SIX

Predictable Disasters

AI and the Future of Crisis Response

Tarek Ghani and Grant Gordon

The greatest barrier to achieving many of the Sustainable Development Goals (SDGs) lies in fragile settings characterized by extreme poverty, weak institutions, and ongoing vulnerability to natural and human-made disasters. Given current trends, complex emergencies may become even more challenging over the next decade, however, artificial intelligence (AI) holds the potential to transform crisis response to both save and improve many lives.[1] In order to realize that promise, crisis response policymakers will have to prioritize ongoing and new AI investments based on a sophisticated understanding of risk and return.

Bending the curve to meet the SDGs in fragile settings will require new tools and radical improvements in the impact, scalability, or cost-effectiveness of current practices—an ambitious goal that can be supported by the exponential growth in promising machine learning applications. In turn, harnessing AI for crisis response requires a clear-eyed understanding of the conditions under which

1. We use the term "artificial intelligence" to refer to automated processes using algorithms to make inferences from data with self-directed learning and adaptation, including, but not limited to, machine learning applications. We use "data science" to describe the broader set of capabilities necessary to implement machine learning projects. We use "crises" to reference both human-made and natural disasters, and we distinguish between the two types as relevant.

We thank Homi Kharas, John McArthur, and Izumi Ohno for editorial wisdom; Megan Roberts for insightful peer review; and Alex Diaz, Stuart Campo, Leila Toplic, Robert Kirkpatrick, Heather Roff, Sarah Sewall, Nathaniel Raymond, Leonardo Milano, Philip Tetlock, Robert Trigwell, Cameron Birge, and Jay Ulfelder for valuable conversations.

machine learning can improve outcomes as well as a framework for when and how to effectively integrate machine learning into organizations.

As we describe below, AI is reshaping our ability to anticipate, respond to, and recover from crises. It increases visibility and access to areas that have historically been inaccessible; it expands capacities to identify and predict crises and their evolution; and it enhances the effectiveness and efficiency of resource allocation and optimization during response efforts.[2] AI does this by strengthening the accuracy and precision of what we know, the speed with which we know it, and the ability to continuously optimize decisions that require analyzing many fast-changing variables simultaneously.[3]

Machine learning applications have already begun to transform three key functions of crisis response policy and programming, which we expect to accelerate over the coming decade. First, machine learning is helping decisionmakers continuously assess the risks of new and ongoing crises, particularly in the domain of natural disasters where data is rich, scientific modeling of underlying causes is advanced, and events are frequent enough to support robust feedback loops. Second, humanitarian and governmental crisis responders are increasingly using machine learning to improve targeting, intervention selection, and service delivery. And third, machine learning is streamlining the mobilization and prepositioning of resources for first responders, with current applications ranging from anticipatory financing for disasters to optimizing the logistics behind delivering humanitarian aid.

What is a vision of the future of crisis response in which AI breakthroughs have been successfully scaled? It is one in which crisis response actors increasingly know where and how crises will happen, and crisis policymakers have the information required to resource and launch efforts to prevent and mitigate these crises. It is a world in which, when unavoidable crises do unfold, financing is immediately released to provide life-saving assistance to those affected based

2. We focus here on "first-generation" machine learning applications that seek to structure, automate, and inform crisis response decisions at macro- (for example, national), meso- (for example, sub-national), and micro-levels (for example, individual). To address first-order questions in the crisis response field, we limit ourselves to identifying, predicting, and optimizing crisis response rather than the number of downstream applications that also shape crises, ranging from the use of automated image processing by drones to AI-based precision agriculture to reduce the impact of climate change.

3. All of this is made possible by general improvements in machine learning algorithms and computational processing power, which support AI-based innovation in any field. But like data, which is context-specific and often a limiting factor, the most valuable AI also needs a feedback platform where interventions and predictions are tested against reality and continuously improved. While basic machine learning applications simulate such feedback by splitting a single dataset into "testing" and "training" components, a frequently occurring decision problem informed by continuous stream of data on relevant inputs and target outcomes is an ideal condition for AI applications.

on preagreed-upon triggers. And, should crises continue, targeted and context-specific aid can be consistently delivered in record time, at scale, and at a radically lower cost, helping to save lives, reduce suffering, and speed recovery.

Before this vision for AI in crisis response can be realized, several key conditions must first be met. To start, data quality, consistency, and coverage need to improve. Absent these improvements, the accuracy of machine learning predictions will be constrained to the limitations of existing data coverage, which too often reflect biased understandings of the world and structural inequities.[4] As data deserts continue to shrink amid a growing range of sources, the quality of AI predictions will also improve. Next, decisionmakers must identify and develop a range of feedback platforms to enable rigorous testing of new machine learning approaches relative to current practices. Last, frameworks and tools for the ethical and accountable use of AI technologies must be created or strengthened in relevant institutions to protect against potential abuse and harm.[5] Together, these necessary steps will help provide the operational architecture needed to help effectively, efficiently, and safely integrate AI into the crisis response field.

While some advocates have previously promoted AI as a crystal ball to predict and prevent global crises, this aspiration obscures the political and organizational constraints that shape crisis response as well as the type of decisions that such predictions can influence. Decisionmakers should consider the political and technical feasibility of any investment in AI as well as its expected impact, conscious that leveraging the impact of any potential application is predicated on identifying the types of decisions amenable to machine learning. It also requires investing in the capabilities or partnerships to deepen machine learning expertise and rigorously assessing the impact of machine learning applications relative to current practices.

Too often, discussions of AI descend into polarized caricature. While techno-utopianism often promotes non-testable platitudes or inflated aspirations of a single project to change the world, techno-pessimism can often fall into a similar trap of developing sweeping generalizations from isolated examples of failed projects. We aim to move beyond these dichotomies by articulating a framework to understand where the expected risks and returns are highest from AI in crisis response. To develop and situate this framework, we analyze current use cases of machine learning and explore the boundaries of their application. Overall, we recommend policymakers adopt a portfolio investment approach to AI that adjusts potential benefits against common risks of political or technical

4. See, for example, Glandon and others, and Toplic.

5. See, for example, NetHope's Artificial Intelligence (AI) Ethics for Nonprofits Toolkit, https://solutionscenter.nethope.org/artificial-intelligence-ethics-for-nonprofits-toolkit.

infeasibility and assesses potential impact relative to current practices. Our goal is to help move the debate from "Does AI change everything?" to "When and how can specific tools most usefully augment or transform current practice?"

We focus on the potential impact AI holds for crisis response as generally reflected in Sustainable Development Goal 16's objective of promoting peaceful and inclusive societies. This framing allows us to specifically address those populations furthest behind in the SDGs—across and within countries—who are most often the primary beneficiaries of crisis response efforts. However, it limits our ability to address how AI is changing many of the long-term drivers of crisis—including poverty, inequality, and economic opportunity—captured in other SDGs and covered by other chapters in this volume. And while we focus on the implications of AI for crisis policymakers in governments and international organizations, we also attempt to highlight implications for those most directly affected by crisis—including concerns about individual autonomy, consent, and privacy, that are central to discussions and decisions regarding AI.

Crisis and Opportunity

Since the start of this century, natural and human-made disasters have levied a rising toll in lives, livelihoods, and social stability. From 2000 until 2019, the UN recorded over seven thousand major natural disasters that claimed over 1.2 million lives, and affected roughly 4 billion additional people, many individuals more than once.[6] Over those same years, the Uppsala Conflict Data Program recorded over 940,000 fatalities from organized violence, including fifty-four active state-based conflicts in 2019—the highest number since the end of World War II.[7] In 2020 alone, UNHCR counted 82.4 million people forcibly displaced worldwide due to conflict, natural disasters, and related disruptions, nearly doubling the 43.3 million estimate just ten years earlier.[8] And as we write this article in September 2021, over 221 million people have been infected globally with COVID-19 and more than 4.5 million have died so far, with those figures expected to grow considerably before the pandemic ends.[9]

6. United Nations Office for Disaster Risk Reduction, "Human Cost of Disasters 2000–19," www.undrr.org/publication/human-cost-disasters-2000-2019.

7. Uppsala Conflict Data Program https://ucdp.uu.se/downloads/brd/ucdp-brd-conf-201-xlsx.zip. The 940,000 total reflects best estimates of "battle-related deaths" defined as fatalities "caused by the warring parties that can be directly related to combat" and thus exclude indirect deaths due to disease, starvation, criminality, or attacks directed at civilians. For more details, see Pettersson and Öberg.

8. UNHCR, "Global Trends 2020: Forced Displacement in 2020," www.unhcr.org/flagship-reports/globaltrends/.

9. Johns Hopkins University & Medicine, Coronavirus Resource Center, https://coronavirus.jhu.edu/map.html.

In listing humanity's global challenges, the UN's SDGs 1.5 and 11.5 highlight the need to prevent and mitigate the compounded harm posted by frequent disasters and complex emergencies. However, the impact goes far beyond these two goals, as it is difficult to imagine achieving many SDGs—from ending extreme poverty and hunger to expanding education and healthcare to promoting peace and protecting the environment—without better crisis response. Before the impact of COVID-19 was clear, researchers had identified the set of countries where the SDGs were already on track to fail based on current trends in climate change and ongoing vulnerability to political instability: by 2030, between two-thirds and 80 percent of the world's poor are likely to live in fragile and conflict-affected countries.[10] This development frontier grows larger if one includes subnational hotspots in otherwise prosperous countries: one Brookings estimate identified 840 poverty hotspots across 102 countries expected to host 1.7 billion people in 2030.[11]

There is a growing recognition that weak institutions and systemic poverty make fragile contexts more vulnerable to both political crises and climate hazards, and slower to recover from both types of shocks.[12] In short, we observe a self-reinforcing cycle in which the hardest development cases remain persistently behind, even while emerging and developed economies continue to progress. And so, improving the ability of crisis response to prevent and mitigate the impact of disasters in fragile settings could imply substantial development returns for those furthest behind.

Trends in modern crises and response efforts suggest the challenge may grow even more complex. For one, the impacts of the climate crisis are increasingly visible in fragile settings and projected to become even more severe this decade across Central America, the Middle East, the Sahel, and Southeast Asia. Wherever climate hazards intersect with weak governance, the challenges likely to face crisis responders will multiply exponentially. Whether fleeing human-made or natural disasters, displaced populations will place increasing pressure on urban centers and neighboring countries, contributing to a regionalization of each local crisis as societies strive to adapt. Meanwhile, international institutions may continue to struggle to rebuild confidence in the face of major power rivalries, while on-the-ground coordination challenges keep growing from an influx of governmental and non-state actors. Moreover, many challenges in crisis response may be further exacerbated by the long-term humanitarian and economic impact of COVID-19.[13]

10. Corral and others, 18; OECD, 3.
11. Cohen, Desai, and Kharas, 210.
12. See, for example, World Bank and Ghani and Malley.
13. See Ghani.

The last two decades have seen exponential growth in the data available on crises as well as the tools needed to effectively and efficiently analyze this information.[14] Historically, data on crises were limited to frontline reports and individual observations. In the 1990s and early 2000s, survey data from crisis--affected populations, operational data from implementers, and historical cross-national datasets were used to inform decisionmaking. Over this past decade, the informational floodgates have opened, releasing high-frequency, granular data from social media, satellite and sensors, call detail records, and so on. Consequently, crisis-related data is rapidly improving in quantity, frequency, granularity, and structure—and thus creating new opportunities to anticipate, respond to, and recover from modern crises.[15]

Recent estimates indicate that, from 2015 to 2022, the number of people online will triple to 6 billion. By 2030, 7.5 billion people will likely use the internet—90 percent of the projected world population six years of age and older.[16] Mobile phones will drive that growth, especially in developing countries, where the aftershocks of disasters can be more severe: from 2020 to 2025, mobile subscribers are projected to grow from 5.2 billion to 5.8 billion, but mobile internet users will grow twice as much, from 3.8 billion to 5 billion as smartphones become more widely available.[17] The impact for data collection—even in the most challenging settings—are already myriad: researchers are conducting phone surveys at unprecedented scale with Voice over Internet Protocol (VoIP), proxying for population movements with cell phone–related location data, and creating new measures of activity and sentiment from social media.[18] Standard early warning indicators like commodity price changes are now more easily tracked with online databases.[19] The cost of remote sensing imagery—from satellites, planes, or low-flying drones—continues to fall. Even remote audio sensors are used in settings like Syria to warn civilians of incoming aircraft or nearby gunfire.[20]

These changes have vastly increased the ability to integrate AI into crisis response, catalyzing a number of high-profile projects that have raised public expectations around AI's potential. To help separate hype from reality, we provide a simple framework that outlines our assessment of the underlying promise and pitfalls for integrating AI into crisis response. Specifically, we highlight how the risks to AI's breakthrough potential in crisis response vary by complexity and

14. Gleditsch, 301–14.
15. For useful overviews, see Panic and Pauwels.
16. See Morgan.
17. GSMA, 6.
18. Relevant examples include Flowminder, Orange Door Research, and Premise.
19. See, for example, Cavallo, Cavallo, and Rigobon.
20. See, for example, Hala Systems.

timeline of the crisis. We rank order risks in our framework but leave returns unspecified, enabling policymakers to apply a risk-adjusted weighting to potential returns for any particular investment in AI in crisis response.

Figure 6-1 categorizes potential AI applications in crisis response across two dimensions. The horizontal axis represents the crisis timeline, ranging from prevention to response. The vertical axis represents crisis complexity, ranging from less complex natural disasters, such as floods, to relatively more complex human-made disasters, such as conflict.[21] We argue that the expected value of any AI investment is driven not only by returns (for example, size and scale of potential impact) but also by the risks posed by technical and political barriers of integrating AI into crisis response efforts. Technical feasibility reflects characteristics that make a crisis context more or less amenable to AI implementation: frequency of the event, data availability and quality, and modeling complexity. For example, natural disasters are often high-frequency events with better data and more reliable scientific models, given the underlying natural processes involved when compared to complex disasters such as civil war. Political feasibility relates to the ability of crisis policymakers to act on the predictions AI can help improve.

Figure 6-1. Feasibility of AI Applications in Crisis Settings

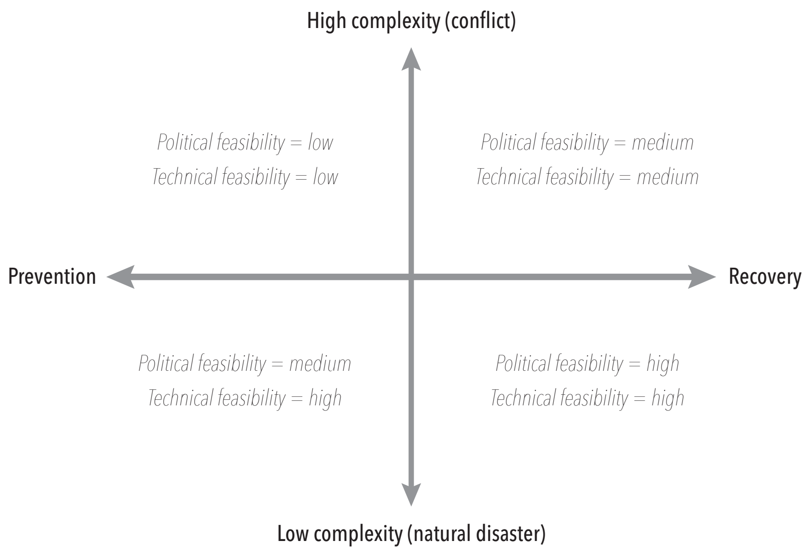

21. For ease of interpretation, we limit the framework's time horizon to one year before and one year after crisis onset. While prevention is desirable over a longer time period, multi-year time spans provide less help in thinking about the decision constraints faced by crisis policymakers and how best to invest in response in AI applications.

This is driven by the incentives around the timing of an intervention—because, unfortunately, response is often easier to mobilize than prevention—and the extent to which a crisis is driven by natural causes, which can often facilitate quicker local and international cooperation.

This framework has several important implications. First, the political feasibility of crisis response tends to increase as a crisis evolves and expands beyond the realm of prevention. The technical feasibility of response also increases for more complex crises as more data enables model refinement. Second, it is often more politically and technically feasible to respond to low-complexity crises than high-complexity crises, though for different reasons. On the one hand, cooperation on natural disaster response is often a win-win for political leaders relative to the contentious politics of conflict prevention and mitigation. And on the other hand, data and modeling of natural disasters is far more advanced than of human-instigated disasters.

After combining these feasibility constraints, we observe the highest risk--adjusted returns for AI investments in the bottom-right quadrant with natural disaster relief efforts, and the lowest risk-adjusted returns in the top-left quadrant with conflict prevention efforts. While crisis response actors should continue pursuing promising, high-return AI investments beyond natural disaster relief efforts, we would urge them to probe the technical and political constraints that may impede the success of those efforts early on in the process.

There are also some important caveats to highlight when considering integrating AI into crisis response, given fundamental concerns and constraints regarding data in crisis settings. Above all, having more data—and more complex data—does not unequivocally imply more useful data. Data that is inaccurate, imprecise, or biased can undermine analysis and response. This is true across crisis-affected countries—for example, there is more data available for Jordan than South Sudan—as well as within countries—for example, social media feeds reflect areas with connectivity. This data unevenness is a major limiting factor and often disadvantages the most marginalized and least digitally connected communities, regularly requiring policymakers to be savvy consumers and communicators of the pitfalls inherent to data-driven analysis. Moreover, collecting and analyzing these new types of data require capabilities not always well represented in crisis response organizations, which in turn demand new organizational investments or partnerships. These are nontrivial but soluble challenges that should be assessed and prioritized early in each new application of machine learning.

There are also a growing set of privacy, security, and ethical challenges around crisis data management.[22] Creating and operationalizing ethical guidelines for

22. For an overview of ethical considerations in applying AI to conflict-related crisis response, see Pauwels.

AI in crisis response is crucial to appropriately safeguarding individuals and following a no-harm principle, particularly when projects are serving vulnerable and marginalized populations. Safeguarding concerns are real and a number of AI projects in the humanitarian space have demonstrated the potential negative consequences of these methods. While much work remains to be done, emerging frameworks for crisis response provide guidance on how to ensure privacy, accountability, safety, and security, as well as a number of key ethical principles.[23] Again, the responsibility falls to all crisis response actors to ask hard questions and demand clear answers about the ethical use of crisis data, so that risks are properly balanced against potential returns.

Prevention and Mitigation

As digital data sources proliferate, machine learning can improve how policymakers anticipate and monitor new crises. Indeed, in the notable case of flood warning and mapping, that promise is already real. But as local contexts and disaster types vary greatly in terms of reliable signals, the crisis response field remains a long way from a crystal ball for crises. For example, in January 2020, well-informed observers understood the risks were high for a global pandemic, famine in Yemen, and war in Ethiopia—but few could say with confidence when and how those risks might unfold. And with risks proliferating nearly as fast as data, decisionmakers need help deciding what data to pay attention to and when.

Digital technologies have transformed the information available to crisis policymakers, offering unprecedented insights from disaster settings even as increasing data availability complicates their efforts to separate signals from noise. In an ideal-case scenario, crisis responders know exactly which data streams to analyze—and how—to better anticipate the likelihood, impact, and profile of an emerging disaster. Given their high frequency and natural processes, floods provide one such example where data scientists have made noteworthy progress—and with clear relevance to multiple SDGs, given the potential of flooding to close markets, disrupt food security, shutter schools, and spread diseases. The EU's Global Flood Awareness System (GloFAS) produces daily flood forecasts and monthly seasonal outlooks using weather data and hydrological models, while UN researchers developed a machine learning approach to processing satellite imagery that reduced the time to develop a flood map for emergency response teams by 80 percent.[24] More generally, scientists are working to translate UN

23. This includes the Humanitarian Data Science and Ethics Group's Framework for the Ethical Use of Advanced Data Science Methods in the Humanitarian Sector and The Harvard Humanitarian Initiative's Signal Code.

24. See European Commission, "Global Flood Awareness System," www.globalfloods.eu/; Nemni and others.

climate projections into impact models that can offer localized forecasts to show how climate change will affect critical sectors such as water, agriculture, or forestry, and nonprofit organizations are using climate models to strengthen physical infrastructure and population resilience against earthquakes and typhoons.[25]

But while climate hazards often lend themselves more readily to data collection and sophisticated modeling approaches, other common disaster types, such as pandemic disease, food insecurity, forced displacement, or deadly conflict are more directly influenced by hard-to-predict behavioral processes. So far, predictive modeling of the onset and evolution of complex emergencies has had more modest success than predictive modeling of natural disasters, but noteworthy efforts continue to make encouraging progress. And as many human-made disasters have gradual onsets followed by cycles of intensification, decline, and often relapse, attempts to mitigate low-level or persistent crises may contribute to prevention efforts over time.

The UN's OCHA-Bucky predictive model of COVID-19 spread and mitigation in humanitarian crises is one informative case. Developed as a collaboration between OCHA's Centre for Humanitarian Data and the Johns Hopkins Applied Physics Laboratory, Bucky provides humanitarian decisionmakers with subnational, four-week projections of the likely spread of the current pandemic in key fragile countries to inform resource planning and facilitate sophisticated scenario analysis of non-pharmaceutical interventions, such as changing social behavior, limiting movement, increasing healthcare access, or prioritizing medical care to vulnerable groups.[26] By combining pre-pandemic data on subnational demographics, intra-regional mobility, and social contact norms with regular updates on local case counts and global disease characteristics, Bucky provides insights for humanitarian responders in contexts such as Afghanistan, Iraq, and South Sudan. While the model is only robust to the accuracy and completeness of underlying data inputs, it goes beyond more standard "dashboard"-style exercises through continuous refinement and the ability to explore counterfactual scenarios. As such, it both complements and incentivizes the types of fundamental data investments necessary for sustained progress on AI adoption.

The devastating scale of COVID-19 has helped spur progress on disease prediction models, such as Bucky, that could generate increased demand for—and investment in—the data, algorithms, and feedback platforms needed not just for COVID-19 and guarding against future pandemics, but also addressing common

25. See, for example, the Inter-Sectoral Impact Model Intercomparison Project, www.isimip.org/; Build Change, https://buildchange.org/.

26. Center for Humanitarian Data, "OCHA-Bucky: A COVID-19 Model to Inform Humanitarian Operations," https://centre.humdata.org/ocha-bucky-a-covid-19-model-to-inform-humanitarian-operations/.

diseases like malaria—which was responsible for over 400,000 fatalities in 2018 alone.[27] At the time of writing this article, the Biden administration released a USD $65b plan for pandemic preparedness, including $3.1b for an early warning detection system.[28] These advances underscore the potential to leverage data science for this type of modeling. Moreover, methods recently developed, for example, to model the spread of COVID-19 in Cox's Bazar refugee camps in Bangladesh using open-source census datasets, locations of potential gathering places, and other information on daily movements, could be adapted and expanded with cell phone–related location data to help decisionmakers as far away as Uganda prioritize antimalarial bed net distribution and determine where to concentrate spraying for mosquitoes.[29] This can be helpful for targeting, even where malaria is endemic and transmission is year-round rather than subject to outbreaks. Other crises may benefit as well. Efforts to predict food shortages and prioritize the allocation of resources both between and within countries to prevent potential famines could receive fresh attention if the COVID-19 prediction models can help raise crisis policymakers' expectations of what insights forecasting approaches might offer elsewhere.[30]

But reliable predictions will most likely lag behind for challenges such as forced displacement and deadly conflict, given issues of data and modeling, even though those crises may become more frequent as climate change exacerbates food insecurity, water scarcity, and resource competition.[31] Forced migration data is comparatively scarce in developing countries. The most notable and systematic data sources are the IOM's Displacement Tracking Matrix and the Internal Displacement Monitoring Centre's Global Internal Displacement Database. The UNHCR and World Bank also launched the Joint Data Center on Forced Displacement with the goal of improving data on forced migration.[32] Moreover, decisions on migration frequently involve a set of push-and-pull factors such as climate impacts or political instability in origin countries and economic opportunities and social freedoms in destination countries. While these factors are

27. World Health Organization, "Malaria Fact Sheet," www.who.int/news-room/fact-sheets/detail/malaria.

28. StatNews. "The White House wants $65 billion for an 'Apollo'-style pandemic preparedness program," September 3, 2021. www.statnews.com/2021/09/03/biden-wants-65-billion-for-apollo-style-pandemic-preparedness-program/

29. See UN Global Pulse, "Modeling the Spread of COVID-19 and the Impact of Public Health Interventions in Cox's Bazar and Other Refugee Camps." www.unglobalpulse.org/2020/10/modelling-the-spread-of-covid-19-and-the-impact-of-public-health-interventions-in-coxs-bazar-and-other-refugee-camps/.

30. See, for example, Andrée and others.

31. For a review of AI in the human security field, see Roff.

32. Sarzin.

challenging to disentangle, much less predict, some researchers have made modest progress here using online search or advertising data.[33] However, existing migration prediction models often tend to be limited to country-year aggregate predictions that can raise concerns with decisionmakers given a tendency of those models to overweight preexisting trends and underpredict large shocks.[34]

Deadly conflict is subject to data limits similar to those of forced migration, though the spread of cell phones and social media may help future researchers better understand how and why violence breaks out. Current conflict prediction efforts are limited by the availability of high-quality input data from around the world, along with the empirical challenge that the onset of civil conflict is a low-frequency event.[35] For instance, one promising public effort combining unsupervised and supervised machine learning methods to analyze newspaper articles and predict conflict risk is limited to English-language sources, as natural language processing is unavailable for key dialects in many relevant conflict-affected countries, not to mention a standardized set of high-quality news sources across such settings.[36] Another recent rigorous analysis using high-quality conflict micro-data from both Colombia and Indonesia had success only in identifying persistent, subnational high-violence hot spots—and not new outbreaks or escalations of violence.[37]

While conflict early warning remains an elusive goal for crisis analysts given the complexity of political systems, progress continues to be made in predicting and preventing deadly violence at the micro-level. For example, Hala Systems has used remote audio sensors to warn civilians of incoming aircraft or nearby gunfire in settings like Syria. Two promising future areas to watch here are efforts by the UN to track hate speech online and via radio stations, both of which are being led by UN Global Pulse, the big data initiative of the UN Secretary-General.[38] Notably, the latter effort involving radio stations is a creative solution to the challenge of measuring activity in remote, marginalized, and often digitally disconnected communities, and could be used to target conflict resolution and policing efforts.

33. See, for example, Bohme, Groger, and Stohr, and Palotti and others.

34. See, for example, Milano.

35. While several governments have begun applying machine learning to conflict prediction and analysis, few details are available, and assessing the quality and policy application of these efforts remains challenging.

36. Mueller and Conflict Forecast, www.conflictforecast.org.

37. Bazzi and others.

38. For initial results of related previous efforts, see UN Global Pulse, "Exploring the Effects of Extremist Violence on Online Hate Speech," www.unglobalpulse.org/project/exploring-the-effects-of-extremist-violence-on-online-hate-speech/ and "Using Machine Learning to Analyse Radio Content in Uganda," www.unglobalpulse.org/project/pilot-studies-using-machine-learning-to-analyse-radio-content-in-uganda-2017/.

Finally, machine learning can improve mitigation efforts such as resource mobilization and prepositioning for responders by providing more accurate and precise predictions of where and when assistance will be needed. For example, in the emerging field of anticipatory disaster financing, the Centre for Disaster Protection and others have prominently advocated for a reorientation toward risk-based financing approaches based on contingency planning and prespecified triggers.[39] While one of the highest-profile anticipatory finance efforts—the World Bank's Pandemic Emergency Financing Facility—was recently shut down after widely shared criticisms of its slow and modest disbursements as COVID-19 spread around the world, it would be a mistake to discard a much-needed alternative to chronically underfunded humanitarian appeals because the financial parameters and operational arrangements of one example were poorly calibrated.[40] Much like the market for terrorism insurance after the 9/11 attacks, it will take time and government investment in order for the anticipatory disaster financing market to mature.[41] As of July 2020, the UN's Central Emergency Response Fund (CERF) disbursed US$5.2 million after a GloFLAS prediction of high probability of flooding in Bangladesh—the fastest CERF allocation in history and the first one to take place before peak flooding.[42] And beyond anticipatory finance, machine learning already delivers large cost savings by optimizing humanitarian aid delivery systems for agencies like WFP.[43] As the predictive models discussed above continue to improve, useful applications are also likely to emerge for the prepositioning of humanitarian resources.

Relief and Recovery

In addition to supporting crisis prevention and mitigation, data availability and machine learning have generated a sea change over recent years for crisis response efforts. OCHA's Centre for Humanitarian Data now houses over seventeen thousand humanitarian data sets as well as a specific catalogue of predictive models in the humanitarian sector.[44] Moreover, a recent analysis of predictive analytics in

39. See Guidance Notes for Highly Effective DRF, www.disasterprotection.org/guidance-notes-for-highly-effective-drf.

40. See Clarke.

41. Michel-Kerjan and Raschky.

42. Center for Humanitarian Data, "Anticipatory Action in Bangladesh before Peak Monsoon Flooding," https://centre.humdata.org/anticipatory-action-in-bangladesh-before-peak-monsoon-flooding/.

43. World Food Program, "Palantir and WFP Partner to Help Transform Global Humanitarian Delivery," February 15, 2019, www.wfp.org/news/palantir-and-wfp-partner-help-transform-global-humanitarian-delivery.

44. Center for Humanitarian Data, "Catalogue of Predictive Models in the Humanitarian Sector," https://centre.humdata.org/catalogue-for-predictive-models-in-the-humanitarian-sector/.

the humanitarian space commissioned by the then UK Department for International Development identified forty-nine different projects predicting the who, what, where, or when of crises.[45] This has created an informational foundation upon which machine learning has transformed three core areas of humanitarian relief and recovery: (1) targeting, the identification of how to allocate resources across crisis-affected populations; (2) intervention selection, the determination of which services should be provided; and (3) service delivery mechanisms, the ways in which core services are provided to clients. While these efforts do not necessarily increase the amount of resources allocated to a given crisis, they have already begun to influence how existing resources are allocated within a crisis. While relief and recovery efforts have demonstrated that AI can be integrated well at the project level, they have yet to be effectively scaled through broader adoption of practices and approaches.

One of the most promising areas for AI to transform humanitarian relief and recovery is in targeting aid delivery. Assessments to determine individual- and population-level needs and vulnerabilities are launched after every crisis and integrated into every project implemented in response. Often these are time- and capital-intensive processes limited in scale and precision by traditional data sources and analytic methods. Leveraging satellite imagery, cell phone records, and other administrative data, machine learning applications can systematically automate, at scale, the assessment process to more effectively and efficiently understand "who has what" and prioritize "who gets what" in crisis response.

A number of noteworthy pilot efforts have demonstrated AI's ability to improve this type of targeting in recent years. In 2014, GiveDirectly, an organization that provides direct cash transfers to the world's poor, developed algorithms to process satellite imagery and detect different types of roofing in Uganda, which was highly correlated with household economic characteristics, to enhance targeting cash payments.[46] Rather than conducting time- and labor-intensive village-level surveys to identify household-level poverty data, these algorithms automated the selection process, increasing the speed and cost-effectiveness of targeting and therefore increasing the ability to provide relief to a greater number of individuals.

More recently, to provide the most vulnerable Togolese citizens with cash support to weather the health and economic consequences of COVID, Joshua Blumenstock of UC Berkeley and co-authors used deep learning algorithms to process satellite images and phone usage data to map extreme poverty and target the transfers accordingly.[47] Using detection approaches for the satellite images

45. Hernandez and Roberts.

46. DataKind, "Using Satellite Imagery to Find Villages in Need," www.datakind.org/projects/using-the-simple-to-be-radical/.

47. Joshua Blumenstock, "Machine Learning Can Help Get COVID-19 Aid to Those Who Need It Most," *Nature*, May 14, 2020, www.nature.com/articles/d41586-020-01393-7.

in conjunction with call record data to estimate wealth and income, the project aimed to augment slower, survey-based methods that would traditionally identify needs.[48] In the article that formally assesses these methods, Aiken and co-authors note that compared to standard approaches, leveraging satellite and phone data reduced targeting errors by between 4 to 21 percent to many of the poorest citizens.[49]

This type of targeting has also been deployed outside of crisis contexts, with similarly large gains. IDInsight, an organization dedicated to reducing poverty through data and evidence, demonstrated that machine learning approaches could be used to identify out-of-school girls in parts of rural India, building an algorithmic model that increased the ability to locate between 50 percent and 200 percent more out-of-school children at the same cost as historically conducted individual household surveys.[50] These approaches have not only offered the ability to accelerate processing, but have also widened the pool of recipients by leveraging larger population datasets. In future crises, these datasets can be used as a foundation for targeting assessments, further increasing efficiencies and the ability to quickly and nimbly respond to crises.

While advances in targeting have helped answer the question of who needs what, AI has also provided the ability to help determine what is needed. A key advance in applications of machine learning is the ability to dynamically identify needs and optimize crisis response based on what works in a specific program context. In the first-ever adaptive experiment implemented in a humanitarian context by the International Rescue Committee in Jordan, academic researchers developed a machine learning algorithm to allocate different types of employment support services to Syrian refugees and vulnerable Jordanians based on their individual characteristics and how those different support services have generated impact.[51] This type of "precision social service delivery" optimized the specific package of services provided to each individual in order to maximize their individual outcomes. The data-driven approach generated a 20-percentage point improvement in the probability that refugees and vulnerable Jordanians were offered a job and were in formal wage employment six weeks after receiving targeted support.

Improvements in data availability and machine learning have provided the ability to better understand individual needs and create individual-level relief packages. Médecins sans Frontières created a machine learning–based application that allows nonexpert clinicians in low-income settings to identify antibiotic

48. Aiken and others.
49. Aiken and others.
50. Brockman and others.
51. Caria and others.

resistance using image processing methods and create bespoke treatment regimes for clients.[52] Similar approaches have been used to provide educational services in emergency contexts. Prior to COVID-19, 75 million children were out of school in crisis-affected countries, representing half of the world's out-of-school population. With only 2 percent of humanitarian funding allocated to education in emergencies and a lack of educational infrastructure and trained workforce, AI has been used to develop individually tailored learning experiences delivered through tablets to provide bespoke educational content to children across multiple learning levels.[53] Can't Wait to Learn is a digital game-based learning software developed by War Child Holland, which allows children to learn at their own pace and level. A quasi-experimental analysis of the program demonstrated that this approach led to significant improvements in math and literacy skills as well as psychological well-being.[54] While EdTech has clear limitations relative to in-person instruction, these individually focused services demonstrate how AI is increasingly transforming core service provision during the relief and recovery phases of crisis.[55]

Moreover, as the digital revolution sweeps through crisis-affected contexts, and given that refugees are increasingly displaced in middle-income countries, data availability and machine learning will likely continue accelerating innovation in service delivery.[56] Especially in the wake of COVID-19, digital delivery is emerging as an alternative cost-effective way to provide services in crisis contexts, including access to healthcare, cash to meet basic needs, educational content for children, and job platforms in local labor markets. Naturally, AI is shaping implementation models, enabling remote health consultations, supporting software to deliver and track cash, creating educational content, and algorithmically matching individuals to job opportunities.[57] For example, SkillsLabs is an example of a software and machine learning–based approach that largely helps refugees navigate labor markets and match into jobs in the EU.[58] Similar platforms have been established for Syrian refugees in Jordan.[59] A crucial challenge to leveraging AI for job-matching platforms in low- and middle-income countries is the availability of job opportunities and lack of evidence on what individual-level characteristics predict high-quality matches. Nonetheless, these

52. See Médecins Sans Frontières, Antibiotic Resistance, www.doctorswithoutborders.org/what-we-do/medical-issues/antibiotic-resistance, and Google (2019).

53. See, for example, Education Cannot Wait, www.educationcannotwait.org/the-situation/.

54. Brown.

55. Rodriguez-Segura and Crawfurd.

56. Devictor.

57. GSMA.

58. See Skills Lab, www.etf.europa.eu/en/projects-activities/projects/skills-lab.

59. See, for example, ILO Skills Platform, www.ecsjo.com/.

platforms demonstrate that even in these settings, there are gains in employment outcomes to using algorithmic matching.[60]

Over the past two decades, information provision has emerged as a key component of humanitarian service delivery.[61] Providing accurate, timely, and precise information at scale to crisis-affected populations enables them to make informed individual decisions about the context and how to respond. Here, too, AI has enhanced the ability to deliver information. For example, the Norwegian Refugee Council has begun using automated chatbots to provide Venezuelan migrants in Colombia with details on their rights within the country.[62] Information provision may be one of the areas most amenable to machine learning applications, as it seeks to provide high-frequency data over existing digital platforms.

AI holds the potential to transform the operational and financial model of how the humanitarian sector responds to emerging crises. The humanitarian response system is driven by what Stefan Dercon and Daniel Clarke call the "begging bowl" problem: a crisis breaks out, humanitarian responders deploy and make the case that aid is needed, and donors aim to overcome a collective action problem to finance response.[63] In practice, this can generate major delays between the advent of a crisis and when humanitarian aid is unlocked. This dynamic is, in part, driven by the inability to accurately and precisely identify when crises will break out and the consequent distrust by donors of needs assessments given the incentives for responders to potentially inflate humanitarian needs.[64]

Real-time data flows and machine learning applications will increase the ability to objectively identify and measure crises as they unfold, opening up opportunities to move into risk-based financing and reshaping how humanitarian response is delivered. Instead of an operational infrastructure grounded in post-hoc fundraising and service delivery, a future humanitarian system could orient around an operational structure that flexibly increases capacity for rapid response as a crisis worsens.[65] The Danish Red Cross and International Federation of Red Cross and Crescent Societies, for example, recently launched the first volcano catastrophe bond, which would release large tranches of funding for disaster response according to a tiered trigger structure.[66] Recently, UN OCHA

60. Caria and others.
61. Greenwood.
62. Toplic.
63. Clarke and Dercon.
64. Konyndyk.
65. Talbot, Dercon, and Barder.
66. See Reuters, "Danish Red Cross launches volcano catastrophe bond," www.reuters.com/article/us-volcano-insurance-bond/danish-red-cross-launches-volcano-catastrophe-bond-idUSKBN2BE00J/.

launched a pilot program in Somalia to explore how these types of instruments can be adapted for drought and sudden-onset emergencies.[67]

Harnessing Breakthrough Potential

Harnessing AI's breakthrough potential requires decisionmakers to recognize that machine learning applications are no panacea but do offer real opportunities to save and improve the lives of those affected by crisis. It requires moving beyond broad debates over whether or not AI is useful to instead embrace systematic analyses of the conditions under which machine learning enhances or detracts from current practice. It requires moving beyond the dichotomy of quantitative versus qualitative data to an approach that identifies and integrates the comparative advantages of each type of data as available. And it requires moving beyond vague theories of change to concrete assessments of breakthrough potential in impact, scale, and cost-effectiveness in specific contexts and plans to develop required new capabilities.

At the ecosystem-level, benefiting from the potential returns of AI requires investing in data quality and coverage, launching feedback platforms that take result-based learning seriously, and strengthening data ethics standards in governance frameworks. At the organizational level, it requires identifying decisions that can integrate machine learning applications, establishing capabilities, and assessing the impact of these new approaches. From an organizational perspective, the challenges of integrating machine learning into crisis response are not different from the broader question decisionmakers face when they determine if and how to invest in new capabilities: (1) what are the anticipated returns from a new approach; (2) how should insights from a new approach be integrated into organizational decisionmaking processes and culture; (3) should new capabilities be in-housed or developed through partnership, and so on. Of course, answers to these questions might vary by organization: for example, smaller, agile organizations may be able to integrate new technologies more quickly, whereas larger incumbents may prove to be later adopters at scale.

One noteworthy point in pursuing AI's breakthrough potential in crisis response is that investing in machine learning applications is about everything except the algorithm itself, which is often off-the-shelf technology. Instead, crisis response actors seeking to apply AI will have to develop the data sources, conceptual models, and feedback platforms to implement machine learning applications. Overcoming organizational barriers to adoption is thus an inevitable step

67. United Nations CERF, "CERF and Anticipatory Action," https://cerf.un.org/sites/default/files/resources/CERF_and_Anticipatory_Action.pdf.

for unlocking AI's breakthrough potential. Some issues, like establishing operational procedures or acquiring technical capacity, are straightforward once leaders see a clear value proposition from AI. Other constraints, including cultural resistance to quantitative analysis or ethical concerns over privacy and security of data sources, will require more nuanced attention. Overall, the achievements of data-driven efforts at the UN, such as the Centre for Humanitarian Data and UN Global Pulse, provides hope that progress will continue.

As crisis policymakers consider how to invest in AI, one common concern is that the use of AI in crisis response currently resembles a disparate set of projects rather than a cohesive portfolio. In part, this stems from AI tracking the natural arc that many new technologies take: several seemingly uncoordinated projects are launched to assess potential before more sophisticated programs or strategies take shape. Moreover, unlike the private sector's rapid adoption of AI applications, mechanisms for scaling AI in crisis response only move as fast as the governance approaches needed to support them.

While the danger of AI "pilotitis" may loom large for crisis response, there are three promising pathways to scale and sustainability. First, specific projects should be templatized for reuse, replicated across contexts, and most important, translated into global frameworks. For example, the Stanford Immigration Policy Lab has developed an algorithm to optimize where refugees are resettled within countries. This now needs to be piloted, tested, and replicated across countries, and if the impact and cost-effectiveness is confirmed, the United Nations and its member states could create a between-country matching system. Second, multilateral institutions should invest in public goods infrastructure and governance for data science and machine learning, including but not limited to climate and conflict prediction models. The Center for Humanitarian Data is a critical step in this direction, but more is needed, including articulating an agenda for how AI can meet the goals for preventing and mitigating crises. Third, machine learning should be integrated as a standard priority in donor behavior: data science should be added as a standard budget line in all projects above a certain funding threshold, funding windows that support AI-proven projects should be launched, and support to enhance state and local capacity to use AI should be made available.

As a range of crisis actors seek to integrate AI to save and improve lives across many different disaster contexts, we endorse a risk-adjusted AI investment approach that acknowledges where political or technical obstacles may impede success. With that said, the crisis response field would be best served by a portfolio of efforts that includes a mix of both high-risk, lower-return bets and low-risk, higher-return initiatives. As noted above, each category of crisis response has its comparative advantages and drawbacks, with prevention and mitigation more exposed to the uncertainty associated with political decisionmaking and relief and recovery more likely to focus on micro-level interventions that need to reach scale for significant impact. Achieving breakthrough potential in crisis response interventions over the next decade will not be simple or linear. After all, while the disasters may be predictable, what works best to prevent or respond to them is clearly not.

References

Andrée, Bo Pieter Johannes, and others. 2020. "Predicting Food Crises." World Bank Policy Research Working Paper No. 9412.

Bazzi, Samuel and others. 2019. "The Promise and Pitfalls of Conflict Prediction: Evidence from Colombia and Indonesia." NBER Working Paper No. 25980.

Blumenstock, J. E., G. Cadamuro, R. On. 2015. "Predicting Poverty and Wealth from Mobile Phone Metadata," *Science*, v. 350, no. 6264.

Aiken, Emily and others. "Machine Learning and Mobile Phone Data Can Improve the Targeting of Humanitarian Assistance." NBER Working Paper No. 29070.

Böhme, Marcus H., André Gröger, and Tobias Stöhr. 2020. "Searching for a Better Life: Predicting International Migration with Online Search Keywords," *Journal of Development Economics*, no. 142.

Brockman, Ben, and others. 2019. "Can Machine Learning Double Your Social Impact?" *Stanford Social Innovation Review.* February 20.

Brown, Felicity, and others. Forthcoming. "Can't Wait to Learn: A Quasi-Experimental Evaluation of a Digital Game-Based Learning Programme for Out-of-School Children in Sudan." *Journal of Development Effectiveness.*

Cavallo, Alberto, Eduardo A. Cavallo, and Roberto Rigobon. 2014. "Prices and Supply Disruptions during Natural Disasters." *Review of Income and Wealth*, v. 60, no. S2.

Caria, A. Stefano, and others. 2021. "An Adaptive Targeted Field Experiment: Job Search Assistance for Refugees in Jordan." University of Oxford, Department of Economics.

Clarke, D.J. "Now is not the time for the World Bank to step back on pandemic financing," August 13, 2020. www.disasterprotection.org/latest-news/now-is-not-the-time-for-the-world-bank-to-step-back-on-pandemic-financing.

Clarke, D. J., and Dercon, Stefan. 2016. *Dull Disasters? How Planning Ahead Will Make a Difference*. Oxford University Press.

Cohen, Jennifer L., Raj M. Desai, and Homi Kharas. 2019. "Spatial Targeting of Poverty Hotspots." *Leave No One Behind: Time for Specifics on the Sustainable Development Goals*, edited by Homi Kharas, John W. McArthur, and Izumi Ohno. Brookings.

Corral, Paul, and others. 2020. "Fragility and Conflict: On the Front Lines of the Fight against Poverty." World Bank.

Devictor, Xavier, Do, Quy-Toan, Levchenko, Andrei A. 2020. "The Globalization of Refugee Flows." World Bank Policy Research Working Paper No. 9206.

Ghani, Tarek. 2021. "The Sting in COVID-19's Tail." *Foreign Affairs*, January 26.

Ghani, Tarek, and Robert Malley. 2020. "Climate Change Doesn't Have to Stoke Conflict." *Foreign Affairs*, September 28.

Glandon and others. "Despite the hype, do not expect big data to replace traditional surveys anytime soon," January 7, 2021. www.3ieimpact.org/blogs/despite-hype-do-not-expect-big-data-replace-traditional-surveys-anytime-soon/

Gleditsch, Kristian Skrede, and others. 2014. "Data and Progress in Peace and Conflict Research." *Journal of Peace Research*, v. 51, no. 2.

Google. 2019. "Accelerating Social Good with Artificial Intelligence: Insights from the Google AI Impact Challenge."

Greenwood, Faine, and others. 2017. "The Signal Code: A Human Rights Approach to Information During Crisis."

GSMA. 2018. "Mobile for Humanitarian Innovation: Landscaping the Digital Humanitarian Ecosystem."

Hala Systems. Inc. 2019. "Protect Everything that Matters." https://georgetown.app.box.com/s/qm2raov0qlh7kodt4ubl2p9zj5879x4v.

Hernandez, K., and T. Roberts. 2020. "Predictive Analytics in Humanitarian Action: A Preliminary Mapping and Analysis." K4D Emerging Issues Report 33. Brighton, England: Institute of Development Studies.

Konyndyk, Jeremy. 2018. "Rethinking the Humanitarian Business Model." Washington, D.C.: Center for Global Development.

Michel-Kerjan, E., and P. A. Raschky. 2011. "The Effects of Government Intervention on the Market for Corporate Terrorism Insurance." *European Journal of Political Economy*, no. 27 (Supplement 1).

Milano, Leonardo. "Reviewing the Danish Refugee Council's Foresight Model," October 20, 2020. https://centre.humdata.org/reviewing-the-danish-refugee-councils-foresight-model/.

Morgan, Steve. "Humans on the Internet Will Triple from 2015 to 2022 and Hit 6 Billion," Cybercrime Magazine, July 18, 2019. https://cybersecurityventures.com/how-many-internet-users-will-the-world-have-in-2022-and-in-2030/.

Mueller, Hannes, and Christopher Rauh. 2020. "The Hard Problem of Prediction for Conflict Prevention." Center for Economic Policy Research.

Nemni, Edoardo, and others. 2020. "Fully Convolutional Neural Network for Rapid Flood Segmentation in Synthetic Aperture Radar Imagery." *Remote Sensing*, v. 12, no. 16.

OECD. 2018. *States of Fragility*. OECD Publishing.

Pauwels, Eleonore. 2020. "Artificial Intelligence and Data Capture Technologies in Violence and Conflict Prevention." Global Center on Cooperative Security. September.

Palotti, Joao, and others. 2020. "Monitoring of the Venezuelan Exodus through Facebook's Advertising Platform." *PLOS One*, v. 15, no. 2.

Panic, Brank. 2020. "Data for Peacebuilding and Prevention Ecosystem Mapping: The State of Play and the Path to Creating a Community of Practice." Center on International Cooperation. October.

Pettersson, Therese, and Magnus Öberg. 2020. "Organized Violence, 1989–2019." *Journal of Peace Research*, v. 57, no. 4.

Reuters. "Danish Red Cross launches volcano catastrophe bond," March 21, 2021. www.reuters.com/article/us-volcano-insurance-bond/danish-red-cross-launches-volcano-catastrophe-bond-idUSKBN2BE00J

Rodriguez-Segura and Crawfurd. "What Works in EdTech," Center for Global Development, September 3, 2020. www.cgdev.org/blog/what-works-edtech.

Roff, Heather. 2017. *Advancing Human Security through Artificial Intelligence*. London: Chatham House.

Sarzin, Zara. 2017. "Stocktaking of Global Forced Displacement Data." World Bank Policy Research Working Paper No. 7985.

Talbot, Theodore, Stefan Dercon, and Owen Barder. 2017. "Payouts for Perils: How Insurance Can Radically Improve Emergency Aid." Washington, D.C.: Center for Global Development.

Toplic, Lelia. "AI ethics: 5 considerations for nonprofits," July 22, 2020. www.nethope.org/2020/07/22/ai-ethics-5-reasons-why-nonprofit-engagement-is-key/.

———. "AI in the Humanitarian Sector," Relief Web, October 8, 2020. https://reliefweb.int/report/world/ai-humanitarian-sector.

World Bank. 2020."Poverty and Shared Prosperity 2020: Reversals of Fortune." Washington, DC: World Bank. https://openknowledge.worldbank.org/handle/10986/34496.