Dynamic Early Warning and Action Model: A policy evaluation tool

Alexandra Malaga, Hannes Mueller,
Christopher Rauh, Alessandro Ruggieri and
Benjamin Seimon.

Contents

This material has been funded by UK aid from the UK government; however the views expressed do not necessarily reflect the UK government's official policies.

Introduction

The last decade has seen an explosion in the capabilities of quantitative forecast methods for armed conflict prediction. These are currently widely deployed for analytical purposes. However, their use for prevention poses challenges as policy options are typically not evaluated in systems that integrate quantitative forecasts.[1]

The Dynamic Early Warning and Action Model (DEWAM) is an attempt to fill this gap. It combines two modules:

1.
Forecasting module

▪ Uses machine learning and text data to generate accurate forecasts of outbreak (risk of armed conflict) and intensity (no. of fatalities).[2]

▪ Targeted detection of 'hard-to-predict' onsets i.e. outbreaks of conflict in countries with a long history of peace.

2.
Dynamic decision-making module

▪ Integrates the forecasts to support decision-making.

▪ Allows for the simulation of specific policies and policy combinations.

▪ Quantitatively estimates returns to policies to facilitate internal debate at strategic and operational levels.

The key trade-off that policymakers face in prevention is that, first, acting on low risks means that resources are wasted on crises that never materialise but that, second, waiting for crises to materialise means that countries get stuck and suffer negative dynamics (conflict trap).

The DEWAM was designed to help policymakers make that trade-off between acting on low risks and failing to prevent. This document will use Nigeria as a case-study for demonstrating the usefulness and outputs of the DEWAM.

[1] See the taxonomy provided by FP21 in the appendix, Figure 11.

[2] National (170+ countries) and subnational (60,000+ grid cells) conflict predictions are updated monthly and are freely available at https://conflictforecast.org/. Note that the armed conflict (outbreak) prediction is the likelihood of future armed conflict. However, for countries already in conflict, this is less informative. The intensity forecast predicts the number of fatalities which can provide insight into the potential for escalation/de-escalation.

Methodology

1. Forecasting module

As a primer, this module uses machine learning methods. These are statistical techniques that identify patterns and trends in historical data to produce estimates about what might happen in the future. These methods are extremely flexible and have come to dominate forecasting. But, they require a lot of data to learn due to this flexibility.

For conflict prediction, naive approaches use historical violence as a proxy for future violence. However, predicting a continuation of conflict for countries currently in conflict, or a continuation of peace for countries currently at peace, is of limited benefit to policymakers. The most useful, and most difficult, cases to identify are where the country has a long history of peace but is suddenly becoming susceptible to violence. These cases are rare. Therefore, it is necessary to go far back in time to provide the machine learning algorithm with sufficient instances of outbreaks to learn from.

This is the key innovation of our forecasting model - the use of news text data reaching back to the 1980s to capture outbreaks in these hard-to-predict cases. Over 5 million newspaper articles are geo-located and broken down into a set of 15 distinct topics.[3] We find:

▪ Text is particularly powerful for capturing increasing risk in countries with a long history of peace.

▪ Topics such as *"military conflict"* or *"judicial abuses"* are positively associated with risk, i.e. with more news stories on police violence or the movement of guerilla groups, the risk increases.

▪ Other topics such as *"power and negotiation"* or *"economics"* are negatively associated with risk. When their share in reporting falls, a country is more likely to be at risk.

[3] See examples of text topics and topic-specific vocabulary in the appendix, Figures 12 and 13.

Methodology

2. Dynamic decision-making module

This module provides a framework that enables the forecast to support policy-makers in a practical way. We seek to answer key questions such as:

- At which countries should attention and resources be targeted?
- At what point should a country office escalate risk warnings to attract attention?
- How does a strategic view on policies change the optimal decisions taken?
- Should the focus be on prevention, de-escalation or a combination of the two?
- How can policy-makers combine policies most effectively to build pathways out of violence?

In order to effectively answer these questions we introduce a model that is built on top of the forecasting module.

Figure 1: Building blocks of the dynamic decision-making module

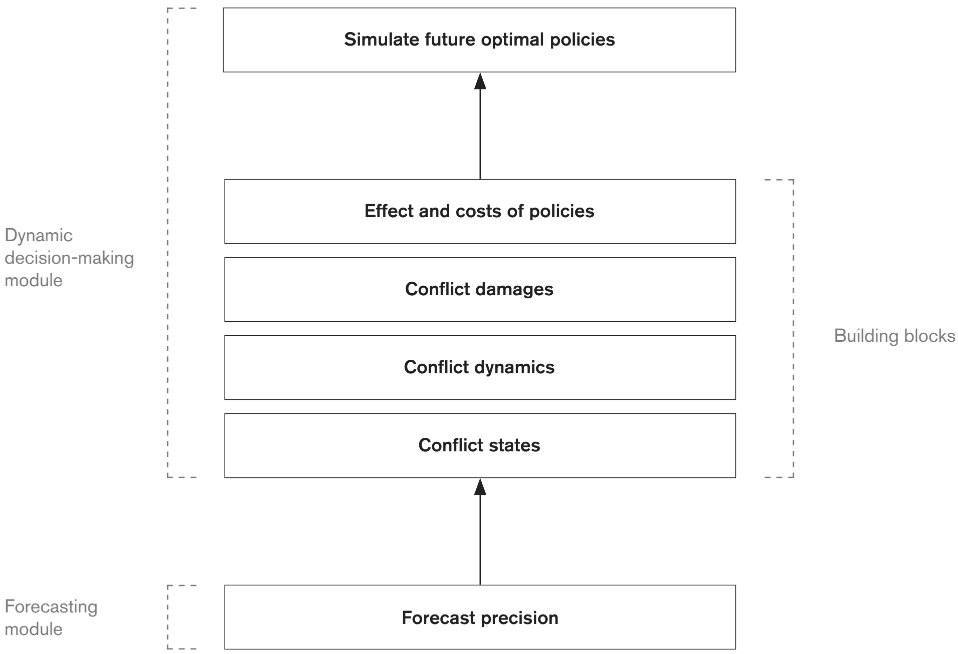

The foundation of the dynamic decision-making module is what we call **conflict states.**

Building block 1: Conflict states

By examining the dynamics of conflict in the past, we distil the cycle of peace and conflict into a spectrum of 12 states.[4] In this way, every country is assigned to a state for every month.[5] It is easiest to think of states as summaries of conflict situations. In general, the higher the state, the more severe the situation. Figure 2 reports a simple characterization of the states using the forecasts from module 1 and the actual intensity of violence.

Figure 2: Descriptive summary of conflict states

State	Predicted likelihood of armed conflict outbreak	Predicted number of fatalities (per 1mm inhabitants)	Observed number of fatalities (per 1mm inhabitants)	Number of observations
1	0.95	0.09	0.0	7981
2	2.56	0.15	0.0	5013
3	7.17	0.16	0.0	3658
4	14.14	0.12	0.0	2307
5	22.24	0.50	0.0	1009
6	20.52	1.16	0.2	482
7	38.49	0.25	0.1	1606
8	41.25	5.48	1.5	258
9	50.29	8.92	0.1	356
10	76.06	1.74	0.2	1443
11	85.77	6.67	5.9	2297
12	84.91	53.51	84.0	431

- **States 1-4, stable peace:** In these states we expect no fatalities resulting from conflict and a low risk of future conflict. For example, Iceland and Switzerland have spent the entire time-span in state 1.

- **States 5-10, elevated risk:** These states generally capture countries that are on the precipice of conflict or have recently experienced violence and are susceptible to recurrence. We expect low levels of deaths in these states, but increased risk of an outbreak. For example, in the two months directly before the invasion of Crimea, Ukraine was in state 6 and 8.

- **States 11 and 12, intense conflict:** These states represent ongoing violence, but are differentiated by its intensity. On average, a country experiences 6 deaths per 1mn inhabitants when in state 11, compared to 84 deaths per 1mn when in state 12. Nigeria and Mexico have spent extended periods in state 11, whilst Afghanistan, Syria and Central African Republic have suffered the most from being in state 12.

[4] The states are generated using a statistical method (Hidden Markov Model) using the forecast as the primary input. The number of states is determined via experimentation.

[5] The latest model is trained on data from January 2010 to April 2023.

Building block 1: Conflict states

Figure 3 illustrates these states using the state history for Nigeria and a selection of comparable countries.

Figure 3: State visit comparison, Jan 2010 to April 2023

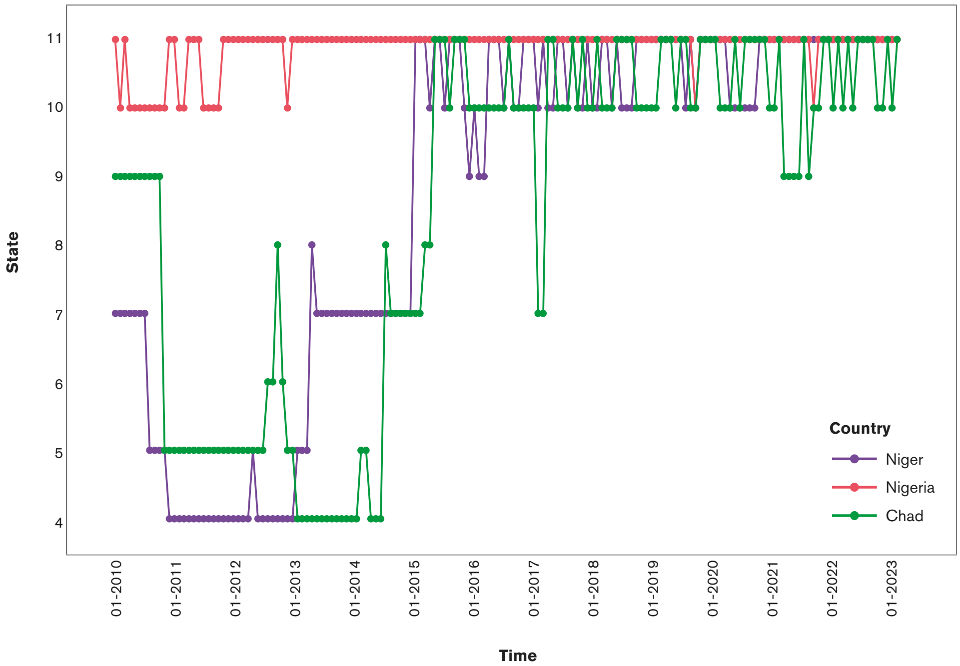

	State	1	2	3	4	5	6	7	8	9	10	11	12
Country	Nigeria	0	0	0	0	0	0	0	0	0	18	142	0
	Niger	0	0	0	25	8	0	27	1	3	27	69	0
	Chad	0	0	0	16	25	3	9	4	15	45	43	3
	Cameroon	0	21	3	25	5	0	1	1	0	4	94	6
	C.A. Republic	0	0	0	0	0	0	0	0	10	9	73	68

Most frequent state

2nd most frequent state

3rd most frequent state

Building block 2: Conflict dynamics

The next step is to model conflict dynamics via a **transition matrix.**[6] It represents the likelihood that a country moves from one state (the row) to any another (the column) within the next three months. This is the engine of our modelling approach since it permits going beyond static representations of the world to dynamic ones.

Figure 4: Transition matrix

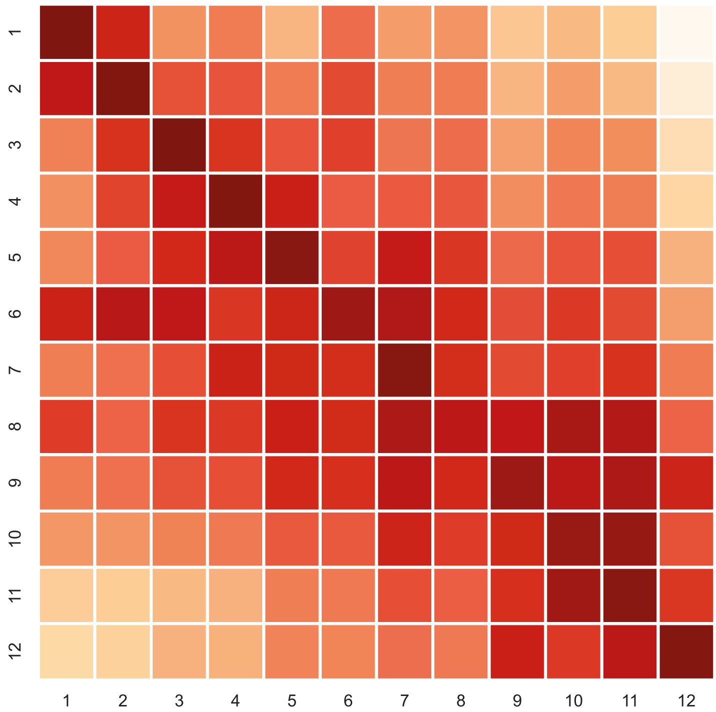

Figure 4 shows this transition matrix visually. The darker the shade of red, the higher the probability of this transition. We see that:

- **States are sticky:** Notice the deep, dark red hue across the diagonal of the matrix. States are most likely to transition to the same or similar state. For example, the likelihood of starting in state 1 and ending up in state 1 three months later is over 93%.

- **Peace is stable:** Notice the light orange area in the top right corner. This shows it is rare for countries to start in states 1, 2, 3 or 4 and transition to state 11 or 12. Instead, countries usually pass from stable peace into a period of elevated risk, rather than directly to violence. For example, imagine you are in state 4. The likelihood of moving to any of states 5-10 is 9%, whilst the likelihood of moving to any of states 11-12 is 0.2%.

- **The conflict trap:** Similarly, countries rarely transition from state 11 or 12 to state 4 or lower - the light orange area in the bottom-left corner. In order to escape conflict, they generally experience extended periods of elevated risk before settling into a peaceful state.

[6] The transition matrix is derived from the state transitions for all countries for the period January 2010 to April 2023.

Building block 3: Conflict damages

So far we have summarised the historical dynamics of conflict via conflict states and a transition matrix. The essence of this element is to capture how the conflict states are associated with measurable outcomes.[7] We focus on the relation between states and GDP growth, fatalities, displacement, and overseas development assistance (ODA).[8]

First consider a static view of the world as shown in Figure 5. The numbers in these charts should be interpreted as: **When I am in state X, what is the expected effect on outcome Y?** In general, a higher state is related to greater reductions in GDP growth. Fatalities become substantial in state 11, whilst countries in state 12 suffer extremely intense violence.

Figure 5: Static effect of states on outcomes

(a) Average GDP growth loss by state, yearly

(b) Average number of fatalities by state, monthly

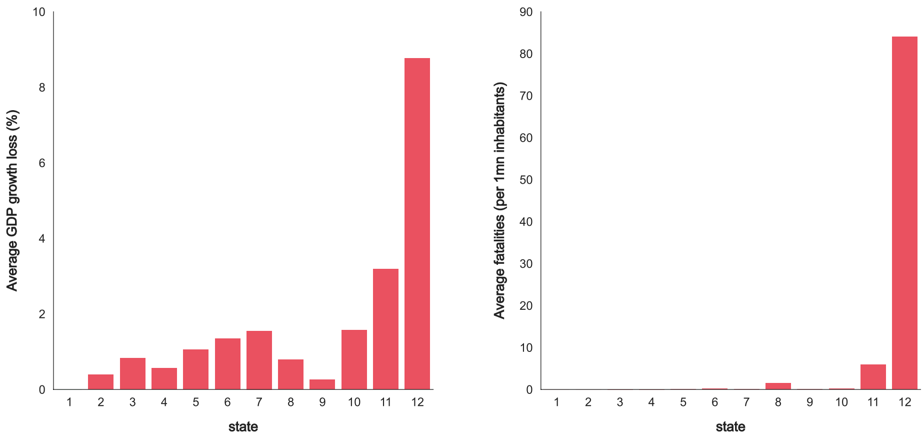

However, viewing conflict as a static issue is like playing chess without thinking ahead. The charts in Figure 6 show the same outcomes but from a dynamic perspective i.e. we account for all possible futures for every state. This answers the question: **When I am in state X, what is the expected effect on outcome Y in the future?** Take state 9 as an example, where in the static case we expect GDP to fall by only 0.02% and 0.1 fatalities per 1mn inhabitants. However, it is associated with significant future losses - GDP losses equivalent to 4x current GDP and 148 fatalities per 1mn inhabitants. Why? Because your future outlook contains a relatively high possibility of moving to state 11 or 12 (and getting stuck in the conflict trap) compared to states 1-4.

Figure 6: Dynamic effect of states on outcomes

(a) Present value of future economic losses (relative to current GDP)

(b) Present value of future fatalities (per 1mn inhabitants)

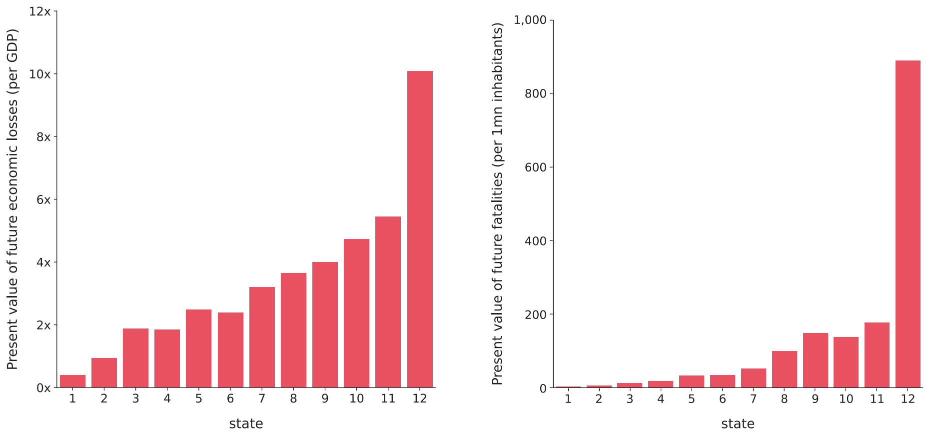

Figure 6 is therefore a way to communicate the key point of the DEWAM: acting now to prevent future damages requires us to understand when these future damages are coming closer, even if they are not visible yet.

[7] These associations are obtained via OLS regression analysis. We claim nothing causal about these relationships.

[8] ODA is categorised as spend only relating to emergency response and peace/security.

Building block 4: Effect and costs of policies

At this point we have modelled the dynamics of conflict via the states and their connection to measurable outcomes. The final step is to model policy interventions.

An intervention is defined via the transition matrix. Effective policies change the dynamics of conflict - they decrease the likelihood of moving to a higher state (towards conflict) and increase the likelihood of moving to a lower state (towards peace).[9, 10] This is portrayed visually in Figure 7:

Figure 7: Change in probability of transitions under policy

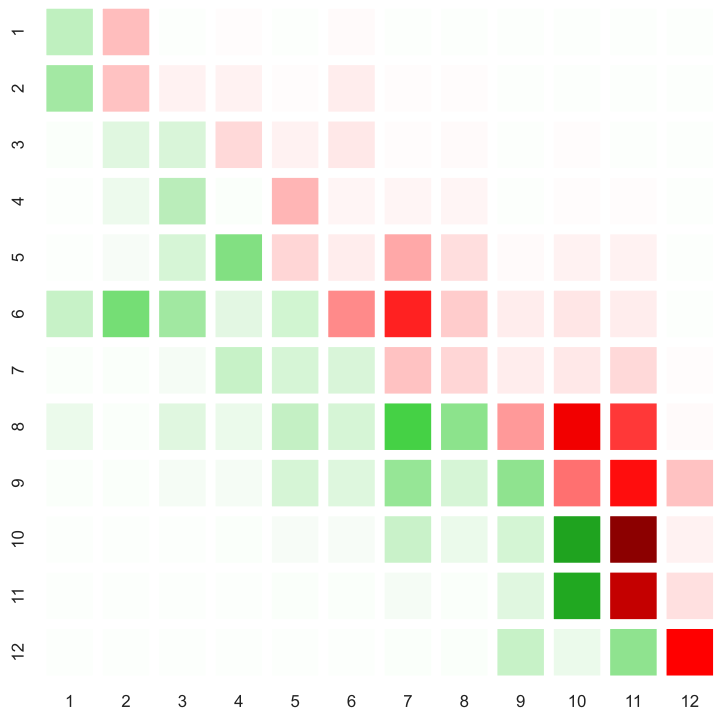

Lastly, we need to associate interventions with a financial cost in each state. This is challenging given the lack of available data on the monetary cost of different policies. Through close collaboration with policy-makers, we have derived cost estimates with two core assumptions:

- **Fixed cost:** Countries are assumed to be investing in policies to mitigate conflict, irrespective of what state they are in. This is proportional to the size of the population.[11]

- **Variable cost:** Costs are assumed to rise in line with the expected number of fatalities in each state.[12] In other words, the more deaths resulting from a hypothetical conflict, the higher the intervention cost.

For context, policies in state 4 interventions might relate to institution building at a cost of $145mn per month for a country with the population size of Nigeria. State 12 interventions would cost $10bn and would imply the mediation for a ceasefire, a massive deployment of peacekeeping troops and/or Disarmament, Demobilization and Reintegration (DDR) initiatives.[13]

[9] The simulations test a range of policy effectiveness levels (2, 5, 10 and 25 per cent). The more effective a policy, the more probability mass is moved.

[10] Keen observers may notice that the probability of staying in the same state does not change uniformly across the states. In the case that the total likelihood of moving to higher states exceeds the total likelihood of moving to lower states, the probability of staying where you are increases. The vice versa is also true.

[11] Expected spend of $0.25 per person.

[12] This requires making an assumption of how much costs rise for each additional fatality. As part of our simulations we test two values: $40,000 per fatality and $200,000 per fatality.

[13] Stated costs for Nigeria assume a variable cost assumption of $200,000 per fatality.

Results

Bringing all these aspects together allows us to simulate policy interventions. The results are reported as a benefit cost ratio (BCR), which can be interpreted as the dynamic long-run return per $1 spent.[14, 15] We also report uncertainty estimates in the form of standard deviations. A BCR less than 1 implies that the intervention is not cost effective.

Keep in mind that the BCR takes into account the imprecision of risk forecasts and the ineffectiveness of policies. This means that for low states large benefits of prevention need to compensate for low escalation risks to bring the BCR over $1.

Our method not only takes into account imperfect forecasts but also optimal future policies into all futures. In the calculation of prevention we take into account what a rational policy-maker would do in a future in which the situation escalates.[16] This might sound like a technical point but this ensures that the policy benefits are optimal, even under the assumption that the failure to prevent can be partially offset by later, optimal interventions.

Figure 8: Returns to interventions, Nigeria

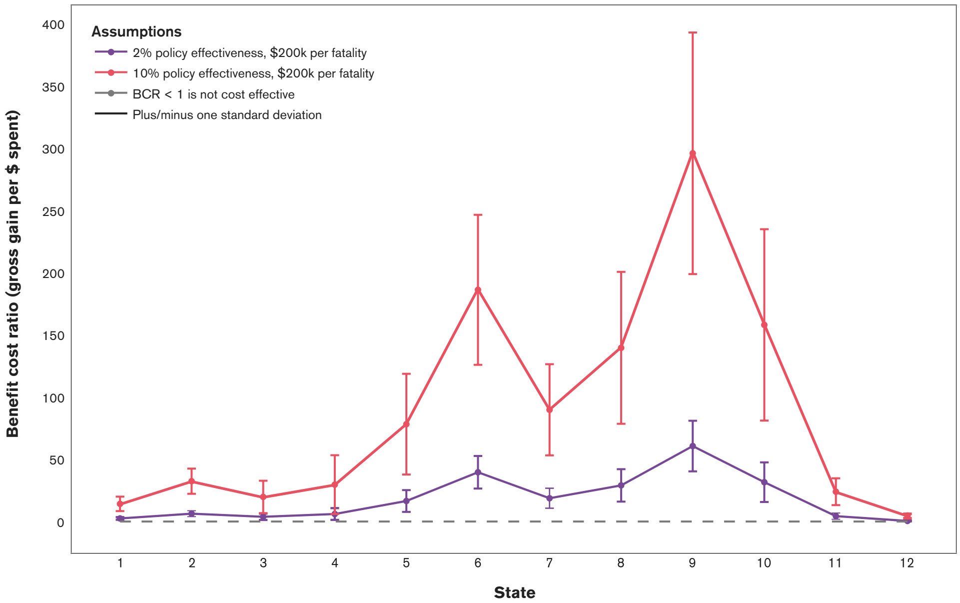

(a) Including GDP growth effect

Cost level	Policy effect / State	1	2	3	4	5	6	7	8	9	10	11	12
$200k pre fataility	2%	3.13	6.96	4.46	6.69	17.14	40.24	19.37	29.68	61.30	32.27	4.98	1.13
	10%	14.84	33.05	20.42	30.31	78.88	186.83	90.44	140.22	296.53	158.64	24.61	5.50

(b) Excluding GDP growth effect

Cost level	Policy effect / State	1	2	3	4	5	6	7	8	9	10	11	12
$200k pre fataility	2%	0.01	0.03	0.05	0.08	0.22	0.42	0.33	0.57	1.48	0.65	0.13	0.05
	10%	0.05	0.14	0.21	0.37	1.02	1.95	1.56	2.74	7.42	3.34	0.68	0.26

[14] The benefit cost ratio is computed as gross gains divided by intervention costs.

[15] Future gains/losses are discounted using a rate of 4%.

[16] We use dynamic programming to conduct simulations. This simulates an infinite number of future months and undertakes a full dynamic optimisation with respect to policy.

Results

Our model suggests that interventions in state 9 are the most cost effective. When including/excluding the effects on GDP growth, interventions in this state could deliver huge returns of $297/$7 per dollar spent. Irrespective of the policy and cost assumptions, returns are markedly lower in states 1-4 and 11-12 relative to 5-10. The intuition for this is as follows:

▪ **Interventions in stable peace (states 1-4):** Interventions are relatively cheap, but potential gains are also low. Think of this as intervening in Sweden - the risk of escalation is so low that policies are not worthwhile. In these cases, the forecast precision is relatively low.

▪ **The case for prevention (states 5-10):** States 5 to 10 represent a continuum of 'preventative' states. Levels of violence are low, but elevated risks of escalation and the huge costs paid in the higher states makes acting cost-effective. This is true even if it is not reacting to open violence, but merely follows an imprecise forecast.

▪ **Interventions in conflict (states 11-12):** Armed conflict is severe in these states leading to significant loss of life and economic growth, high displacement and a large increase in aid requirements. However, interventions have now become extremely costly - think of Afghanistan or Syria.

The webpage conflictforecast.org will publish a list of countries and their states together with their BCR. Just as an illustration, the following countries are currently (May 2023) in state 5: Djibouti, Eritrea, Lebanon, The Philippines, Russia and Rwanda.

Conclusion

We have devised a model that effectively integrates forecasts for evaluating returns to armed conflict prevention and de-escalation policies. This takes into account uncertainties faced by policy-makers including forecast precision, conflict dynamics, policy effectiveness and costs, whilst also capturing damages associated with conflict.

The core benefit of forecasts is that they permit actions to be taken in anticipation of crises. However, strong anticipation means acting on less precise forecasts. Taking action in low-risk environments necessarily requires acting on false positives, i.e. cases where you make an intervention that in retrospect was not necessary. Our model demonstrates that there is a distinct cut-off at which these preventative actions become cost-effective i.e. the benefit of anticipation outweighs the cost of acting on false positives.

One limitation of our approach is that we do not model spatial spillovers. In other words, conflict and the implementation of policies in one country are assumed to have no effect on conflict or outcomes in other countries. We recognise this is a simplification, particularly for countries where there are extensive histories of conflict around borders. We seek to address this as part of our work on modelling policies at the subnational level.

Extension: Subnational policy interventions

Our work with policy-makers has highlighted that policies tend to be specific and localised due to the heterogeneity of conflict drivers and their subsequent outcomes. This motivated the development of a method that enables simulating policy interventions at a more fine-grained geographic unit - the grid cell level. They are a convenient and frequently used method for disaggregating the world map into uniform geographic units.[17]

Figure 9: Grid cell state snapshot, Nigeria

(a) Grid cell states, May 2005

(b) Grid cell states, January 2023

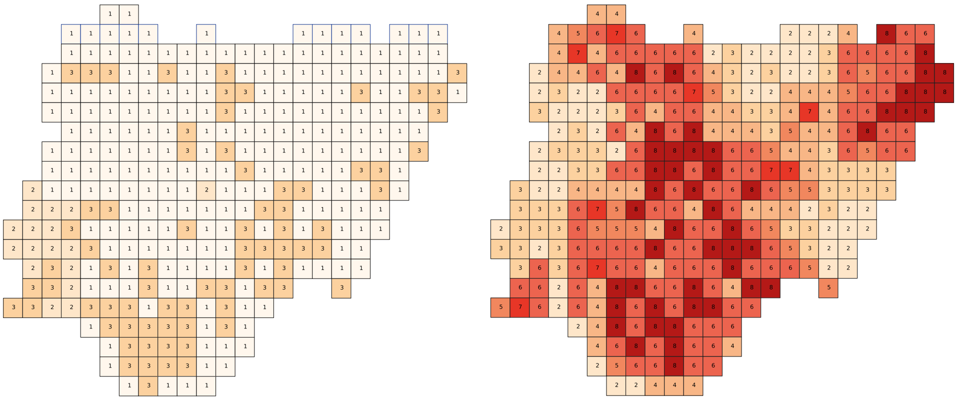

Predicted GDP growth (yearly): 6.4%
Predicted aid: $29mn

Predicted GDP growth (yearly): -1.6%
Predicted aid: $1bn

The building blocks of the model are identical to those used for simulating national policies. However, some modifications are necessary to ensure simulations represent reality as closely as possible:

▪ **Conflict states:** We now have 8 states and every grid cell is assigned a state for every month.[18] States 7 and 8 are the only situations in which violence occurs in a grid cell. Yet, they are differentiated by the amount of violence in directly neighbouring cells. Notice how in Figure 9b state 8 appears in constellations (violent hotspots) whereas state 7 tends to be more isolated (isolated outbreaks).

▪ **Conflict dynamics**: We explicitly model spatial spillovers. Levels of violence in one grid cell directly affect the transition likelihoods of neighbouring cells.

▪ **Conflict damages:** A lack of high-quality data makes accurately tracking outcomes (e.g. economic growth) at the grid cell level impossible. Instead we make a connection between levels of regional violence and measurable outcomes at the national level.[19] In general, the higher the proportion of a country in more risky/violent states, the worse the expected outcome as highlighted in Figures 9a and 9b.

[17] Each grid cell is approximately 55km x 55km.

[18] The latest subnational model is trained on data from January 2010 to January 2023.

[19] Specifically we use shares of population and GDP by grid cell. In this way we can compute the total share of a country's population/economic activity that is in a given state at a given point in time. OLS regressions are then used to associate these shares to measurable outcomes (GDP growth, fatalities, displacement and ODA) at the national level.

Extension: Subnational policy interventions

We simulated potential gains over the next 12 months from a policy to de-escalate violence in Borno, Nigeria.[20] This serves to increase the likelihood of each grid cell in Borno transitioning to peace. In each month and for each grid cell, the intervention is assumed to be successful 25% of the time i.e. triggers movement to a lower state. The rest of the country follows a state trajectory unaffected by the policy, although a pacification of Borno has spillover effects to neighbouring regions. Our model suggests that this policy could save up to 400 lives and prevent $1.4bn in GDP losses between Jan 2023 and Jan 2024 in Nigeria.

Figure 10: Policy intervention in Borno, Nigeria

(a) Grid cells for policy intervention

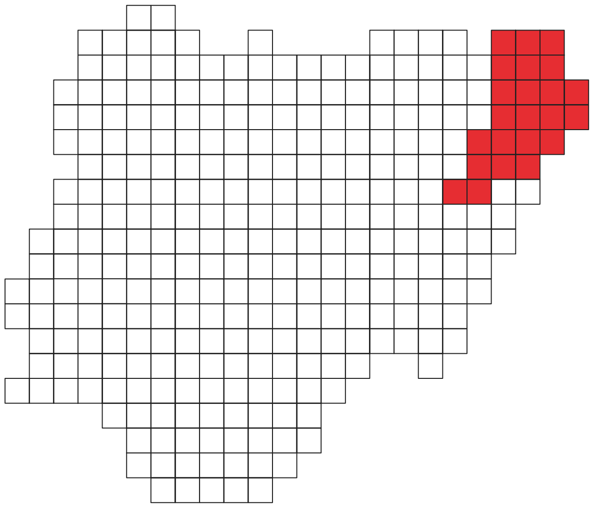

(b) Indicative number of lives saved

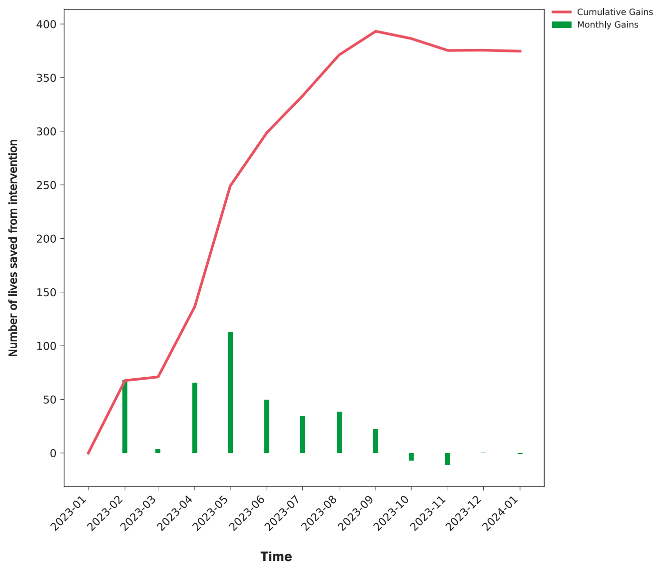

This is a significant development from modeling national policy interventions. It enables evaluating policies which are actually taking place on the ground. Comparing and contrasting the value of different interventions is crucial, and our model can answer questions such as:

- Which regional parts of a country are vulnerable to escalation? What gains could we expect from preventing outbreaks in different regions?

- Is it more beneficial to de-escalate already violent hotspots, or act to prevent these hotspots spreading further?

- How can the timing and location of interventions around borders affect outcomes?

[20] Note that the methodology and results are still experimental. At present we do not assign costs to interventions in given sub-regions, hence results are presented as total gains rather than a benefit-cost ratio. The results shown assume a policy effectiveness of 25%.

Appendix

Four Models for Integrating Forecasting into Policymaking

Figure 11: FP21, Forecasting in Policymaking: Beyond Cassandra, Jan 2022

The Model	What is it?	Strengths	Weaknesses
1. Analytical	**Accurate forecasts are made available to policymakers.**	Least disruptive. Exposes forecasting to policymakers with low stakes.	Unlikely to meaningful change policymakers thinking. Policymakers may cherry pick favorable forecasts. Policymakers not involved in learning process.
2. Early Warning	**Forecast are embedded into early-warning systems focused on high priority issue areas.**	Makes bureaucratic sense. Useful foothold to expand if techniques are useful.	Same weaknesses as first model. Constrains forecasting to discrete issue areas.
3. Policy Evaluation	**Forecasters evaluate likelihood of success for discrete policy options.**	Policymakers incentivized to engage with underlying forecasting logic. Requires policymakers to identify goals of their policies. Shifts gravity of policy debate to evidence, away from ideology. Foundation for active learning.	Will threaten authority of policymakers. Requires more resources, slows down policy process. Jeopardizes objectivity of forecasting. Accuracy of conditional forecasting not well studied.
4. Decision Making	**All policies are presented as testable forecasting questions; forecasting methods supplant the existing process.**	Integrates forecasting at every level of the process. Creates conditions for evaluation of policymaker merit and accountability. Potential for transformative impact on foreign policy.	Unviable in terms of resource requirements and depth of organizational changes. Much more research and policy experience needed to advance this model.

Appendix

Our forecast usus millions of newspaper articles to make the conflict forecast. In our analysis of the content of the newspaper articles we rely on a so called topic model which summarizes the millions of articles and words into topics using unsupervised machine learning. The topic model allows us to calculate 15 topic shares for each country/month we display in the bubbles to the right.

The model assumes that if journalists write about a topic, say politics, they will use a different vocabulary than if they write about another topic, say conflict. The word “congress”, for example, is more likely to be used in an article about politics than in an article about conflict.

Our forecast method re-estimates the topic model every month and so the content of the topics changes across times.

Figure 12: Forecasting module topics, Sudan, April 2024

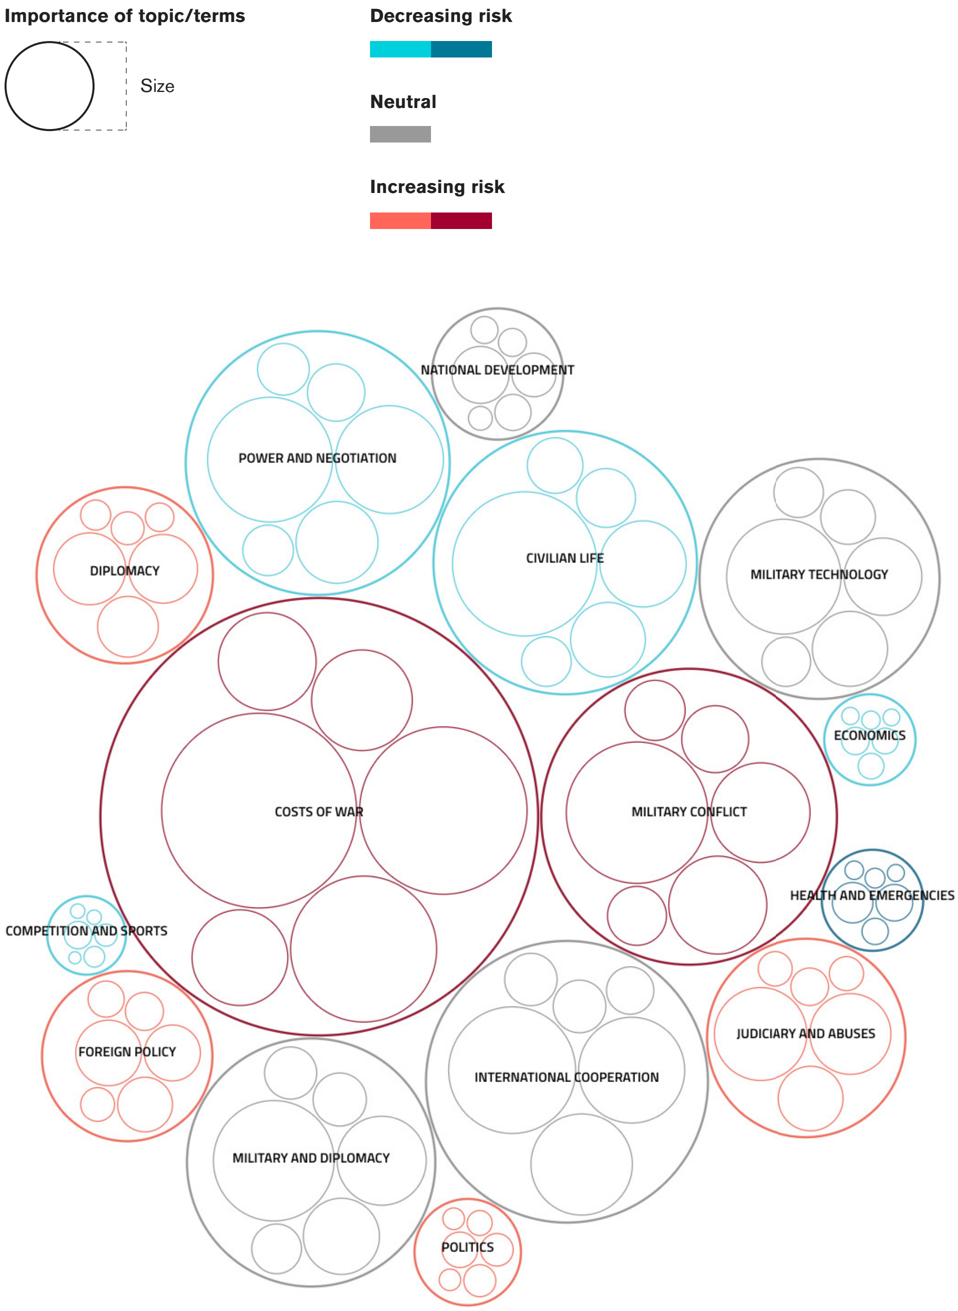

Appendix

Figure 13: Example of topic-specific vocabulary, April 2024

Cost of war

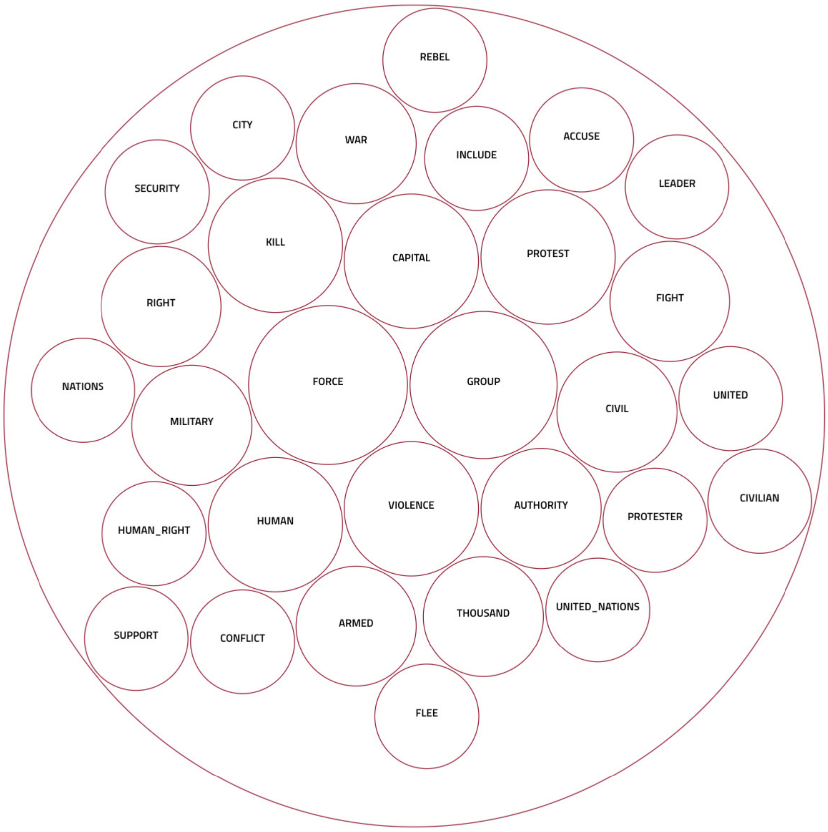

Power and negotiation

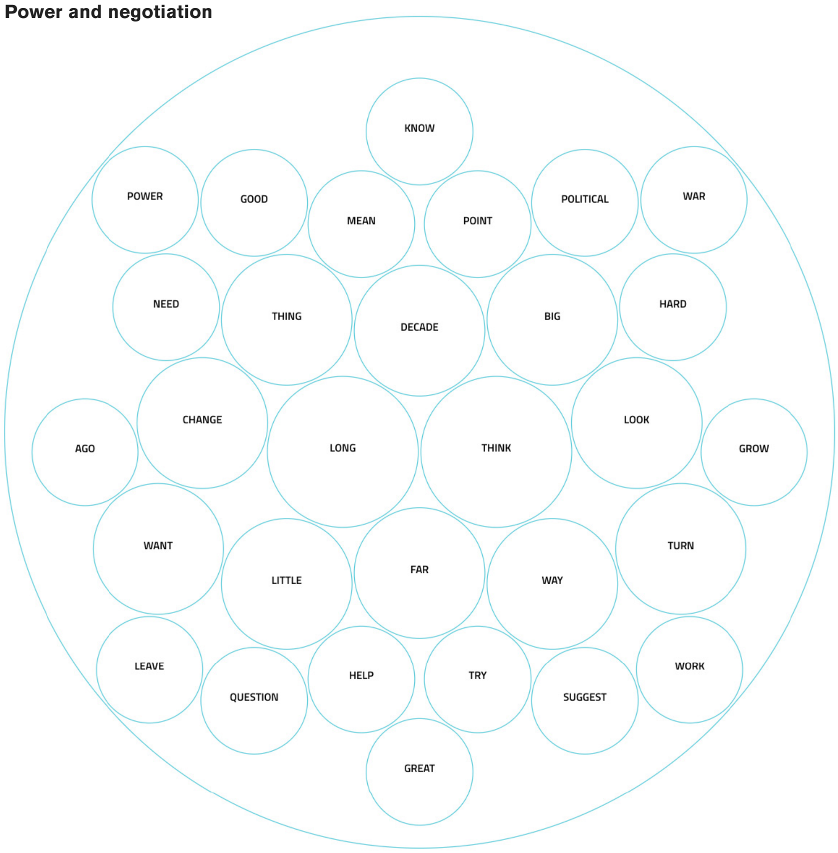

Conflict Forecast

UKaid from the British people

GOBIERNO DE ESPAÑA MINISTERIO DE CIENCIA E INNOVACIÓN

CSIC Consejo Superior de Investigaciones Científicas

More info: conflictforecast.org

Say hello: conflictforecast@gmail.com