exceed a given threshold of violence or not. EWP, for instance, generates forecasts for whether the number of persons killed in mass atrocities will exceed 1,000 in a period of one year or less. EWP also limits its forecasts to *new* mass killings over the next two years. The Violence Early-Warning System (ViEWS) project is limited to binary predictions, but also covers continued conflict. It currently predicts whether at least 25 people will be killed in a given country in a month.[13] (for details, see this article).

Although highly useful, have these projects chosen the most useful target for their predictions? Below, I suggest some of the alternatives available to modelers and discuss their relative benefits. I will concentrate on forecasts at the country level, but similar conclusions apply to forecasts at a more detailed geographical resolution. The discussion applies equally well to projects primarily interested in violence against civilians as to those forecasting confrontations between two armies. In the discussion, I will focus on the number of documentable direct deaths from violent actions as a measure of conflict intensity, but a similar logic holds for the number of indirect deaths or the number of people forcibly displaced.

When defining a prediction target, choices along four dimensions are particularly important for modelers:

1. *Dichotomous* v. *continuous*-variable forecasts: Whether the intensity of the conflict will exceed a threshold or not over a given time period v. seeking to estimate the likely fatality count over the period.
2. *Temporal resolution*: Whether the most useful time period for aggregating fatalities is a calendar year, a month, a week, or even a day.
3. *Single-period* v. *multiple-period* future: Given the time period, whether the forecasts cover only the next month or year, or a series of consecutive future periods.
4. Whether to forecast *onsets of conflict* only in places where there is no ongoing violence, or *incidence of conflict* also including the likelihood of a continuation of violence.

The most useful prediction target, in my mind, is to forecast the likely number of fatalities in each of a series of consecutive, rather short time periods (e.g., months), including both new and ongoing bouts of violence. I will refer to this target as 'continuous multi-period incidence' below, or CMPI for short. Despite the fact that no existing early warning systems define their targets in this way, there are several reasons for why CMPI is better than the alternatives.

First, CMPI is a **more general target**. When first generated, such forecasts can be easily recoded to any of the simpler targets. Forecasts of a gradual escalation of violence up to a given threshold is a forecast of an onset, and a forecast of a total of 1,000 deaths over a 12-month period is a forecast of major conflict. One forecasting model, then, can cover multiple purposes and user preferences.

Second, CMPI **avoids arbitrary coding decisions**. A typical threshold is 1,000 deaths per calendar year. For state-based conflict in the year 2020, this threshold just about includes fighting in the Tigray province in Ethiopia as well as in Northern Nigeria, but excludes violence in both Mali and Burkina Faso. Lowering the threshold to 500 deaths per year would include these, but still exclude fighting in Cameroon. Are these conflicts really fundamentally different from each other? Lowering the threshold even further, say to 25 deaths per year, means including lots of cases that are currently relatively quiet, such as Algeria, South Democratic Republic of the Congo, and Sudan, risking that the early warning system is watered down and unable to distinguish between small crises and major disasters. Similarly, coding 'onset' of conflict or mass killing adds another set of arbitrary decisions. If we define onset as a year with at least 1,000 deaths after at least one year with less than 1,000, Mali and Burkina Faso are at risk of conflict onset just as much as Ghana and Mauritania. A model forecasting a CMPI target that seeks to predict the number of fatalities in a given future time period irrespective of what happened in the most recent period avoids all these arbitrary decisions.

Third, CMPI makes **better use of available information**. Typically, prediction models are algorithms that relate past observations of the prediction target to past risk indicators. When defining the target dichotomously with a target of 1,000 deaths per year, the model cannot use any characteristics that separate the obviously volatile cases of Mali and Burkina Faso from the much calmer cases of Algeria and Namibia to improve the model. With a continuous target, all the information about past violence feeds into the model and improves accuracy. When including information also on ongoing violence, an even larger amount of information is available to train the models.

Fourth, by using a short time period such as the month as the time unit, a CMPI model is **better able to update quickly** using new information as it becomes available. Conflict data projects such as the Uppsala Conflict Data Program and Armed Conflict Location & Event Data Project (ACLED) release updates monthly or even daily, and early warning systems that keep pace with these updates are much better placed to provide timely warnings.

Fifth, using multiple CMPI gives a **better representation of the total impact of a crisis over time**. Even though conflicts can escalate swiftly, political violence tends to continue for multiple years when first set in motion. When using CMPI as a prediction target, early warning systems are able to both provide timely warnings for the next couple of months as well as estimating the total amount of violence aggregated over the next few years.

Finally, a CMPI target allows a **more nuanced handling of uncertainty**. Good early warning systems support their models by test runs where the model generates predictions for historical data that were not used to train the model. With a dichotomous prediction target, the only uncertainty metrics that can be derived from such test runs show the extent to which a dichotomous prediction coincides with the dichotomous classification of the historical observations. When predicting the expected number of fatalities, it becomes possible to also assess the magnitude of prediction errors. In addition to providing information to users about the performance of the model, it also gives model developers much more powerful tools to improve their systems. For instance, consider two models that forecast 200 and 900 deaths in a country in a month, respectively. With a dichotomous prediction target evaluation, evaluation metrics are unable to show that the latter model came much closer to a threshold of 1,000 deaths than the former. With a continuous target, evaluation metrics can easily show the superiority of the latter, and, provided this model does well also in a majority of the cases, modelers have tools to select better performing models. Moreover, CMPI early warning systems do not have to settle with a point prediction but can also generate a probability distribution for the number of fatalities—e.g., in addition to a 'best' estimate of say 600 deaths, present the prediction as a five percent probability of more than 1,000 deaths, 50 percent of more than 500, and 95 percent of more than 100 in a given time period. (In principle, such a probability distribution can also be generated for dichotomous targets, but these are less intuitive than the continuous-target counterparts.)

All in all, these advantages strongly point to moving towards CMPI prediction targets. From the user side, there are hardly any downsides, since a CMPI prediction can be simplified to any metric preferred by the user. From the modeler side, CMPI targets are considerably more complicated to set up and maintain. It is far from being unfeasible, though—the ViEWS project has run a model with a dichotomous multi-month incidence since 2018 and is currently developing an extension to forecast the expected number of fatalities in a given country month.

Researchers' Dilemmas in Getting Stakeholders to Prevent Mass Atrocities

Victor O. Okorie, Obafemi Awolowo University

Introduction

The 2021 Sudikoff Interdisciplinary Seminar Series on Genocide Prevention identified several ways to reduce the warning-response gap. The identified pathways range from shedding more light on genocide triggers to developing a common language to ensure that warnings are not lost in translation.[16] The rationale of these suggestions is that a higher degree of clarity and a common tongue would easily get critical stakeholders to act promptly to prevent mass atrocities. A clinical precision in pinpointing and communicating the drivers of an imminent genocide would certainly get apolitical stakeholders to act at the speed of light.

However, in settings where economic and political interests of some crucial stakeholders weigh on the side of continuing atrocities, those researching and reporting on such violence and risk factors are especially challenged. These settings are a Gordian knot. They spawn epistemological, empirical, and ethical entanglements that demand researchers embrace multiple roles. The roles may include, but are not limited to, researcher, ethicist, activist, writer, citizen, presenter, lobbyist, facilitator, advocate, trainer, and media representative. This essay reflects on some of the roles and contexts that present ethical quandaries for researchers interrogating warning signs for genocide.

Researchers: Loyalty, legitimacy, legality, and reality

Researchers are members of different countries, communities, and groups. They have obligations and loyalties to their various groups, cultures, countries, and wider society. Their countries and epistemic communities, respectively, may influence whether a given warning about genocide is perceived to be legal and legitimate. Countries and institutions must give ethical clearance for the collection of data for forecasting genocide to be ethically and epistemologically correct. These gatekeepers, just as research subjects, have the right to withdraw or revoke their consent for ongoing research. The gatekeepers may create ethical logjams, including contexts where individuals are willing to provide evidence but the institutions are unwilling to give ethical approval. This scenario is common in countries wherein precursors of genocide are hidden in the garb of repressive regime, pseudo-democracy, and successionists' agitations.

The quandaries are more complicated in contexts where the involved countries or their alliances are the potential perpetrators of the looming or ongoing genocide. For instance, at a certain stage in the pogrom that almost exterminated the entire Igbo people of Nigeria, the British government sent a team to the region for an on-spot assessment. Given Britain's clandestine interest in the war of aggression, the team

arrived in Nigeria, but stayed in Lagos and prepared its report without stepping into Igbo land. The loyalty of the team members to their country produced an overtly skewed report, which justified blockade policies and anti-humanitarian posture of the Nigerian government and enabled the mass atrocities in Igbo land. In contrast to the report, however, independent observers, like Leo Kuper, refer to the Nigerian government's massacre and systematic starvation of the Biafran secessionist Igbo during the Nigerian civil war as genocide.[17] Kuper's views resonate with those of other scholars such as Achebe, Adichie, and Soyinka.

However, there are a few scholars whose perceptions partly aligned with the fabricated report.[18] The latter category of scholars was, perhaps, driven by loyalty for their country as well as the fear of being punished by the state. For instance, Wole Soyinka, the Nobel laureate, was incarcerated during the pogrom because of his anti-genocide views. Similarly, the risk of being charged with a cultural crime prevented Shahram Khosarvi from giving an audible voice to "the women who cover" in Iran.[19] Researchers focusing on mass atrocities in several contexts are increasingly facing conditions that demand loyalty to the state and legality in-country at the expense of collecting and publishing quality data for predicting mass atrocities, especially in the Global South. Stakeholders who may be interested in atrocity prevention may distrust data collected under such conditions, limiting their willingness to act. In many cases, researchers may take on the role of activists to break the deafening silence of governments in the face of looming genocide. The activist role is not without its own ethical challenges, which are highlighted in the next section.

Researchers and activism: Breaking the deafening silence

The Yiddish title of Elie Wiesel's Holocaust memoir is [Un di Velt Hot Geshvign](#): "And the World Has Remained Silent." "The World Was Silent When We Died" is a chapter in Chimamanda Adichie's *Half of a Yellow Sun*. The mass atrocities that these titles describe clearly illustrate the outcome of stakeholders' silence in the face of looming or ongoing genocide. Activist research may fill that silence, but at the same time presents a new set of dilemmas to researchers. Activist research is variously defined. Hale describes it as "a practice that helps one to better understand the causes of inequality, oppression, violence, et cetera; [it] is carried out, at each phase, in direct cooperation with an organized collective of people who themselves are subject to these conditions; and [it] is used together with the people in question to transform these conditions."[20] For Speed, it is the "overt commitment to an engagement with our research subjects that is directed toward a shared political goal."[21] This form of research presents an array of dilemmas that invokes debate about the responsibilities of researchers to the people they study.

Researchers are ethically obligated to try to understand and represent the perspectives of perpetrators, victims, and witnesses of impending mass atrocities. The dilemma of researchers is striking a delicate balance among the three voices. Nordstrom suggests that the voice of victims should be amplified above others.[21] Yet, in the context of predicting and preventing genocide, it is imperative to note that subjectivities of potential victims are no more important than those of potential perpetrators. This is because accounts of looming mass atrocities framed solely around the narratives of victims may produce many false positives, while those based on the narratives of perpetrators may lead to many false negatives. As such, researchers focused on early warning should not only give ethnographic voice to potential victims, but also show "the ways in which perpetrators might be considered as cultural agents."[23] This implies that to understand unfolding or future atrocities, researchers should explore the broader

cultural values—and the links between those values and violence—of both potential victims and potential perpetrators. This is the foundation of objectivity in ethnographic reports.

Additionally, researchers are inevitable witnesses of impending genocide. The witness role is no less ethically challenging. Researchers through knowledge production are indisputable 'seers' and eyewitnesses of impending atrocities. Genocide, for researchers forecasting mass atrocities, is neither a word nor a figure. It is people who have names and addresses, and are about to be premeditatedly murdered by other humans. It is human faces at the verge of premeditated violent death. The horrific images of impending premeditated violent deaths may live forever in the researchers' organic memories and continue to haunt their dreams. The researchers, thus, bear witness. To be true witnesses, however, they must collect, analyze, and publish evidence with utmost objectivity and reflexivity. Objectivity and self-reflexivity in ethnographic accounts would keep the feet of researchers within the ethical boundaries. Through objectivity researchers fulfil one of the cardinal requirements of science while self-reflexivity allows stakeholders to know the extent to which a given forecast has been produced after the imagination of the researchers, for every work is a portrait of the doer. Thus, objectivity and self-reflexivity in forecasting reports will contribute to making the stakeholders act at the speed of light.

Early Warning of What, For What, and Whom?

Clionadh Raleigh, University of Sussex and ACLED

I was recently on a panel about early warning and conflict trends, and one of the participants was eager to tell the audience that technologies are being developed so quickly that we will soon have excellent ways to predict the exact locations of conflict in the future, and therefore to quickly respond. But this exuberance is unfounded: we know where conflicts are now and it does not seem to make a huge difference to whether they generate attention, let alone a response.

This disconnect illustrates a critical point about the science of early conflict warning and its usability. The working hypothesis of many early warning and prediction models is that this information is necessary for practitioners and governments to know in order for these same communities to engage in short- and long-term conflict prevention. The central issue for the modelers of early warning is to be accurate enough so that its audience may rely on the information. As a scholarly community, we have spent several years debating the metrics and measures of early warning systems in order to achieve higher accuracy. But what if we have labored under the wrong assumptions in our approach to early warning?

This essay makes two points and two recommendations. The first point is that early warning models do not reflect what modern political violence actually looks like. The models use crude, macro definitions of political violence to be explained by similarly aggregated structural variables at national levels. Pursuing a global model exacerbates these problems as conflict takes on very different modalities in response to domestic political environments. This means that these models miss key insights about political violence, including local drivers, volatility in time, and variation across space of political violence. In short, part of the reason that predicting the future of political violence is so difficult is that we have inaccurately described the present and sought a consistency in explanations where none exists.

The second point is that the modelling methodologies currently popular within early warning studies require a scale and focus that limits their actual usability to those who require this information for early action. High intensity events, and those whose parameters are defined by academic theories rather than on-the-ground actions, are rare and difficult to model under the best of circumstances. Further, responders are normally more concerned with low level threats and their effect on humanitarian response, access, and diffusion. Models are also unable to integrate timely, often qualitative, information into predictions. And yet, subtle and slight variations in the political and social environment are precisely what lead to the forms of violence that early responders need to know about.

For these two reasons, I recommend the following changes in our approach and debates: responders are better served with more detailed studies of specific conflict environments and circumstances. These should be supplemented with understandings of national and subnational government resilience and

authority structures. Information on the number, form, and agendas of active violent groups and their coordination and fragmentation should be prioritized. Further, as an alternative to complicated early warning and prediction models, researchers can provide more accurate and usable information by analyzing the parameters of conflict intensity in combination with its violence volatility. This second recommendation reflects a reality about the information needed in unstable environments, rather than using models that test conflict theories and hoping they are useful in practice, despite all evidence to the contrary.

Which conflicts?

Conflict emerges and grows from local contexts, and it adapts to every environment. At the event level, that conflicts are happening in one location and time but not others is part preparation, part location draw, part opportunity, and part random bad luck. Events may be delayed because of a lack of ammunition, combatant injury, a wrong turn, a peace offering, the weather, or many other factors. At the conflict level, violence can emerge in communities over mistakes, elite competition generated in the capital but fought in the peripheries, local authority contests, preparation for larger violence, criminality with larger ambitions, and many other reasons. The emergence and the patterns, the groups, and the motivations, are not automated or formulaic. There is no conflict 'system' which has rules that apply homogeneously.

The characteristics of conflicts, the groups involved, the focus of competition, and the capacity for violence by each group can change. What is competed over (territory, authority, control, resources, positions) and the scales on which it can occur (households, communities, groups, offices, regions, states, borders) vary and multiply. Conflict adapts to every political circumstance: it becomes institutionalized if the state is the strongest entity and has a monopoly on its means; it becomes fractured if subnational competition is encouraged by national level governments; it takes on new branding to maximize support and recruits (e.g., some new Islamist groups); it assumes forms that allow for maximum survival and benefits (e.g., cartels in some environments; rebels in others); violent groups are kept on retainer by elites who require 'armed labor'; and multiple conflicts occur in many places within the same state for different reasons.

No early warning system can reflect these complexities. Conflict is not an aberration or a failure of government relationships and institutions, development levels, or the physical environments; it is a reflection and adaptation to those systems. Conflict models must begin with these assumptions and try to model what happens in environments where they are aware of the terms, scale, and agents of competition. Current early warning models often suggest a consistency in causality, and a reluctance to address scales of violence and their dimensions.

Are models capable of capturing reality?

None of the characteristics of conflict noted above bode well for a generalizable model. For responders and governments, the risks that they are primarily exposed to are low level and often nonlethal. But most existing models use high intensity events to predict and use the previous occurrence of those same events over the longer-term period (of possibly over fifty years in the past) to assess the future. The scale and focus on models is out of sync with the needs of those requiring early warning.

Prediction and forecasting models provide early warnings for a range of conflict and violence types, including genocide and/or politicide[24], coups[25], election violence[26], and conflict or political violence.[27]

Most projects produce lists or assessments, ranking which countries are at the greatest risk of conflict in the future. Most projects have near universal (state-based) coverage. Even at the subnational level, these projects at their best can currently tell us where conflict continues, as their usability is curtailed by the limitations of general information, substituting counts of events or fatalities per state as indicators of instability. It is less helpful for people to know that Somalia will be violent tomorrow, or that Botswana likely will not be. It is useful to know what form and type of variability is in violent environments.

If early warning models promise anything, it is a generalizable formula for conflict; therefore, specifics are of limited use. But better prediction is only possible through improved understanding of a conflict environment, and its possible trajectories. This is achieved by knowing more about the political dynamics of the country, the agendas and alignments between government and non-state agents, what is worth 'competing over,' and how that competition is likely to play out. It is not a homogenous model, it is knowledge about domestic politics.

What would be helpful for practitioners, governments, and communities?

More detailed studies of individual conflicts and circumstances, national and subnational government resilience and authority structures, and violent group fragmentation. In reality, governments, communities, and organizations need to assess immediate risk in order to make decisions in real time. Responses to violence warnings are very rarely to prevent violence or intervene early, but to prepare and plan for the dimensions of a specific conflict. This is the reality of working in unstable environments. Mitigating violence and peacebuilding at anything higher than specific communities is not something that many are capable of, nor are these the goals of organizations or governments. This suggests that usable information must be localized, specific, and problem oriented, not generalized, nationalized, or projected into the future.

So, what specifically should they be warned about?

The risk of exposure to political violence is made up of the continuity of conflict intensity plus the volatility of that violence. This creates a four-way classification system that designates the range of all political violence and its consistency (high and low event levels, high and low volatility). This information can suggest whether programming can continue, whether roads will be accessible, and whether an organization needs more or less security staff. It is usable by those who need it, which is the most important result of any warning.

The Psychology of Risk: Implications for Communicating and Acting Upon Early Warnings and Ongoing Atrocities

Paul Slovic, Decision Research and University of Oregon

The aim of communicating risk of mass atrocities through an early warning is to motivate individuals and their governments to stifle the precursors of violence while they are still in an early stage. But this is easier said than done. Over the course of the past century, national decision makers often have been informed of imminent or ongoing genocides and mass atrocities yet have chosen not to intervene. As Samantha Power observed, America's record is one of strong abstract support for principles and ideals opposed to genocide but of little or no action when a real-world situation arises that calls for immediate, effective intervention to prevent a looming genocide or halt one that is ongoing.[28] If we do not act when we see abuses happening right now, imagine how hard it may be to motivate action in advance, when there is no visible harm—only warning signs that point to the possibility, not the reality, of catastrophe. Climate change and COVID-19 provide vivid examples of the dire consequences that result from failure to heed early warnings.[29]

It is important to heed warning signs and intervene early as it is more difficult to stop mass violence once it begins. But to inform and motivate potential actors effectively, risk communication must overcome psychological obstacles to intervention in ongoing genocides that have been identified by research.[30] It is likely that these psychological mechanisms would challenge the effectiveness of early warning as well. What follows is a brief overview of these obstacles and their implications for communication. References to this research can be found in two collections.[31]

Risk communication

Communication about risk has been essential to survival since humans lived in caves and it has been studied scientifically for more than forty years, in close association with research on risk perception. Early studies showed that the public's judgments of risk often differed greatly from judgments by technical experts, who attributed these differences to public ignorance and irrationality. Experts aimed to set the public straight through education and what came to be known as risk communication.[32] However, subsequent studies painted a different picture, showing that the public had a rich conception of risk that took into account important qualitative factors such as catastrophic potential, trust in authorities, risk to future generations, etc., that were typically excluded from experts' risk assessments.[33] Risk communication, which originally consisted of experts lecturing laypeople, was restructured to recommend two-way dialogue which respected the important value-laden qualities of risk that were important to people.[34]

Recently, research on risk perception has aimed to understand why people and their governments, who place a high value on protecting individual lives, so often turn a blind eye to protecting large numbers of

people threatened by genocides and mass atrocities. Three cognitive limitations have been identified and shown to pose severe obstacles to effective response to atrocities. They are psychic numbing, a sense of inefficacy, and a decision making bias known as the prominence effect.

Psychic numbing

Consider two questions. First, how **should** we value the protection of human lives? And second, how **do** we value the protection of human lives?

Regarding the first question, if we believe that every life is essentially equally important, then the value of protecting those lives should increase in a straight line as the number of lives at risk increases, as shown in Figure 1a. This is a simple process of addition. When additional losses of life threaten the extinction of a people, as in the case of genocide, the very next life at risk is even more valuable to protect than the life before it, causing the value line to curve upward as in Figure 1b.
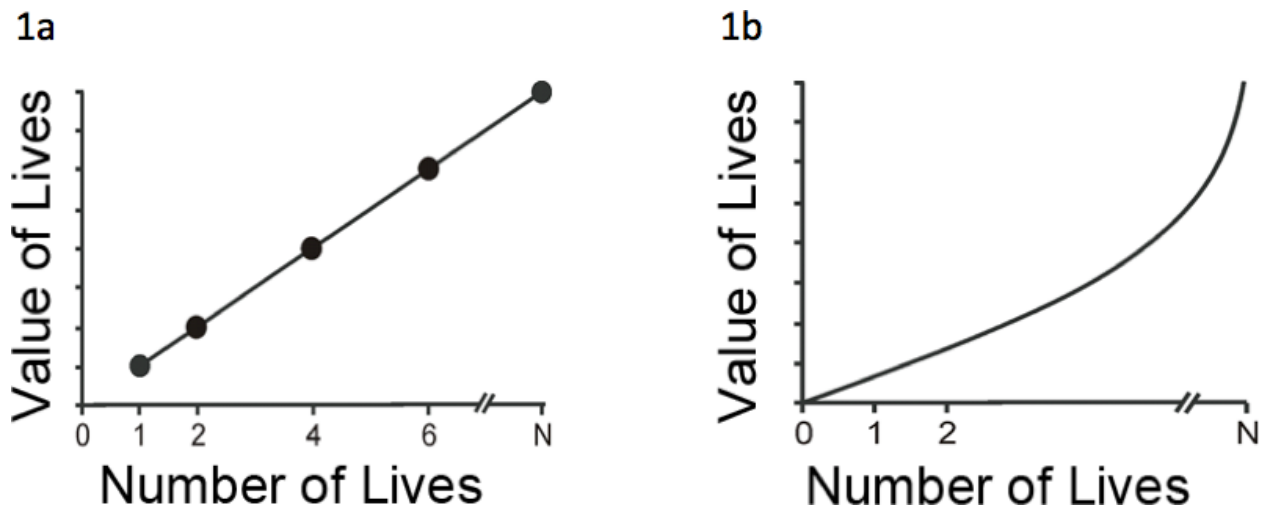

Figure 1. Two normative models for valuing lives as the number at risk increases. Adapted from Paul Slovic, "'If I Look at the Mass I Will Never Act': Psychic Numbing and Genocide," *Judgment and Decision Making* 2 (2007): 79–95, Figs. 2 & 3.

Figure 2 illustrates how people **do** tend to value increasing numbers of people at risk, based on research. The value curves reflect the fact that intuitive judgments and feelings often override the analytic judgments that underlie Figures 1a and 1b. Our feelings often follow an underlying *arithmetic of compassion*[35] that is nonrational and deadly.
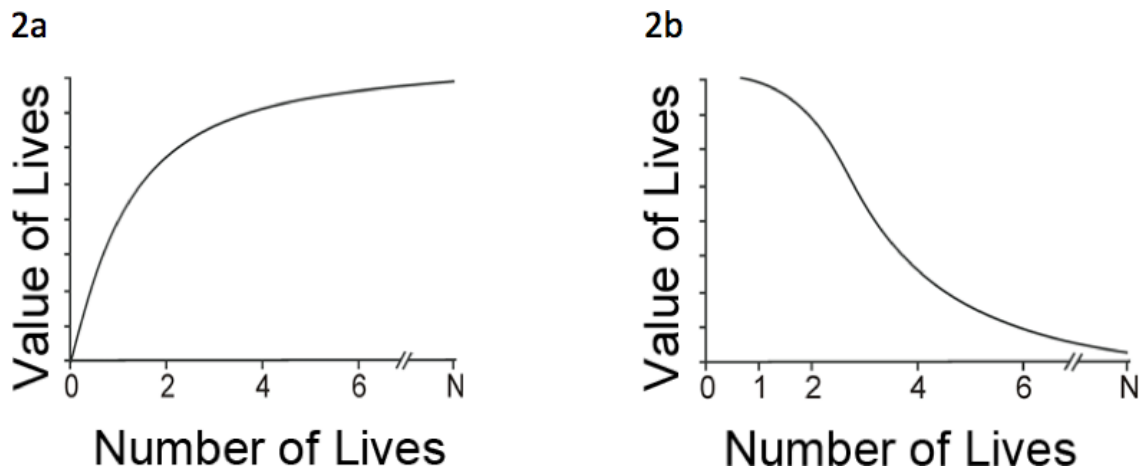
Figure 2. Two normative models for valuing lives as the number at risk increases. Adapted from Paul Slovic, "'If I Look at the Mass I Will Never Act': Psychic Numbing and Genocide," *Judgment and Decision Making* 2 (2007): 79–95, Figs. 2 & 3.

Figure 2a shows that the biggest change in value occurs with the first life, going from zero to one. On an emotional level, we care greatly about protecting single lives, something known as "the singularity effect."[36] But as the numbers increase, "psychic numbing" begins to desensitize us. Two lives do not feel twice as valuable to protect as one life. In fact, as the number of lives at risk increases, the additional lives quickly seem to add less and less value and the curve flattens. This curve in Figure 2a also implies that a life that is so valuable to protect if it is the first or only life at risk loses its value when it is part of a larger tragedy, with more lives endangered. Figure 2b is even more problematic. It implies that the value of lives does not always increase monotonically with magnitude as in 2a, but sometimes decreases, perhaps collapsing towards zero when many lives are at stake. Both ways of valuing life shown in Figure 2 enable mass killing. Numerous laboratory experiments have demonstrated psychic numbing and compassion collapse.

Pseudoinefficacy

We help others not only because they need our help but because we get a good feeling, sort of a warm glow, when we help them. The trouble is that it does not feel as good to help someone when our attention is drawn to the fact that there are others whom we are not able to help. In such situations, research has shown that many people then simply give up and do not even help those they **can** help, a phenomenon labeled *pseudoinefficacy*.[37] This is wrong! Just because we cannot fix a problem completely does not mean we should do nothing.

The prominence effect

Decision makers typically find choosing between actions that involve competing values and objectives quite difficult. Research has demonstrated that decision makers often resolve such conflicts by choosing in a way that satisfies certain "prominent objectives" that are inherently more defensible even when these objectives are not inherently more important.

Slovic proposed the prominence of political and national security objectives, obviously highly defensible, over humanitarian lifesaving (less defensible) as an explanation for the reluctance of the US government to intervene to prevent genocides and mass atrocities in Syria and elsewhere.[38] For example, early in the Syrian war, while acknowledging "very real and legitimate" humanitarian interests in Syria—some 80,000 people had already been killed, and millions had lost their homes—President Obama said his "bottom line" had to be "what's in the best interest of America's security."[39]

Implications for risk communication and action

Here are a few suggestions based on the research described above.

For those who are potential or actual victims of atrocities, if you can, let the world know you as individuals! Communicate your stories so others may appreciate your hopes and your dreams, your struggles and your successes. It is important that the world not think of you as statistics.

For government officials and members of NGOs, do what you can to enlist and support journalists, writers, filmmakers, and social media experts to describe the warnings and the stories of abuse and atrocities to the outside world.

For citizens, who care greatly about the protection of individuals but feel powerless to stop mass atrocities, do not let psychic numbing stop you from doing whatever you can do. When you see statistics of atrocities, try to imagine the lives of some of the individuals represented by the numbers. As Holocaust survivor Abel Herzberg said, "There were not six million Jews murdered; there was one murder, six million times." And do not succumb to a feeling of powerlessness just because you cannot help everyone. Appreciate the fact that even partial successes resulting from warnings and other interventions can save whole lives. And amplify your ability to make a difference by joining and supporting NGOs that are dedicated to humanitarian causes and are doing heroic work.

Finally, for those in positions of power, recognize the limitations of existing laws and institutions that have repeatedly failed in their promise to ensure that "never again" would genocides and mass atrocities be allowed to happen. Think creatively to design policies and international laws that will compel governments to protect large numbers of endangered people with a degree of intensity that respects the great importance placed on protecting individual lives.

About the Authors

Quscondy Abdulshafi is a program specialist in democracy, human rights, and governance with over nine years of experience working in the East and Horn of Africa. He was the founding member of the Darfur Student Movement of 2003, a student organization against the genocide and ethnic cleansing in Darfur, and the winner of the Civil Society Leadership Award from the Open Society Foundations in 2016. Abdulshafi is now a Regional Advisor at Freedom House's Africa Program, and a US board member for Peace Direct. The opinions here are of Quscondy Abdulshafi and not of his affiliate institution.

Chiara De Franco is an associate professor of international relations at the University of Southern Denmark and co-director/director of Research of the Centre for War Studies. De Franco's research—broadly conceived—builds on constructivist and practice theoretical approaches to international relations and international political communication. She is the principal investigator of Protection Complexity, an international research program funded by the Danish Research Council. The project investigates how the EU, UN, and African Union practice protection of civilians in Mali, Central African Republic, and Somalia. For a full project description and list of collaborators visit www.protexproject.eu. De Franco is currently also working towards the development of a new narrative approach to international practices. De Franco is the author of *Warning about War* (winner of the 2021 International Studies Association [ISA] Annual Best Book Award and of the 2021 ISA/ICOMM 2020 Annual Book Award), *Media Power and the Transformation of War*, and several academic articles, and co-editor of *Ending War: a dialogue across disciplines* and *Forecasting, Warning and Transnational Risks*.

Victor O. Okorie holds a joint PhD in development and anthropology, as well as a MA in cultural anthropology from the University of Wisconsin-Madison. He also holds a Master of Philosophy in agricultural extension and rural sociology and a Bachelor of Agriculture from Obafemi Awolowo University, Ile-Ife, Nigeria, where he currently is a senior lecturer for both undergraduate and graduate courses. He is a recent recipient of an African Peacebuilding Network individual research grant (2014) sponsored by the Social Science Research Council, and a former fellow of the World Council of Social Science (2013-2014). Okorie's current research activities interrogate how embodied risks and uncertainty associated with disasters manifest medically as emotion-related illnesses such as pre-disaster anxiety disorders, post-traumatic stress disorder, depression, and anxiety among vulnerable populace in West Africa. He also explores the nexus between various methods of knowledge production and public uptake of scientific information in the context of disasters and risk reduction.

Clionadh Raleigh is a professor of political geography and conflict at the University of Sussex and the executive director of ACLED. Raleigh founded ACLED in 2005 as part of her PhD work, and it is now an independent NGO with 200 staff collecting and analyzing information on political violence and demonstrations across the globe. Her research concentrates on the dynamics of conflict and violence, African political environments, and elite networks.

Håvard Hegre is a Dag Hammarskjöld Professor of Peace and Conflict Research at Uppsala University and research professor at PRIO. Hegre's research interests revolve around how democratic institutions can prevent armed conflict, and how the relations between institutions and incentives for the use of organized violence are contingent on socioeconomic changes such as poverty reduction, economic diversification, and education. He is also interested in forecasting political events such as armed conflict.

Paul Slovic studies judgment and decision processes with an emphasis on decision making under conditions of risk. His work examines fundamental issues such as the influence of affect on judgments and decisions. He also studies the factors that underlie perceptions of risk and attempts to assess the importance of these perceptions for the management of risk in society. His most recent research examines psychological factors contributing to apathy toward genocide, politicized violence, and decision making pertaining to nuclear war. He no longer does classroom teaching but does advise students in their research. For further information visit his websites: www.decisionresearch.org and arithmeticofcompassion.org.

Mollie Zapata was a researcher with the Simon-Skjodt Center for the Prevention of Genocide at the United States Holocaust Memorial Museum from 2017-2021. She was responsible for conducting quantitative and qualitative research on a wide range of issues related to genocide prevention, with a primary focus on the Center's Early Warning Project. Previously Mollie had been working in natural language processing analytics at Protagonist, an international affairs consulting firm, for various government and foundation clients. Her other work experience includes the US Institute for Peace, the Institute for Inclusive Security, the Enough Project, and the National Geographic Society. Mollie holds an MA from The Fletcher School at Tufts University and a BA in international affairs from Boston University.

Notes

[1] Seminar attendees included scholars and practitioners who have worked on early warning and forecasting projects for the African Union, the ViEWS, the Dallaire Institute, the European External Action Service, ACLED, and multiple academic institutions.

[2] "Peoples under threat 2019," Ceasefire Centre for Civilian Rights and Minority Rights Group International, 2019, https://reliefweb.int/report/world/peoples-under-threat-2019-role-social-media-exacerbating-violence.

[3] Thessa Lageman, "Remembering Mohamed Bouazizi: The man who sparked the Arab Spring," *Al Jazeera*, December 17, 2020, https://www.aljazeera.com/features/2020/12/17/remembering-mohamed-bouazizi-his-death-triggered-the-arab.

[4] Jon Emont, "'Facebook Protesters' Helped Sudan Drive Out Bashir," *The Wall Street Journal*, April 12, 2019, https://www.wsj.com/articles/social-media-movements-helped-drive-out-african-leaders-11555085194; "Sudan's Omar al-Bashir mocks 'Facebook protesters'," *BBC*, January 31, 2019, https://www.bbc.com/news/world-africa-47075513.

[5] Sarah Marsh, "The Facebook group that staged first in Cuba's wave of protests," *Reuters*, August 9, 2021, https://www.reuters.com/world/americas/facebook-group-that-staged-first-cubas-wave-protests-2021-08-09/.

[6] "IS foreign fighters: 5,600 have returned home - report," *BBC*, October 24, 2017, https://www.bbc.com/news/world-middle-east-41734069.

[7] Stephanie Nolen, "India's ethnic clashes intensify within social-media maelstrom," *The Globe and Mail*, August 23, 2012, https://www.theglobeandmail.com/news/world/indias-ethnic-clashes-intensify-within-social-media-maelstrom/article4496392/.

[8] Sabina Mihelj and César Jiménez-Martínez, "Digital nationalism: Understanding the role of digital media in the rise of 'new' nationalism," *Nations and Nationalism* 27, no. 2 (2021): 331-346).

[9] Paul Mozur, "A Genocide Incited on Facebook, With Posts From Myanmar's Military," *The New York Times*, October 15, 2018, https://www.nytimes.com/2018/10/15/technology/myanmar-facebook-genocide.html.

[10] "Community Standards," Facebook, https://www.facebook.com/communitystandards/introduction.

[1] United Nations, Human Rights Council, *Report of the independent international fact-finding mission on Myanmar* (New York, NY: UN Headquarters, 2018).

[2] This research has been conducted together with Christoph O. Meyer (King's College London) and Florian Otto (Control Risk). The main findings are discussed in some length in *Warning about War: Conflict, Persuasion and Foreign Policy* (Cambridge: Cambridge University Press, 2019).

[3] For details, see Håvard Hegre et al., "ViEWS2020: Revising and evaluating the ViEWS political Violence Early-Warning System," *Journal of Peace Research* 58, no. 3 (2021): 599-611.

[4] Fatality numbers are from https://ucdp.uu.se/exploratory, filtered to the year 2020 and state-based violence.

[5] Ibid.

[6] Tahlia Mullen, "Taking stock of early warning for atrocity prevention: 2021 Sudikoff Interdisciplinary Seminar Series on Genocide Prevention Communicating Risk Rapporteur's Report," US Holocaust Memorial Museum, 2021, 2.

[7] Leo Kuper, *Genocide: Its Political Use in the Twentieth Century* (New Haven: Yale University Press, 1981), 45.

[8] D. F. Atidoga and Abubakar Ishaq, "Was Genocide Committed against the Igbo Nation of South Eastern Nigeria during the Civil War? The Law of Genocide on Trial," *GLR*, no. 1 (2018): 183.

[9] Shahram Khosravi, *Young and Defiant in Tehran* (Philadelphia: University of Pennsylvania Press, 2011), 34.

[20] Charles R. Hale, "What is Activist Research?," *Social Science Research Council* 2, no. 1-2 (2001): 14.

[21] Shannon Speed, "At the Crossroads of Human Rights and Anthropology: Toward a Critically Engaged Activist Research," *American Anthropologist* 108, no. 1 (2006): 71.

[22] Carolyn Nordstrom, *Shadows of War: Violence, Power, and International Profiteering in the Twenty-first Century*, vol. 10 (Berkeley: University of California Press, 2004), 123.

[23] Neil L. Whitehead, "Introduction: Humanistic approaches to violence," *Anthropology and Humanism* 34, no. 1 (2009): 6.

[24] Jay Ulfelder, "A multimodel ensemble for forecasting onsets of state-sponsored mass killing," in APSA 2013 Annual Meeting Paper; Benjamin E. Goldsmith and Charles Butcher, "Genocide forecasting: Past accuracy and new forecasts to 2020," *Journal of Genocide Research* 20, no. 1 (2018): 90-107.

[25] "CoupCast - Political Forecasting," One Earth Future, available at: https://oneearthfuture.org/activities/coup-cast.

[26] Clayton Besaw and Matthew Frank, "ELection VIolence (ELVI) Indicators and Forecast Dataset," One Earth Future, 2021, available at: https://oefdatascience.github.io/REIGN.github.io/

[27] Including Håvard Hegre et al., "ViEWS: A political violence early-warning system," *Journal of Peace Research* 56, no. 2 (2019): 155-174; Matina Halkia et al., "The Global Conflict Risk Index: A quantitative tool for policy support on conflict prevention," *Progress in Disaster Science* 6 (2020); Foundation for Partnership Initiatives in the Niger Delta, "An Integrated Conflict Early Warning and Early Response System: Manual for Data Collection and Analysis," 2019, https://fundforpeace.org/2020/02/12/an-integrated-conflict-early-warning-early-response-system/.

[28] Samantha Power, *A Problem from Hell: America and the Age of Genocide* (New York: Basic Books, 2002).

[29] Howard Kunreuther and Paul Slovic, "What the Coronavirus Curve Teaches Us About Climate Change," *Politico,* March 26, 2020, https://www.politico.com/news/magazine/2020/03/26/what-the-coronavirus-curve-teaches-us-about-climate-change-148318.

[30] Robin Gregory, Michael Harstone, and Paul Slovic, "Improving Intervention Decisions to Prevent Genocide: Less Muddle, More Structure," *Genocide Studies and Prevention* 11, no. 3 (2018): 109–127.

[31] Paul Slovic, ed., *The Perception of Risk* (London: Earthscan, 2000); Paul Slovic, *The Feeling of Risk: New Perspectives on Risk Perception* (London: Earthscan, 2010).

[32] Alonzo Plough and Sheldon Krimsky, "The Emergence of Risk Communication Studies: Social and Political Context," *Science, Technology, & Human Values* 12, no. 3/4 (1987): 4–10.

[33] Paul Slovic, "Perception of Risk," *Science* 236 (1987): 280–285.

[34] National Research Council, *Understanding Risk: Informing Decisions in a Democratic Society* (Washington, DC: The National Academies Press, 1996); Paul Slovic, "Going Beyond the Red Book: The Sociopolitics of Risk," *Human and Ecological Risk Assessment* 9 (2003): 1181–1190.

[35] See https://www.arithmeticofcompassion.org.

[36] Johanna Wiss, et al., "The Influence of Identifiability and Singularity in Moral Decision Making," *Judgment and Decision Making* 10 (2015): 492–502.

[37] Daniel Västfjäll, Paul Slovic, and Marcus Mayorga, "Pseudoinefficacy: Negative Feelings from Children who Cannot Be Helped Reduce Warm Glow for Children who Can Be Helped," *Frontiers in Psychology* 6 (2015): Article 6.

[38] Paul Slovic, "The Prominence Effect: Confronting the Collapse of Humanitarian Values in Foreign Policy Decisions," in *Numbers and Ierves: Information, Emotion, and Meaning in a World of Data,* eds. Scott Slovic and Paul Slovic (Corvallis, OR: Oregon State University, 2015), 53–61.

[39] The White House, "Remarks by President Obama and President Park of South Korea in a Joint Press Conference," May 7, 2013, https://obamawhitehouse.archives.gov/the-press-office/2013/05/07/remarks-president-obama-and-president-park-south-korea-joint-press-confe.

Acknowledgements

UNITED STATES HOLOCAUST MEMORIAL MUSEUM

The Simon-Skjodt Center for the Prevention of Genocide of the United States Holocaust Memorial Museum works to prevent genocide and related crimes against humanity. The Simon-Skjodt Center is dedicated to stimulating timely global action to prevent genocide and to catalyze an international response when it occurs. Our goal is to make the prevention of genocide a core foreign policy priority for leaders around the world through a multi-pronged program of research, education, and public outreach. We work to equip decision makers, starting with officials in the United States but also extending to other governments, with the knowledge, tools, and institutional support required to prevent—or, if necessary, halt—genocide and related crimes against humanity.

The assertions, opinions, and conclusions in this report are those of the authors. They do not necessarily reflect those of the United States Holocaust Memorial Museum.