"If You Only Have a Hammer": Optimal Dynamic Prevention Policy*

Hannes Mueller Christopher Rauh Alessandro Ruggieri
Ben Seimon

January 2, 2025

Abstract

How will better forecasting change policy decisions in high-stakes scenarios such as armed conflict prevention? We study the gains of improving forecasting for a policymaker who faces a recurring risk and has the choice between preventive early actions and de-escalating late actions. We build a Markov model where early and late interventions are the solution to an optimal stopping problem and the timing of interventions depends on the ability to forecast future states. Quantitatively, we study the role of forecasting for armed conflict prevention using a large panel of countries. The benefits of prevention are substantial but critically depend on the systematic use of forecasting. The information rent of using a forecast is larger than 60% of GDP. In line with the theory, we find that de-escalation policies reduce the incentives for prevention, whereas prevention increases incentives for de-escalation.

*Corresponding author: Mueller, h.mueller.uni@gmail.com. Mueller: IAE (CSIC), Barcelona School of Economics and CEPR; Rauh: IAE (CSIC), University of Cambridge, Barcelona School of Economics, CEPR, and PRIO; Ruggieri: CUNEF Universidad; Seimon: Fundació d'Economia Analítica. We acknowledge financial support from two UK FCDO grants, German Federal Foreign Office funding, the European Research Council project ERC-AdG 101055176 (ANTICIPATE), the Keynes Fund Cambridge, and IMF project funding. Mueller acknowledges support from the Spanish Government and the Spanish Ministry of Economy and Competitiveness, through the Severo Ochoa Programme for Centres of Excellence in R&D (CEX2019-000915-S) and the Ministry of Science Innovation and Universities (PGC2018-096133-B-100). Rauh acknowledges support from AEI/MICINN (ATR2023-144291). We are deeply grateful for help on an earlier version of this project by Rafael Hortala-Vallve. We thank Bruno Conte Leite, Luis Ignacio, and Alexandra Malaga for excellent research assistance. Finally, we thank policymakers for their openness in numerous conversations with us and seminar participants in Cambridge, Bank of Spain, Barcelona, Bologna, Lucerne, and Navarra for their feedback. All errors are our own. The title refers to the quote by Abraham Maslow: *"If the only tool you have is a hammer, you tend to see every problem as a nail."*

JEL Classification: O1, F5, F1, L8
Keywords: prevention, forecasting, machine learning, dynamic optimization, armed conflict, power sharing, information gain, artificial intelligence

1 Introduction

Prediction-policy problems are common. From high inflation, banking crisis, epidemics, crime, and climate disasters to armed conflict, predicting the right timing and targeting of a policy could be as important for overall efficacy as the average treatment effect of the chosen policy tools conditional on good targeting (Kleinberg et al. 2015). The last decade has seen an explosion in the capabilities of artificial intelligence (AI) to achieve better predictions thanks to model improvements and new sources of data. However, the use of predictions for policy targeting poses challenges, and public policies are typically not evaluated in systems that integrate quantitative forecasts based on machine learning. This raises the questions of what the gains of having a forecast are. Further, we do not know how a forecast affects trade-offs between late and early action.

This article is an attempt to help fill this gap. It proposes a framework for integrating forecasts into a dynamic model of decision making that can be solved with standard optimization tools and applies it to policies concerning armed conflict. We build the framework from forecast data which we summarize using a Hidden Markov Model. The fundamental building blocks of our model are latent Markov stages, each with a different payoff, and a transition matrix, capturing the dynamics across stages.[1]

Picture a benevolent policymaker who decides in the dynamic setting to intervene depending on the observed two stages: the good and the bad stage. We model policies as costly actions that increase the likelihood of staying in or transitioning to the good stage. More specifically, we define *prevention* as a reduction in the probability of transitioning from the good to the bad stage and *de-escalation* as increasing the probability of transitioning from the bad to the good stage. Using the most parsimonious model, we show theoretically that, ceteris paribus, prevention tends to crowd in de-escalation but de-escalation can crowd out prevention.

Then we treat an expansion of this framework as a laboratory to analyze quantitatively the role of forecasting in policymaking. A natural application is the prevention and de-escalation of armed conflict. We model forecasts in this framework as a differentiation of stages into sub-stages with differing probabilities of transi-

[1]We refer to Markov states as stages rather than states to avoid confusion with nation-states in the context of conflict.

tioning to bad stages. Under *full information* the policymaker can differentiate the true dynamics between stages and can act accordingly. The policymaker under *partial information* cannot observe the full set of the true sub-stages. Based on the different choices resulting from these two information sets, we define an *information rent* as the difference in resulting welfare.

Conflict prevention is regarded as one of the key policy challenges of the international community in the 21st century. Increasing evidence in the academic literature points towards the fact that avoiding or escaping the so-called *conflict trap* would have large developmental effects.[2] But, in his first address at the Security Council, the United Nations Secretary-General (2017) emphasized that it had proven very difficult to persuade decision-makers to make prevention their priority. In a world experiencing an increasing number of open and intense armed conflicts, it is difficult to muster the resources for prevention as this means spending resources upfront to avoid humanitarian disasters in an uncertain future.

In light of the quasi absence of conflict prevention efforts in the real world, we bring our framework to the data under the assumption that policymakers currently operate under partial information. Though conflict risk during peace is generally relatively low, some peaceful stages are characterized by higher conflict risk than others. We estimate dynamics in the true model using monthly conflict forecast data for 168 countries over the period 2010-2024 in a Hidden Markov Model.[3] We use a specific observed conflict de-escalation policy - the adoption of power-sharing agreements - to calibrate the partial equilibrium model. Power-sharing agreements are often backed externally through mediation, foreign aid, and peacekeeping forces and have been shown to be relatively effective in reducing violence. However, they are almost exclusively implemented in high risk stages with open or recent armed conflict. This gives us the opportunity to estimate policy effectiveness levels and some aspects of intervention costs from these policy bundles. To derive an estimate of the prevented conflict damages, we associate the different conflict risk stages in the model to economic growth, refugees, Overseas Development Assistance (ODA), and fatalities. This gives us a full description of stages, transitions, and damages under partial information.

[2] See Collier et al. (2003), Rohner and Thoenig (2021) and Margalef and Mueller (2023).

[3] The webpage `https://conflictforecast.org/` provides monthly updates of forecasts with a yearly horizon.

We then study the policymaker's decisions under full information and demonstrate that, under conservative assumptions about intervention costs, current forecasting systems are able to generate cost effective opportunities for early intervention. Benefits from prevention are larger than their costs and approach the benefit-cost ratio of interventions in open conflict. We also find strong evidence for the mechanisms we identify in the theory; namely de-escalation in conflict dampens the incentives for prevention, whereas prevention strengthens the incentives for de-escalation. Finally, we derive the information rent. The differentiation into low and high-risk sub-stages increases the incentives to prevent transitions into high-risk stages with a profound impact on welfare: we find an information rent from forecasting of 60% of GDP with a standard deviation of 16%.[4]

Our paper builds on a large literature studying conflict. The first strand is a literature studying the causes of conflict and derives policy implications for preventing or de-escalating it.[5] In this sphere our work relates most closely to work studying the role of political institutions and political power sharing.[6] In a theoretical approach, Laurent-Lucchetti et al. (2024) analyze the informational rent of democracy in the context of bargaining between ethnic groups. Instead, we evaluate the dividend generated by access to information within the context of conflict prevention. A second large literature is concerned with forecasting conflict.[7] We bring these two branches together and model forecasts to optimally target policies. The literature most closely in spirit to our work are recent quantitative approaches that model economic development, policy, and conflict jointly. Couttenier et al. (2024) focus on the spatial interaction between violence and development. In contrast, our focus is on the incentives for prevention versus de-escalation, and how forecasts interact with these incentives. Thoenig (2023) studies the interplay between trade and diplomatic negotiations for which strategic behavior is of first order. Our model contains no strategic interactions but introduces a new model of

[4]This may be a lower bound given that in terms of conflict damages we do not model the potential costs a conflict might have on neighboring countries, which have been estimated to be substantial (Federle et al. 2024).

[5]For overviews see Blattman and Miguel (2010), Blattman (2023), Rohner (2024).

[6]See Acemoglu and Robinson (2001), Besley and Persson (2011b), Francois et al. (2015), or Mueller and Rauh (2024).

[7]See Goldstone et al. (2010), Chadefaux (2014), Mueller and Rauh (2018), Bazzi et al. (2022), Hegre et al. (2022).

how forecasts shape the information set and optimal policy.

The calibration of our model and stage space is specific to conflict. Nonetheless, we believe that the framework can be applied elsewhere. The policy prediction problem in the context of conflict is similar to those in climate change, macroeconomic policy, and financial crisis management. Weitzman (2007) discusses how the presence of tail risk calls for earlier rather than later action in the context of climate change. In macroeconomics, Bernanke and Mishkin (1997) discuss extensively how targeting inflation requires reliable forecasts in order to time policies correctly.[8] Similarly, financial regulations aim to preempt crises, while post-crisis measures often involve costly bailouts and interventions. This dynamic is analyzed in works by Bianchi (2016) and Jeanne and Korinek (2020), which examine the balance between crisis prevention and response. One aspect that makes these problems similar is that the literature is increasingly experimenting with machine learning when forecasting but the resulting forecasts have to date not been introduced into models to derive optimal policy.[9] We contribute a dynamic framework with policy interventions that builds on forecasts derived from machine learning models.

The prediction model relies on topics derived from six million newspaper articles.[10] The economics literature has long integrated data gained from text to model expectations and uncertainty. Our project is related to the literature on news shocks in Macroeconomics.[11] Ramey (2011) constructs a government spending indicator derived from news to capture changes in the present value of government spending. Auerbach and Gorodnichenko (2012) argue the fiscal multiplier depends on the general state of the economy. This leads them to model different policy re-

[8]Beaudry et al. (2023) provide insights into this discourse, underscoring the implications of delayed response to persistent inflationary pressures. Acharya et al. (2022) argue that while minor shocks may be absorbed without aggressive policy measures, more substantial disturbances necessitate preemptive action to avert a transition to an undesirable steady state—a perspective echoed in the work of Fatás and Singh (2024).

[9]Central banking, for example, has extremely similar policy debates to armed conflict prevention but forecasting using machine learning is most heavily used in feature generation and forecasting (Joseph et al. 2024, Liu 2024). Existing studies do not typically conduct out-of-sample forecasts in models that analyze policy effects - a critique that goes back to at least Sims (1980).

[10]The data is described in detail in Mueller et al. (2024a).

[11]For a review see Section 4 of Ramey (2016). See Tetlock (2007), Baker et al. (2016), Hansen and McMahon (2016), Hassan et al. (2019), Ochs (2021), and Aruoba and Drechsel (2024) for literature using text to capture risk and uncertainty.

sponses for recessions and expansions.[12] Having monthly forecast data from several machine learning models based on news text for 168 countries allows us to model a rich latent space with twelve Markov stages and we can compare what it means to have access to the news-based forecast which helps distinguish some of the least risky of the twelve stages.

This paper proposes a new approach to prediction-policy problems through our model of forecasts as sub-stages. In the terminology of Kleinberg et al. (2015), our policy experiment is only changing the prediction quality available to the policymaker, not the available treatments.[13] We show that this change in the information environment can lead to a dramatic shift in implemented policies and large economic gains.

2 A Stylized Model of Policy Interventions

In this section, we present a simplified model of policy interventions. This model mainly illustrates the insights one can gain from using a Markov chain model to analyze the dynamic benefits of policy interventions.

We start with a two-stage model featuring a "bad" stage and a "good" stage with dynamics governed by transition likelihoods. The key novelty is to assume that policies affect the transition likelihoods, i.e. target the dynamics of the system. Using a Markov chain allows us to model the stochastic transitions between more and less desirable stages in a straightforward manner, while also enabling the simulation of the long-term consequences of policy interventions. Throughout we will use terminology motivated by our application of civil war prevention. We conceptualize the bad stage as war and the good stage as peace, reflecting the conditions policymakers seek to influence. Also, we will use the terms *state* and *stage* interchangeably.

We then expand the model to a more general framework suitable for our quantitative applications. In Section 2.3 we introduce a model of forecasts in this context and the different information sets this introduces. Forecasts can be modelled

[12] Ochs and Rörig (2022) and Goulet Coulombe (2024) use random forests to capture the state-dependent effects of policies.

[13] This is a simplification, we discuss how targeting and treatment effects are not strictly separable.

through the introduction of additional sub-stages as they allow policymakers to distinguish situations that are indistinguishable without the forecast. Based on this model we introduce the idea of an *information rent*.

2.1 Two-stage Model

A country can be in one of two stages s: peace ($s = 1$) or conflict ($s = 2$). Transitions between stages are defined by a Markov chain. If a country is in peace, it may stay in peace with probability p_{11} while it might experience an outbreak of conflict with probability $(1 - p_{11})$. If a country is in conflict, it may stay in conflict with probability p_{22} while it might transit back to peace with probability $(1 - p_{22})$. In other words, p_{11} is the persistence of peace, and p_{22} is the persistence of conflict.

The policymaker can engage in interventions I. Interventions during peace are called prevention ($\phi_1 = 1$), whereas those during conflict are de-escalations ($\phi_2 = 1$). Interventions, summarized by the policy vector (ϕ), come at cost I_s depending on the stage s.

To capture the idea of prevention we assume that peace is associated with no damages while conflict is associated with damages which provides a payoff of $-D$ per period spent in that stage. Summarizing damages in vectors gives

$$\mathbf{D} = \begin{bmatrix} 0 \\ -D \end{bmatrix}.$$

The key assumption of our model is that the policymaker has a model of the real underlying dynamics of the world which can be described by a Markov chain model in which transitioning in and out of conflict are stable transition likelihoods. The likelihoods are given by past conflict dynamics and a forecast system and are assumed to be fixed.

In a simple two-stage case the transition matrix without intervention $\mathbf{T}_n$, i.e., $\phi_s = 0 \ \forall \ s$, can be written as

$$\mathbf{T}_n = \begin{bmatrix} p_{11} & 1 - p_{11} \\ 1 - p_{22} & p_{22} \end{bmatrix}.$$

As outlined before, we call *prevention* an action that increases the probability of peace tomorrow when the country is in peace today. We call *de-escalation* an action that increases the probability of peace tomorrow when the country is at war today. A more general way to describe these two different kinds of policies is to say that prevention moves probability mass from the more costly stage in the future towards the less costly stage in the present. De-escalation moves probability mass from the costly stage in the present towards a less costly stage in the future. We will call both actions an *intervention*. Formally, we can write the transition matrix $\mathbf{T}(\phi_1 = 1, \phi_2 = 0)$ with prevention as follows:

$$\mathbf{T}(\phi_1 = 1, \phi_2 = 0) = \begin{bmatrix} p_{11} + \tau\,(1 - p_{11}) & (1 - p_{11})\,(1 - \tau) \\ 1 - p_{22} & p_{22} \end{bmatrix}.$$

where $\tau \in [0, 1]$ denotes the effectiveness of the intervention. Note that we assume that prevention shifts probability mass proportionally to the existing transition probabilities without intervention. Specifically, we assume that preventive action follows the logic of reducing escalation likelihoods proportionally so that if the risk is higher ($1 - p_{11}$ is large) more probability mass will be shifted.

This assumption can be justified by the fact that actions are taken before transitions are realized. An action that replaces the lottery of the transition matrix with the certainty of staying in peace but only works with probability τ yields expected costs $(1-\tau)(p_{11}V_1 + (1-p_{11})V_2) + \tau V_1 = (p_{11}V_1 + (1-p_{11})V_2) - \tau(1-p_{11})(V_1 - V_2)$ which is illustrating that the effect of τ is proportional to $1 - p_{11}$. Avoiding the lottery has larger benefits if it has higher downside risks.

De-escalation increases the probability of stage 1 tomorrow when the system is in stage 2 today:

$$\mathbf{T}(\phi_1 = 0, \phi_2 = 1) = \begin{bmatrix} p_{11} & 1 - p_{11} \\ (1 - p_{22})\,(1 + \tau) & p_{22} - \tau\,(1 - p_{22}) \end{bmatrix}.$$

Note, that this policy also moves probability mass towards stage 1, an asymmetry which will be important. Again we assume that the intervention has a larger effect if the baseline likelihood of de-escalation is higher. We also assume that the same parameter τ governs effectiveness. Again, the idea is that de-escalation helps avoid the lottery of staying in the bad stage and amplifies the probability of moving to the good stage. We use the normalization of one parameter τ linked to the de-escalation baseline likelihood $(1 - p_{22})$ to discipline our analysis.[14] This makes our model sensitive to the relative costs of prevention vs. de-escalation. Overall, the relative cost-effectiveness of interventions is also governed by the costs of these two types of interventions given by the vector:

$$\mathbf{I} = \begin{bmatrix} I_1 \\ I_2 \end{bmatrix}$$

where $I_1 > 0$ and $I_2 > 0$. Write the present discounted value of stage s as

$$V_s = \max\{V_s^i, V_s^n\}, \quad s \in \{1, 2\},$$

where V_s^i and V_s^n are present discounted values conditional on steady-state policies, i.e. intervention and non-intervention, respectively. Let $\beta \in (0, 1)$ be a discount factor. Then we can write:

$$\begin{aligned} V_1^n &= \beta \left(p_{11} V_1 + (1 - p_{11}) V_2 \right) \\ V_1^i &= -I_1 + \beta \left((p_{11} + \tau (1 - p_{11})) V_1 + (1 - p_{11}) (1 - \tau) V_2 \right) \end{aligned}$$

and

$$\begin{aligned} V_2^n &= -D + \beta \left((1 - p_{22}) V_1 + p_{22} V_2 \right) \\ V_2^i &= -D - I_2 + \beta \left((1 - p_{22}) (1 + \tau) V_1 + (p_{22} - \tau (1 - p_{22})) V_2 \right) \end{aligned}$$

where we assume that τ is small enough to ensure that $(1 - p_{22})(1 + \tau) < 1$ and $p_{22} - \tau (1 - p_{22}) > 0$.

The distribution of probability mass is visible in both intervention scenarios. Under prevention, the probability mass $\tau (1 - p_{11})$ is redistributed from stage 2 to stage 1. Under de-escalation, the probability mass $\tau (1 - p_{22})$ is redistributed from stage 2 to stage 1. Note that the intervention cost is paid before outcomes are realized because decision for interventions is taken before the transitions materialize. The trade-off is therefore an exchange of a certain cost for changes in probabilities p_{11} in preventive action and p_{22} in interventions. This assumption implies that the

[14] Note, we now apply the intervention effect τ as an increase in the likelihood of moving to peace $(1 - p_{22})$. We do this because applying the intervention to the persistence of conflict p_{22} would lead to an unrealistically powerful intervention in practice as p_{22} is very high typically.

policymaker enacts policies to change the course of conflicts, instead of reducing conflict damages (like humanitarian aid).

The assumption that the probability mass redistributed with de-escalation is proportional to $(1 - p_{22})$ and not the persistence of conflict p_{22} merits discussion. We make this assumption to level the playing field between prevention and de-escalation effects with many stages. When expanding to many stages, in intermediate stages we have both a preventative and a de-escalation effect of interventions. But stages are also extremely persistent and if policy effectiveness were assumed to be proportional to this high persistence, the de-escalation effect would always dominate mechanically. The most intractable conflicts, i.e. those with high persistence would be where de-escalating policy is most effective. Instead, we make de-escalation more effective where the default already suggests that there are ways out of the conflict, i.e. where $(1 - p_{22})$ is higher in the two-stage model.

2.1.1 Solution Characterization

Before moving on to our main results it helps to look at the trade-offs posited by the model. Exploring the value functions, it is easy to show that the decision-maker wants to prevent, if and only if

$$I_1 < \beta\tau\left(1 - p_{11}\right)\left(V_1 - V_2\right) \tag{1}$$

and wants to de-escalate iff

$$I_2 < \beta\tau\left(1 - p_{22}\right)\left(V_1 - V_2\right). \tag{2}$$

These conditions illustrate the simple intuition of the model. Policies are implemented if the certain intervention costs I_1 or I_2 lead to a change in expected benefit which is large enough. The expected benefit of interventions is determined by the change in probability mass ($\tau\left(1 - p\right)$ in prevention) multiplied times the change in present value between war and peace $(V_1 - V_2)$. We will assume that it is always true that $(V_1 - V_2) > 0$, i.e. it is always better to be in peace than in war regardless of the interventions taken in equilibrium.

Solving for the present values without either de-escalation or prevention and inserting them into the optimal intervention conditions (1) and (2), we get[15]

$$I_1 < \beta\tau\left(1-p_{11}\right)\frac{D}{1+\beta-p_{11}\beta-p_{22}\beta} \tag{3}$$

and the necessary condition for de-escalation without prevention is

$$I_2 < \beta\tau\left(1-p_{22}\right)\frac{D}{1+\beta-p_{11}\beta-p_{22}\beta}. \tag{4}$$

The equivalent conditions with prevention and de-escalation are derived in the Appendix. The intuition is that de-escalation introduces the intervention costs I_2 in Equation (3) and reduces the persistence of war. Similarly, prevention introduces cost I_1 in Equation (4) and increases the persistence of peace. We will now use these results to highlight important elements of prevention and de-escalation and their interactions.

2.1.2 The Role of the Transition Matrix

One could imagine that prevention becomes unambiguously more likely with an increase in the likelihood of an outbreak of conflict $(1-p_{11})$. Yet, the role played by conflict risk during peace $(1-p_{11})$ for prevention in this model is surprisingly subtle. To see this, note that an increase in $(1-p_{11})$ will unambiguously reduce the difference $V_1 - V_2$ as it makes the present value of peace V_1 relatively worse. This disincentivizes prevention. More intuitively, higher risk also means that prevention moves more probability mass towards peace. If a country has been identified as high risk, less resources will be wasted on false positives where the forecast $(1-p_{11})$ never comes true. The trade-off between these two forces gives rise to our first proposition.

Proposition 1. *High conflict risk during peace, $(1-p_{11})$ makes prevention more likely. A high baseline likelihood of moving from conflict to piece $(1-p_{22})$ makes de-escalation more likely.*

It is straightforward to check this proposition simply from the first derivative of Equation (3) with respect to p_{11}, which is always negative and also holds with

[15] We derive the value functions V_1 and V_2 under different intervention policies in the Appendix.

commitment to de-escalation. The intuition for these results is simply that the effect of redistributing probability mass always dominates. Under the assumptions we made regarding the proportional policy effects, a higher risk of an outbreak of violence $(1 - p_{11})$ means that more probability mass can be moved by prevention and it is therefore always more attractive to intervene when risk is higher. Note that the de-escalation result suggests that a higher baseline likelihood of moving from conflict to peace leads to increased incentives to de-escalate. This result is driven by the fact that policy efforts have a higher effectiveness if they are targeted toward an achievable goal.

2.1.3 Crowding out Prevention

The simple decision rules in Equations (1) and (2) contain the main intuition for how optimal policies interact dynamically. Prevention will tend to make stage 1 more attractive by increasing persistence in peace p_{11}. The persistence of peace increases the distance $V_1 - V_2$. This in turn will make both prevention and de-escalation more attractive. The opposite is true for intervention as it will tend to increase V_2 which will reduce the difference $V_1 - V_2$.

Both these present values are negative, driven by the negative costs of conflict $-D$. There is a simple ranking of $V_1 > V_2$ for $\beta < 1$. Intuitively, the persistence of conflict (p_{22}) makes both present values worse for $p_{11} < 1$. This is because being stuck in conflict for longer leads to higher costs in expectation and as long as there is a danger of an escalation ($p_{11} < 1$) this also affects the present value of peace. The impact of the persistence of peace p_{11} is less obvious from the formulas, but it is easy to show that it positively impacts values. In the Appendix, we insert the present values under different assumptions regarding prevention or de-escalation into condition (3) and then compare the resulting conditions for optimality. This gives rise to our second proposition:

Proposition 2. *De-escalation crowds out prevention in the steady state. For parameter values for which the policymaker wants to engage in intervention in the steady state, there is a subset of parameters where the policymaker only wants to engage in prevention if there were commitments not to de-escalate.*

The proof of this proposition is intuitive given the previous observations. The goal of interventions is to reduce $(1 - p_{22})$ and the steady-state value of V_2 increases with reductions in $(1 - p_{22})$. This means that conflict becomes a "less terrible" stage to be in and, optimally, the effort to prevent this stage can decrease. Preventive action becomes less useful if de-escalation is used. The substitution effect cannot be found in the other direction. If we insert the respective present values into Equation (4) we get our third proposition:

Proposition 3. *Prevention crowds in de-escalation in the steady state. For parameter values for which the agent wants to engage in prevention, there is a subset where the policy-maker only wants to engage in de-escalation because he/she knows there will be prevention in the future.*

The asymmetry comes from the fact that a high value of p_{11} makes being in peace more attractive - the present value of peace, V_1, increases with prevention. This means that bringing a country from war into peace through de-escalation leads to a larger change in present values with higher p_{11}. Formally, the steady state value of V_1 increases with an increase in p_{11} and the difference between $V_1 - V_2$ also increases. This provides a higher incentive for de-escalation, as countries that enter peace will remain there longer.

2.2 S-stage Model

Before moving to a quantitative application, we extend our framework to make it suitable for it. We consider a discrete-time infinite-horizon model environment with S stages, i.e. we assume $s \in S = [1, 2, ..., S]$, ordered by how severe the conflict outbreak is, with $s = 1$ being the state with the least severe conflict flag and $s = S$ being the state with the most severe one. Let D_s denote the damage caused by the risk stage s and let $\mathbf{D}$ be the vector of damages across all S stages.[16]

Countries transit between different stages s over time and, in the absence of interventions, the likelihood of moving between stages is governed by a transition

[16] In the empirical application, we will rank stages based on the estimated long-run discounted potential damages incurred.

function in the absence of interventions, $\mathbf{T}_n$, which is defined as follows

$$\mathbf{T}_n = \begin{pmatrix} p_{1,1} & p_{1,2} & \cdots & p_{1,S} \\ p_{2,1} & p_{2,2} & \cdots & p_{2,S} \\ \vdots & \vdots & p_{s,s'} & \vdots \\ p_{m,1} & p_{m,2} & \cdots & p_{S,S} \end{pmatrix}. \tag{5}$$

In the matrix $\mathbf{T}_n$, a generic entry $p_{s,s'}$ denotes the probability of moving to a state s' conditional on being in a state s. By construction, each of the entries cannot be negative, $p_{s,s'} \geq 0$, and each row sums to one, $\sum_{s'} p_{s,s'} = 1, \forall s \in S$.

Implementing an intervention allows the government to modify the transition function so that probability mass is moved away from more risky states s', with a certain degree of effectiveness $\tau \in (0, 1)$. In particular, we assume that intervening stage s lowers transition likelihoods proportionally for higher stages and increases transition likelihoods proportionally for lower stages. The general idea behind this proportionality is that policies do not operate in a vacuum but need to work with the existing transition probabilities. Policies work just as one would steer a boat on a river - it is easier to navigate the boat in the direction of the current (escalation or de-escalation paths) than to force it in an unnatural direction. Formally, we assume that policy interventions change transition likelihoods as follows:

$$p^i_{s,s'} = \begin{cases} p_{s,s'} + \tau p_{s,s'} & \text{if} \quad s' < s \\ p_{s,s} + \tau \left(\sum_{s'>s} p_{s,s'} - \sum_{s'<s} p_{s,s'} \right) & \text{if} \quad s' = s \\ p_{s,s'} - \tau p_{s,s'} & \text{if} \quad s' > s \end{cases} \tag{6}$$

where τ denotes again the effectiveness of interventions. In other words, transition probabilities $p^{\phi}_{s,s'}$ given policy ϕ can now be written as

$$p^{\phi}_{s,s'} = \begin{cases} p_{s,s'} \text{ if } \phi_s = 0 \\ p^i_{s,s'} \text{ if } \phi_s = 1 \end{cases}.$$

The effects of interventions specified in (6) implies that all interventions (except in stages 1 and S) now can have preventative and de-escalation elements because they at the same can avoid shifting to a worse stage (prevention) and away from a violent stage (de-escalation). Note, we can have $p^i_{s,s} \geq p_{s,s}$ or $p^i_{s,s} < p_{s,s}$ depending on how much probability mass is moved towards stage s in the preventative part of a policy and how much is moved to lower stages in the de-escalating part of a policy.

We assume proportionality for several reasons. First, as explained above, the model follows the logic of interventions taking place before outcomes are realized. Interventions in more risky situations with a higher probability mass $\sum_{s'>s} p_{s,s'}$ will be able to shift more probability mass. We apply this logic to de-escalation as well - existing pathways out of peace get amplified by the intervention. Second, our assumption avoids theoretical solutions where the intervention leads to outcomes outside the support of the transition matrix without intervention. Third, our approach disciplines the model through the transition matrix. This way, we hard-wire the decision model to build on the underlying dynamics determined by the existing conflict dynamics uncovered by the forecast model.

While interventions have both some preventive and de-escalating effects, the character of a policy in stage s will vary as it will be driven by how much probability mass is in higher or lower stages. In other words, the character of policies that can be implemented in a stage will be a simple function of the ranking in present values V_s and the transition matrix. Stages with a lot of probability mass above them will be more preventative.

Finally, we assume that implementing a policy entails a cost I_s, which depends on the stage s. Let $\mathbf{I}$ denote a vector of costs.

The problem of the policymaker can be expressed using a recursive formulation. The policymaker compares the discounted expected values of intervening with a degree of effectiveness τ, against the discounted expected value of not intervening today, i.e.

$$V_s = \max\left\{V^i_s, V^n_s\right\}.$$

V^i_s is the value of intervening, equal to

$$V^i_s = -I_s + \sum_{s' \in \mathcal{S}} p^i_{s',s}[-D_{s'} + \beta V_{s'}]$$

where $\beta \in (0, 1)$ is the discount factor, while $p^i_{s',s}$ are probabilities given by equation (6). V_s^n is the value of not intervening, defined as

$$V_s^n = \sum_{s' \in \mathcal{S}} p_{s',s}[-D_{s'} + \beta V_{s'}].$$

Notice that the difference between V_s^i and V_s^n lies in i) the transition matrix, which varies depending on whether an intervention is implemented or not, and ii) the cost of intervention, $-I_s$ which is incurred conditional on intervention and reduces the current period's payoffs. A solution to this problem is an indicator function for optimal government intervention, ϕ_s, taking value 1 if an intervention takes place in stage s, and 0 otherwise and a value function, V_s, attaining its maximum when ϕ_s is implemented. We summarize all intervention decisions in policy vector ϕ. Under optimal policy ϕ^o, the resulting transition matrix $\mathbf{T}_o = \mathbf{T}(\phi^o)$ is obtained by combining transition matrix (5) with the effect of policy on transitions given by (6).

2.3 The Information Rent from Forecasting

We model the presence of a forecast as the capacity to split stages into sub-stages through the use of data. We label the information set with the larger number of stages the *full information* set, and the information set with fewer stages the *partial information* set. We can simulate optimal policies under different information sets and calculate differences in welfare. We call this the *information rent* of forecasting.

To illustrate this point, consider the stylized two-stage transition matrix without interventions $\mathbf{T}_n$. With a forecast of conflict outbreaks the policymaker would gain a richer information set, e.g., observations in peace would be distinguished between *high outbreak risk* and *low outbreak risk*.

In the quantitative application, we will estimate a full information model with S stages. In the partial information setting, we will instead assume that the policymaker cannot distinguish some stages. The transition matrix under partial information will only have $\tilde{S} < S$ stages and to notate the partial information treatment we will call this transition matrix $\tilde{\mathbf{T}}_n$. Following the same logic as in Equation (6), we can compute the optimal policy vector ($\tilde{\phi}^o$) resulting from decisions under partial information. This means we can derive two optimal policy vectors - one optimal policy vector under full information (ϕ^o) and one constrained optimal policy vector under partial information ($\tilde{\phi}^o$). Across the two simulations, we hold all other exogenous factors constant (damage vector, intervention costs, discount factor). What differs are only the (perceived) transition matrices.

As a consequence of not being able to distinguish between some stages, the policymaker will either always intervene in the merged stages, thereby engaging in untargeted policies, or never intervene, thereby allowing for some costly false negatives to occur. Moreover, given that prevention and de-escalation interact, as shown in Proposition 2, the optimal policies under partial information might differ from those under full information in other stages as well.

To calculate the information gain from forecasting, we compare the total gains from acting optimally under full information (ϕ^o) against the optimal policy under partial information ($\tilde{\phi}^o$), in a world where transitions occur according to the true matrix $\mathbf{T}_n$ and effects of interventions are given by those in Equation (6). In other words, we simulate the effect of acting optimally when a policymaker has access to a forecast (i.e. can differentiate between all stages S) compared to when they do not distinguish between some stages. Let the true present value $V_s^{partial}$ in stage s of decisions ($\tilde{\phi}^o$) made under partial information be summarized by

$$V_s^{\text{partial}} = -D_s - \tilde{\phi}_s I_s + \sum_{t=1}^{\infty} \beta^t \left(T(\tilde{\phi})^t (-D - \tilde{\phi} I) \right)_s . \tag{7}$$

We can now formally define the information rent Π_s given stage s as:

$$\Pi_s = V_s - V_s^{\text{partial}} \tag{8}$$

To derive a useful measure, we will normalize the information rent by GDP and multiply the vector of net present rent $\mathbf{\Pi}$, summarizing the gains for each stage, with the ergodic stage shares under full information.

3 Application: Armed Conflict

We now set up an application of our framework - the optimal dynamic intervention in armed conflict.[17] The application will allow us to expand the state-space,

[17]Work on this approach was directly motivated by a policy project for the U.K. Foreign, Commonwealth & Development Office (FCDO) (Mueller et al. 2022).

analyze optimal policy under different assumptions and highlight the benefits of a better forecasting system through the derivation of benefit-cost ratios and the information rent shown in Equation (8).

In our analysis, we will take the viewpoint of a policymaker that decides on additional resources to be allocated optimally over the conflict cycle *up and above* what already exists. The thought-experiment is one of coordinated policies by local actors and the international community with the aim to prevent and de-escalate armed conflict optimally with the goal to maximize benefits net of costs. We discuss the policies for conflict prevention and de-esclataion in Appendix F.1.

To take our model of risk stages to the data requires defining and calibrating all the elements of the dynamic programming problem. We will first define stages through the forecasting models at `https://conflictforecast.org/` and observable conflict dynamics under *full information*. This will give us a characterization of stages and a transition matrix $\mathbf{T}$ under full information. In light of our theoretical discussion in Section 2.3, we will also propose an alternative *partial information* scenario in which the policymaker does not have the full information that separates stages and, instead, acts in a reduced stage space.

Importantly, we will assume that the current policy process operates under the partial information scenario. We will calibrate the model under this assumption using data on well-researched interventions like power-sharing agreements, peacekeeping interventions and development aid. The assumption here is that policymakers behave as if they are constrained by their information set. The calibrated model can then be used to answer three questions. First, is prevention cost-effective under the parameters values that best describe the current policy regime? Second, are our theoretical findings from the two-stage model carrying through a much more complex stage-space? Third, what is the information rent from introducing policy targeting with forecasts in peaceful stages?

A central motivation for our analysis is that, currently, forecasts are not used in policy. The assumption behind our welfare analysis that this is because forecasts are not available. We make this assumption both part of our calibration exercise and the policy experiment of introducing the forecast to derive the information rent. We defend this assumption in Appendix F.2. However, the lack of systematic use of forecasts in policy targeting also has political economy and organizational reasons. Our findings should therefore be regarded as offering a normative benchmark of potential gains rather than a positive description of feasible steps towards the information rent.

3.1 Risk Stages and Full vs Partial Information

We calibrate our model using the assumption that there are no systematic conflict prevention policies implemented in peaceful countries because policymakers cannot distinguish high and low conflict risk situations. We calibrate policy effectiveness and costs for this partial information world but first need to derive the full information stages.

3.1.1 Conflict Risk Stages

Defining risk stages is part of the decision process. A risk stage is meant to capture the environment a country is in a given month. Identifying different stages is pivotal to characterizing conflict dynamics and simulating interventions. Rather than enforcing arbitrary cutoffs on our risk and intensity forecasts, we use a Hidden Markov Model (HMM), which is a statistical model, to detect latent states and give an 'optimal description' of the stages and their dynamics via various data inputs we provide.[18]

What is the optimal number of stages? Theoretically, it makes sense to think that more stages are always better. However, in practice, we want to derive a parsimonious description of conflict dynamics that is as valid as possible across countries. This means we cannot split forecast and conflict data in too many states as these quickly become unconnected among themselves. In addition, our definition of conflict risk stages also affects the analysis of conflict damages.[19]

We settled on twelve stages under full information after experimentation. This choice is motivated by two criteria, which are 1) allowing sufficient observations in each stage and 2) differentiating between and within situations of stable peace, elevated risk, post-conflict, and during violence. Panel (A) in Table 1 shows the summary statistics of the resulting stages and Panel (B) shows the summary statis-

[18]This was suggested to us by the former chief data scientist of the FCDO Tom Wilkinson whose input was extremely useful for the development of the latent space model.

[19]Damages are estimated using country-fixed effect regressions and these become meaningless when the number of states is too high. See Section 3.2 for more details.

tics for the partial equilibrium model where we merged stages 1 to 5 into a new stage 1.[20] The columns in Table 1 show the variables that underlie the clustering the HMM uses to identify risk stages. Further details of the derivation of the full information world are shown in Appendix B.

One of the key insights from Mueller and Rauh (2022) is that armed conflict risk becomes very hard to detect when the last conflict is sufficiently long ago. This is true for stages 1 to 5 in the full information stage model which (on average) have no recent violence. This is where the risk forecasts of armed conflict outbreaks and conflict intensity play a crucial role. From stage 1 to stage 9 the outbreak model increases up from less than 10% to over 80%. This is because the outbreak model treats ongoing conflicts as extremely high risk for outbreaks but is not able to further distinguish different situations of risk. The predicted intensity of armed conflict reinforces the ordering while adding information that helps distinguish stages 1 to 7 from the later stages. Note, for example, that stage 9 features a very low observed number of fatalities but a very high-intensity forecast in the full information model.

[20]We order the stages according to our present value of costs without intervention, consistent with the model in Section 2. In Appendix Figure B2 we show for the full information model that we hardly sacrifice any predictive performance for the onset of armed conflict by reducing the continuous predicted conflict probability to twelve risk stages.

Table 1: Summary statistics for risk stages

(A) Full information

	(1)	(2)	(3)	(4)	(5)	(6)
	Predictions			Past realizations		
	Likelihood of armed conflict		Battle deaths	Battle deaths	Months since	Share of
Stage	Full model	Text model	per 1mn inhabitants	per 1mn inhabitants	last armed conflict	observations (%)
1	0.01	0.04	0.05	0	321	34.4
2	0.03	0.11	0.03	0	239	15.8
3	0.09	0.18	0.08	0	114	16.3
4	0.16	0.21	0.18	0	143	6.0
5	0.16	0.25	0.46	0	144	2.7
6	0.40	0.29	1.13	0	9	6.1
7	0.47	0.34	2.00	0	17	2.7
8	0.68	0.38	10.00	4	0	1.9
9	0.81	0.49	7.83	0	3	4.9
10	0.90	0.52	24.69	2	0	4.4
11	0.91	0.61	101.42	12	0	3.9
12	0.90	0.62	481.99	131	0	1.0

(B) Partial information

	(1)	(2)	(3)	(4)	(5)	(6)
	Predictions			Past realizations		
	Likelihood of armed conflict		Battle deaths	Battle deaths	Months since	Share of
Stage	Full model	Text model	per 1mn inhabitants	per 1mn inhabitants	last armed conflict	observations (%)
1	0.05	0.10	0.08	0	239	75.2
6	0.40	0.29	1.13	0	9	6.1
7	0.47	0.34	2.00	0	17	2.7
8	0.68	0.38	10.00	4	0	1.9
9	0.81	0.49	7.83	0	3	4.9
10	0.90	0.52	24.69	2	0	4.4
11	0.91	0.61	101.42	12	0	3.9
12	0.90	0.62	481.99	131	0	1.0

Notes: Columns (1) to (5) display the means of the variables indicated in the headings. The stages are the result of an HMM based on these inputs. Predictions, available at `conflictforecast.org`, are derived using a random forest with past violence and topics from six million newspapers summarized using the Latent Dirichlet Allocation as predictors. Predicted and realized battle deaths are transformed as log+1 and months since the last armed conflict using a polynomial of degree specified in Appendix B. The distributions of the inputs are presented in Appendix Figure B1. Column (6) indicates the share of observations in each stage. Entries in the table represent the means at given stages. Stage 1 in partial information is a weighted average of full information stages 1 to 5 where weights are given by the distribution of country and months across stages.

Table 1 also reports the average number of months since the last armed conflict for the different stages. A low number means that the country suffered from recent conflict. High numbers mean a long period of peace indicating stability. This column helps us summarize the stages in three broad groups in full information. We call stages 1-5 *pre-conflict*. Stages 6, 7, and 9 are very high risk and typically post-conflict which is why we call the *post-conflict risk*. Stages 8 and 10 represent relatively low-level violence, while stages 11 and 12 exhibit higher levels of observed fatalities. Stage 12 is capturing the top 1% of risk and violence - the worst possible situation from a humanitarian perspective.

The idea of our information treatment is that without a systematic forecast, it would be difficult to distinguish stages 1 to 5 with the precision of our machine learning model. If such a forecast is not available, the policymaker is therefore not systematically aware of transitions between stages 1 to 5. To simulate this information set we merge stages 1 to 5 under what we call *partial information*. This generates a new super-stage 1 which absorbs 75% of all our observations. In terms of average likelihood, intensity, and months since the last conflict armed conflict, this stage is close to the old stages 1, 2, and 3 because of the higher share of observations these stages have in the full information model. To illustrate the difference between the full and partial information treatment we first turn to the respective transition matrix.

3.1.2 Transition Matrices

The transition matrices and full and partial information, $\mathbf{T}_n$ and $\tilde{\mathbf{T}}_n$, are at the heart of our model. They capture the true dynamics and perceived dynamics between stages in a simple way that allows us to simulate possible futures for each stage under different assumptions.

Figure 1(a) shows a visualization of the transition matrix with no interventions under full information. The y-axis denotes the stage number in month t and the columns denote stages in month $t+1$. The shades of red mark the size of probabilities with the darkest colors indicating probabilities over 90% and light red indicating probabilities below 1%. The most obvious feature is that all stages, except for stage 8, are very stable. The diagonal of the matrix can be seen in dark red.

Under partial information, shown in Figure 1(b), there is no forecast in countries without recent violence so that the policymaker cannot distinguish between

Figure 1: Transition matrices

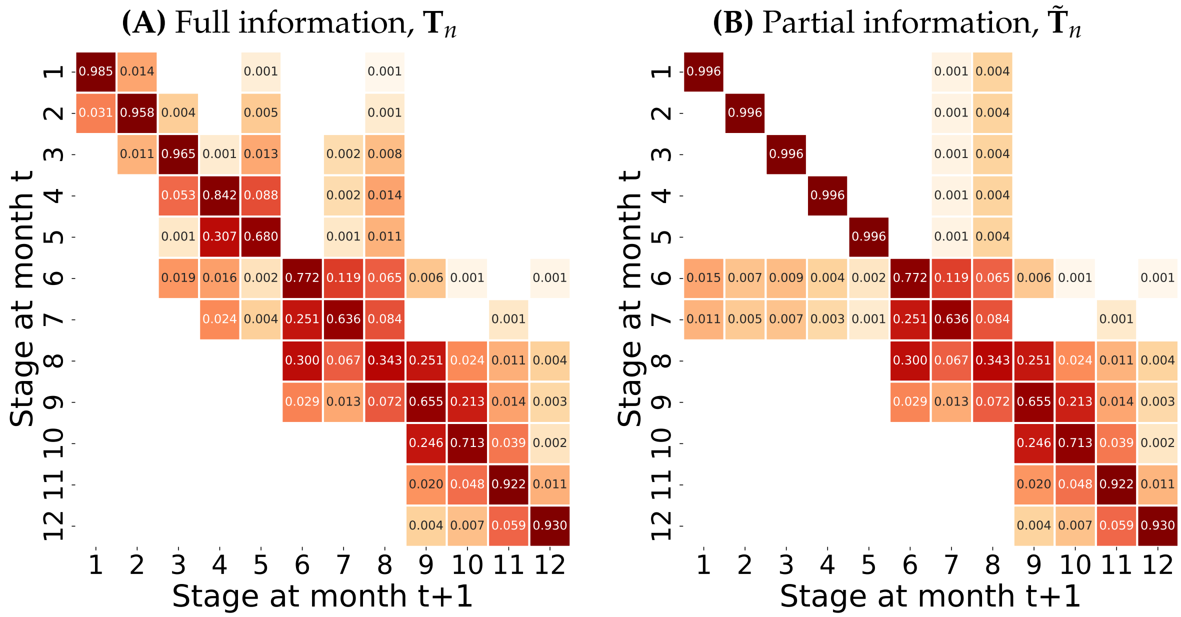

NOTES: In both panels the y-axis depicts the stage in month t and the x-axis the stage in month t + 1. Entries represent the likelihood of moving from a stage in month t to another stage in month t+1. The darker the shade of red, the higher the probability. Panel A shows the transition matrix under full information in the absence of interventions. Panel B shows the perceived transition matrix under partial information in the absence of interventions.

situations of stable and relatively high-risk peace. The figure shows transition matrix $\tilde{\mathbf{T}}_n$ under the assumption of no interventions. The new super stage 1 is visible as a darkish red diagonal with an extremely high persistence and very low risk of escalation into stages 7 and 8.[21] This is a crucial contrast to the full information set in Figure 1(a) where the likelihood of a transition to stages 7 and 8 goes from less than 0.1% in stages 1 and 2 to over 1% in stages 3 to 5.

A common feature in both transition matrices is the asymmetry and the strong variation in the connectedness of the stages. If a country suffers an outbreak by transitioning from stage 6, for example, to stage 10 or 12 there is no way to transition back directly. This is the essence of the conflict trap. Transitioning to higher stages changes the distribution for further transitions dramatically towards worse

[21]Notice, that in stages 6 and 7 there exists a non-zero probability of moving back to any of stages 1 to 5. In the partial information transition matrix, we redistribute this probability according to the weight of each stage in the ergodic distribution of the full information transition matrix without intervention displayed in Appendix Table D.1.

outcomes. In stages 6 to 8 there are very broad possibilities of what could happen, i.e. these are situations of extreme danger of escalation but also provide chances for de-escalation. If an information rent arises from full information this is because the policymaker under full information notices a change in the underlying dynamics where stages 3 to 5 under full information are much more connected upward. However, the presence of an information dividend is not a given. The risks, even in full information, are still relatively minor when compared to the extreme connectedness of stages 6 to 10 with the conflict stages 11 and 12.

3.2 Damages

We focus on four components of conflict damages: the cost of lives through conflict fatalities, the damage to the economy through the reduction of growth in conflict, the cost of humanitarian aid through the stock of refugees that have to leave the country, and official development assistance (ODA) spending on emergency response and peace/security which are directly related to armed conflict. It is important to keep in mind that this is an extensive, yet not exhaustive list of costs and that the overall costs are likely larger. Specifically, we can neither quantify the distress of refugees themselves nor the full costs and benefits to the host communities. Some of these costs, however, are captured by the GDP loss as displacement and economic decline have strong overlaps. This is why we focus on the cost to the UN system from displacement.

We discuss the data sources and the derivation of the damage vectors in Appendix C.1, where we also explain how we translate and summarize damages into monetary terms. Note that our starting point is a panel of 176 countries for which conflict risk predictions are generated and a subsequent risk stage can be assigned. However, our sample is reduced to 168 countries due to missing GDP data.[22] Our results confirm findings in the literature according to which ongoing conflict is extremely damaging to the economy. Stages 11 and 12, in particular, are associated with significant reductions in the growth rate. On average, a country in these stages can expect to experience a yearly loss of GDP per capita equivalent to 5.3% and 9.6% respectively. Importantly, the post-conflict stages (8, 9, and 10) are

[22] Countries which are missing GDP data in our sample period are Bermuda, Hong Kong, Puerto Rico, North Korea, Palestine, Taiwan, Venezuela and Kosovo).

not associated with economic booms so economic damages persist and compound. When simulating countries' trajectories, we draw from damage distributions given by the estimated coefficients and standard errors shown in Appendix C.1.

We split the analysis here by GDP per capita levels. The average static damage for all countries below the median GDP per capita for partial information can be seen in panel a of Figure 2. The equivalent figure for all countries above the median GDP per capita can be seen in panel b. We show the total costs broken down into four categories. First, note the differentiation in the damage composition across the two income groups, particularly in stage 12. For low-income countries, a significant proportion of the damages are driven by fatalities. For high-income countries, damages are driven almost entirely by GDP losses. Irrespective of income group and stage, refugee costs play a minor role. In summary, according to the static view, damages mostly occur in stages 10 to 12. The total monthly cost of conflict spent in stage 12 is close to 16 billion US$ for low-income countries compared to 80 billion US$ for high-income countries.

Figure 2: Perceived static damages by stage, partial information

(A) Below median GDP per capita (B) Above median GDP per capita

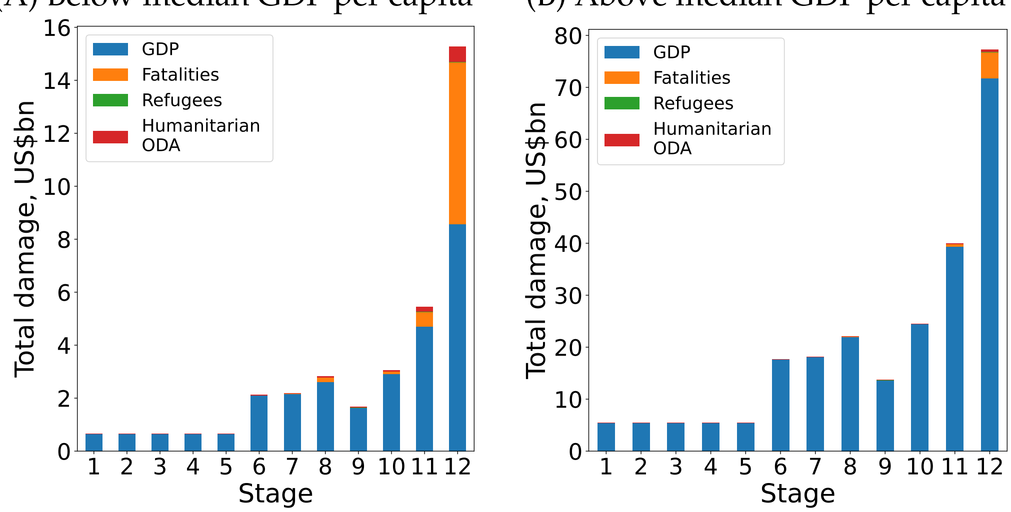

NOTES: Both panels show the average damages D_s in a given period for panel (A) countries with below median GDP per capita and (B) above median GDP per capita. Note that in the partial information case shown here, the perceived future damages in stages 1-5 are always equal since the policymaker cannot distinguish these stages. Therefore, we modify the static damage vector for stages 1-5 to be a weighted average of the static damages in the full information case, according to the in-group ergodic distribution of the full information transition matrix in the absence of interventions.

3.3 Policy Effectiveness and Costs

Effectiveness and costs are two sides of the same coin so that we simply derive a single number of effectiveness which we then hold constant across the entire policy vector. We then discuss different cost functions and conduct robustness checks with these. To have a starting point for τ in the model, we take the estimated treatment effects of power-sharing agreements from Mueller and Rauh (2024) and simulate the implementation of a policy in the stages in which power sharing typically takes place.

Power sharing agreements stand for a larger engagement which is coupled with foreign aid promises, the deployment of peacekeeping forces, and institutional changes. These sorts of interventions are focused on countries with ongoing or recent violence and, accordingly, power sharing agreements are concentrated in stages 9 to 12. In our simulations, we target the average treatment effect of a reduction of 45% in the violence intensity within 18 months (Mueller and Rauh 2024). Our calibration method is discussed in Appendix E. As a result of this, we get an estimate for τ of 16%.

To carry this estimate of policy effectiveness into our model we take the estimates of τ and the transition matrix in Figure 1(b), i.e. under partial information, and apply the shift formulas from Section 2.2. We visualize the resulting shifts in probability mass in Figure 3 in which red squares indicate reductions in probability and green ones indicate increases.

Policies in the conflict trap, i.e. in stages 8 to 12, move most of the probability mass. Their high probabilities above and below the diagonal lead to strong effectiveness. One interpretation of stages with strong transition tendencies toward de-escalation or prevention is that they have an identifiable 'momentum,' and policy interventions can exploit this momentum to either counter and address it in prevention or to re-enforce and help it in de-escalation. This also demonstrates the realism of the assumption of proportional policy effectiveness. Interventions in high stages will be attractive because of the very concrete dangers of further escalations and the urgency for de-escalation they are associated with.

Little is known about the cost of interventions. In one of the few studies Chalmers 2007 that contrast intervention cost estimates with their benefit the numbers vary widely. Chalmers proposes, for example, an intervention package of financial assistance for sustainable governance for Sudan with an estimated cost of 1.3 billion

USD spent over 15 years, i.e. 7.2 million USD per month. This contrasts with the estimated cost of the US intervention in Afghanistan which stands at 2.26 trillion USD over 20 years (SIGAR 2021), i.e. 9.4 billion USD per month.

To give a realistic estimate of intervention costs we model two main elements of the cost function. The first element is the *cost per GDP* which we treat as the baseline expenditure necessary to build a relation with local actors and leave an impact on the local economy. The crucial assumption here is that we assume that the costs of a policy intervention scale with the GDP in the target country. Interventions in all stages are more expensive the higher is its GDP per capita. To get an idea of what is realistic as a baseline cost we consider the total of "economic" and "government and civil society" official development assistance (ODA) spending. As these categories target economic and institutional dynamics they could also be used to try and affect conflict dynamics (Rohner 2024). Since ODA data is yearly, we divide by twelve to get the monthly ODA expenditure on these two categories. We then divide by total GDP to get a per GDP spend. We then take the average ODA per GDP spent across all observations where ODA expenditure on these two categories exceeds 0. As a result, we assume a fixed monthly intervention cost in each stage of approximately 0.1% of GDP. In addition to this baseline cost, we derive a *cost per fatality*, capturing the fact that interventions becomes more expensive with ongoing violence. We find that peacekeeping budgets increase by 14.956mn USD for a one unit change in the log of fatalities per 1mn inhabitants. This estimate is then applied to the mean log of fatalities per 1mn inhabitants by stage. For example, in stage 12, the average log of fatalities per 1mn inhabitant is 4.38, meaning that an intervention costs an additional 65mn USD per month in this stage. Figure 3 shows the resulting cost levels across stages and for different levels of GDP. We observe that GDP is one of the main drivers of costs.

Figure 3: Intervention effect and costs

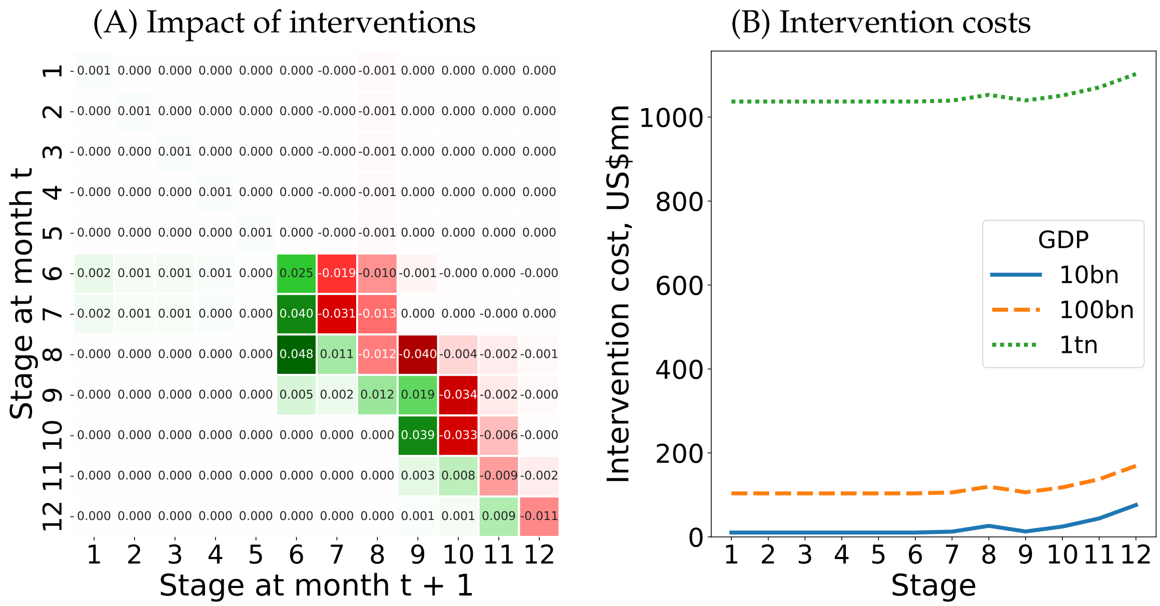

NOTES: Panel A shows the impact of interventions with effectiveness $\tau = 16\%$ on the partial information transition matrix. The y-axis depicts the stage in month t and the x-axis the stage in month $t+1$. Red indicates reductions in probability masses, green an increase. Panel B shows the monthly intervention costs for each of the 12 stages. The blue solid line shows the monthly intervention costs for a country with a total GDP of US$10bn across stages. The orange dashed line shows the intervention costs for a country with a total GDP of US$100bn and the green dotted line for a country with a total GDP of US$1 trillion.

This concludes the estimation and calibration of damages, transition likelihoods, policy effectiveness, and costs using the partial information model. We now take this estimated model and analyze optimal policies under full information.

4 Results

To derive results, we move from the transition matrix under partial information to the transition matrix under full information, keeping policy effectiveness τ and intervention costs the same. Importantly, this changes the usefulness of policy tools in stages 1 to 5 as the degree to which they can be timed according to the underlying risk.

4.1 From Static to Dynamic Damages

The damages from armed conflict are one of the key factors leading to poverty and low average growth in affected countries. Grouping observations into stages allows us to estimate contemporaneous damages D_s associated with a stage s. The dynamic model allows us to look ahead through simple interactions with the transition matrix $\mathbf{T}_n$. The present value of damages in each stage is given a stream of future discounted damages in which intervention never takes place. Formally, discounted future damages in stage s with intervention never taking place anywhere D_s^{PV} are given by:

$$D_s^{PV} = D_s + \sum_{t=1}^{\infty} \beta^t \left(\mathbf{T}_n^t D\right)_s \tag{9}$$

where $(\mathbf{T}_n^t D)_s$ represents the expected damages at time step t, starting from state s, after applying the transition matrix, $\mathbf{T}_n$, t times. Note, that for low discount values all dynamic values will be very close to each other as the starting position matters less and less. In what follows we will contrast the dynamic view derived from our model to the static view which follows the contemporaneous damages D_s.

Figure 4 contrasts static and dynamic damages. In panel (A) we report static damages for the full information model. In panel (B) we report dynamic damages as defined in equation 9. There is a striking contrast in worldview that can be derived from this. If a policymaker is driven by static damages she will put much more attention to countries in stages 11 and 12. When looking ahead, through the perspective of connected stages and the resulting dynamic damages, relative attention shifts away from the conflict stages and instead becomes a continuum.[23] We show in section 4.3.1 that the incentives for prevention naturally fall with the discount rate.

4.2 Optimal Policy under Full Information

We are now ready to report our first main result - the discussion of optimal policies under full information using the calibrated model from partial information. The main question here is whether interventions in lower stages are ever cost-effective in full information for parameters that have been calibrated on policies that are

[23] Notice the relatively low dynamic damages in stages 1 and 2 which are driven by their close association, persistence and disconnect with the other stages.

Figure 4: Static vs dynamic damages by stage, full information

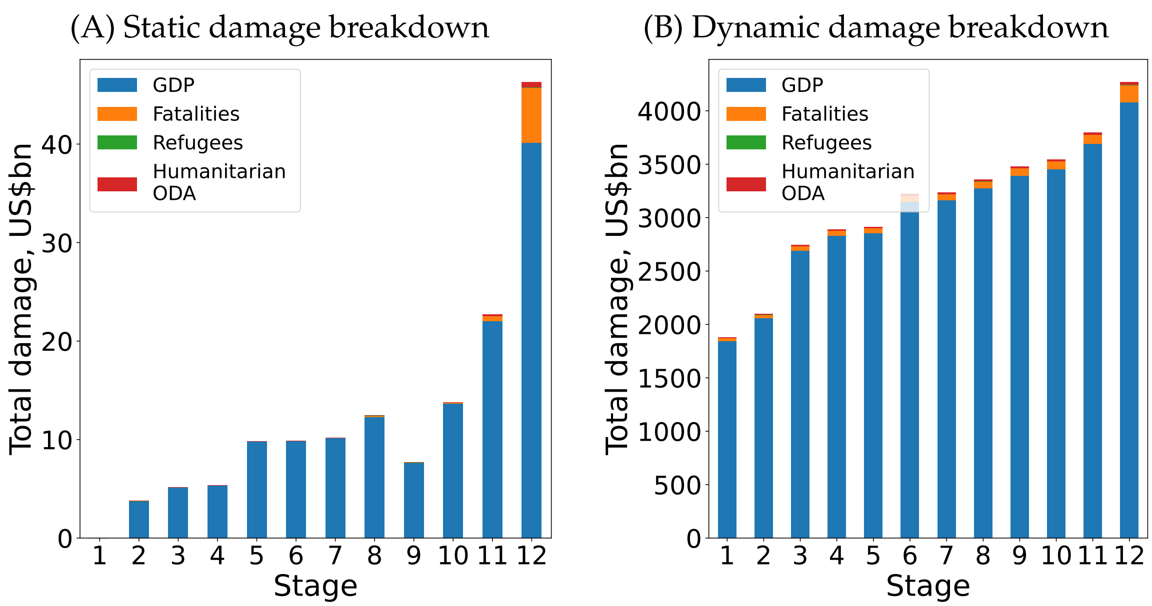

NOTES: The left panel shows the damages in a given period D_s, while the right shows D_s^{PV} defined by Equation (9), which is the present value including all future damages discounted using an annual rate of 4% and assuming a 1% annual GDP growth. Damage estimates in both panels represent the average across all countries in our sample.

common in higher stages.

The results in this section are reported as benefit-cost ratios (BCR), which can be interpreted as the net present value of the long-run return per \$1 spent today.[24] Importantly, the BCRs capture the idea of a one-period deviation. We calculate the net gains as the value function difference between intervening in stage s and not intervening in stage s, while always maintaining optimal behavior in the future. In other words, we take into account what a rational policymaker would do in a future in which the situation might escalate. This ensures that the policy benefits are optimal, even under the assumption that the failure to prevent can be partially offset by later, optimal de-escalations. To calculate gross gains we then add the intervention costs of intervening in stage s. To get to the BCR we then divide this

[24]The benefit-cost ratio is computed as gross gains divided by intervention costs. Future gains/losses are discounted using a rate of 4% per year.

sum by the intervention costs. So formally, BCR_s in stage s is given by

$$BCR_s = \frac{V_s^i + I_s - V_s^n}{I_s}.$$

A BCR less than 1 implies that the intervention is not cost-effective. Remember that the calculation of the BCR takes into account the low predicted risk in countries with no current violence and that policies are not guaranteed to be effective. For low stages for which the predicted risk is very low, large dynamic benefits of prevention need to compensate for low escalation risks to bring the BCR over 1.

To capture uncertainty in our BCR estimates, we take into account the uncertainty from the estimation of the conflict stage model and the OLS regressions used to estimate damages. When estimating the Hidden Markov Model the Viterbi algorithm computes the most probable path and assigns a stage to each country/month observation based on this most likely sequence. However, the Hidden Markov Model also outputs a posterior probability distribution over the stages. In other words, there exists some uncertainty over what stage the country is assigned to. We can therefore sample a stage, for each country/month observation, from the posterior distribution 10,000 times. This gives 10,000 draws of the sequence of stages for each country, from which we can compute 10,000 different transition matrices.[25] We pair each of these transition matrices with a damage vector drawn using the coefficients of the OLS regressions and their respective standard errors.[26] Together, these simulations yield standard deviations around the BCR, driven by the uncertainty of stage assignment and damages incurred.

We have in total 10,000 BCR estimates across all stages for the 168 countries in our sample. First, we compute the average BCR by stage across all countries for each draw. Figure 5 displays the distribution of these mean benefit-cost ratios by stage. To avoid distributional assumptions, we show the 2.5th percentile, mean, and 97.5th percentile. Under full information, it turns out to be beneficial to intervene in all stages except for stage 1. The point estimates of the BCRs turn out to be hump-shaped with respect to stages, with the highest BCR peaking at 15 in

[25] In Appendix D, Figure D1 we show the standard deviations around each of the estimated transition probabilities.

[26] We sample 10,000 times from the parameter estimates of the OLS regressions, assuming a normal distribution where the mean is equal to the point estimate and the standard deviation is equal to the standard error.

stage 8. The point estimate for intervention in stage 1 is close to 1, while in stage 12 the return is still higher than \$5 per dollar spent. The high relative return to interventions in stage 8 is driven by the high uncertainty about the pathways of this stage. Note that uncertainty regarding the categorization of this stage is also high so that a risk averse policymaker would not regard it as particularly attractive to intervene in stage 8.

Figure 5: Benefit-cost ratios

(A) Benefit-cost ratios of interventions (B) Benefit-cost ratios of policy bundles

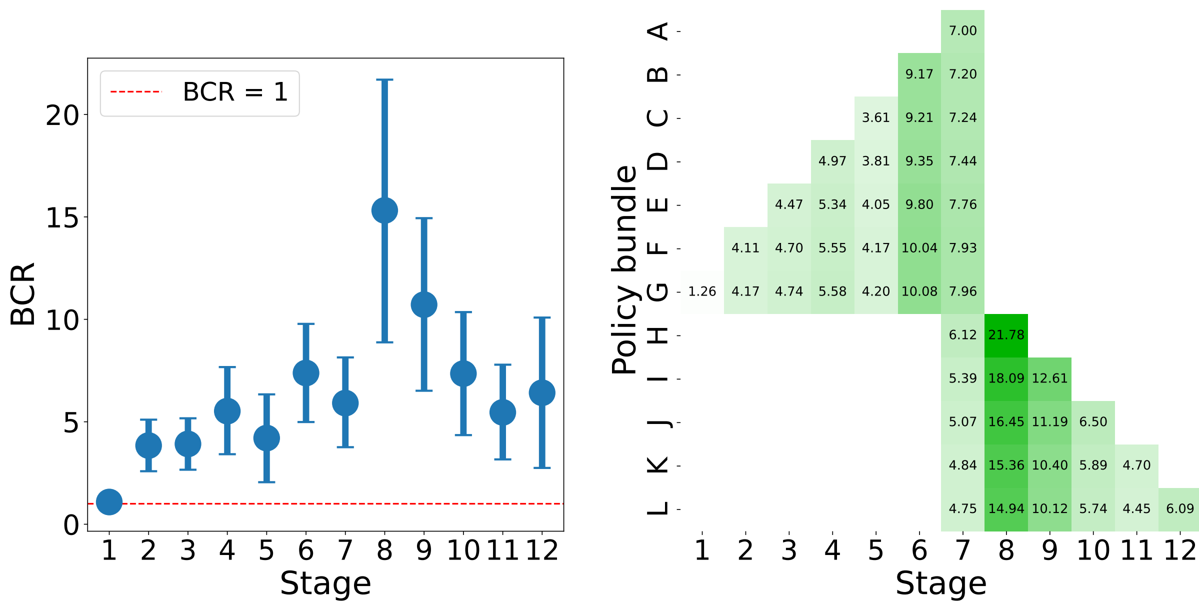

NOTES: Panel A shows reports benefit cost ratios (BCRs) from 10,000 simulations across all countries by stage calculated from equation (4.2). For each simulation we draw a transition matrix and damage vector, calculate the BCR for each country, and average across countries. Point estimates are the mean BCR score. Bars represent the 2.5th and 97.5th percentile of the 10,000 means. Panel B shows BCRs of different policy bundles (indicated by rows). Interventions in white cells are switched off, i.e. the policymakers is not allowed to intervene. The calculation of BCRs does not account for the uncertainty of stage assignment and damages incurred as previously described. Instead, the point estimate of OLS regression coefficients is used to generate D_s, and the full information transition matrix in Panel A of Figure 1 is utilized. Darker shades of green indicate higher BCRs.

Crucially, the point estimates for the BCRs of stages 4 to 7 are not too dissimilar to the BCRs of stages 11 and 12. This is partially a result of the assumptions regarding intervention cost increase with fatalities. But it is to a large degree driven by the dynamic benefits of preventing a fall into the conflict trap. Here is where stage

8 is particularly interesting. When we run community detection algorithms on our full information transition matrix they detect a connected community going from stage 9 to 12 but excluding 8 when we force few communities.[27] This means stage 8 is more connected to stages below it than stages above it. It is just outside the conflict trap but right before it. Figure 5 shows that intervening in this situation maximizes expected prevention benefits.

4.2.1 Complementarity versus Substitutability

We now turn towards analyzing the spill-overs identified in our theoretical results. To investigate this we switch off the possibility to intervene in some stages and re-optimize to find the new BCRs for higher and lower stages.

In Figure 5(B) we simulate the returns to different policy bundles around stage 7 to see whether the BCR for interventions in stage 7 is affected and by how much. We start with row A where interventions are only permitted in stage 7. Row by row we allow more interventions in lower stages to join till policy bundle G where interventions are allowed to take place in stages 1 to 7. We then check how incentives to intervene change in stage 7 as we add substitutes in higher stages.

In policy bundle A, the return to intervening stage 7 is $7 per $1 invested. This increases with every additional lower stage in which interventions are permitted - as predicted by our theory. In other words, prevention in stable peace is complementary to interventions in intermediate risk stage 7. This is because lower stages become even better when intervening in them so incentives to shift to these stages from stage 7 increase. Notice the significant increase in BCRs from 7 to 7.96 when we allow for more preventative policies.

In contrast, allowing for de-escalation in higher stages reduces the BCR in stage 7. This crowding out is substantial, reducing the BCR to below 5 in row L when allowing for interventions in all higher stages. Incentives to intervene in stage 7 fall by close to 33% with interventions in conflict. As outlined in the theory model, the rationale is that if one de-escalates during high risk and raging conflict, violent situations become less of a burden, reducing incentives to prevent them in the first place.

[27]Calculations available from the authors. The algorithms we used were the Leiden Algorithm and the hierarchical Girvan-Newman. The Girvan-Newman algorithm first separates stage 1 from all other stages, then stage 2, and then makes 4 communities 1, 2, 3-8, 9-12

4.3 Information Rent

We now turn to our main policy experiment: the withdrawal of the forecast. Imagine that there is no forecast in countries without recent violence so that the policymaker cannot distinguish between situations of stable and relatively high-risk peace, i.e. between stages 1 to 5. In Equation (8) we defined the information rent from providing this policymaker with a forecasting module. According to our definition, an information dividend is generated if having the full information leads to different decisions than under partial information.

In panel A of Figure 6 we show the distribution of gains from 10,000 simulation draws to assess the outcomes of implementing the optimal policy under partial information in a hypothetical full information world and comparing those outcomes to the optimal full information outcome. The x-axis shows information gain (relative to GDP) for the countries included in the sample. The average gain is 62.2% of GDP (standard deviation of 15.8). Less than 1% of simulations yield a gain of zero.

Given the BCRs in Figure 5 it should be clear that the information rent is coming mostly from the low stages which are generating interventions under full information but little interventions under partial information. Panel B of Figure 6 analyzes the differences in policy choices across stages under full and partial information. Under full information, the policymaker optimally intervenes in over 60% of cases where it would not optimally intervene under partial information. This figure rises to almost 100% in stages 2 to 5. [28]

This dramatic increase in intervention incentives with the information set is surprising. Why are interventions in stage 1 suddenly optimal despite this being the lowest stage? The reason is that stage 1 is most closely connected to stages 2 to 5. This creates the incentive to implement preventative policies in stage 1 to prevent escalation into these stages under full information. Under partial information, this incentive is completely gone as stages 1 to 5 cannot be distinguished. This finding suggests that the finer information granularity generates incentives for policy interventions that are completely absent otherwise. Put differently, having a forecast helps the policymaker formulate policies because it provides new policy targets for conflict prevention in situations that are very far away from con-

[28] Appendix Figure D3 shows that in the partial information case, a policymaker would intervene just 1% of the time in stages 1-6 across all countries and draws. In both partial and full information, it is almost always optimal to intervene in any of stages 6-12.

Figure 6: Distribution of information rent

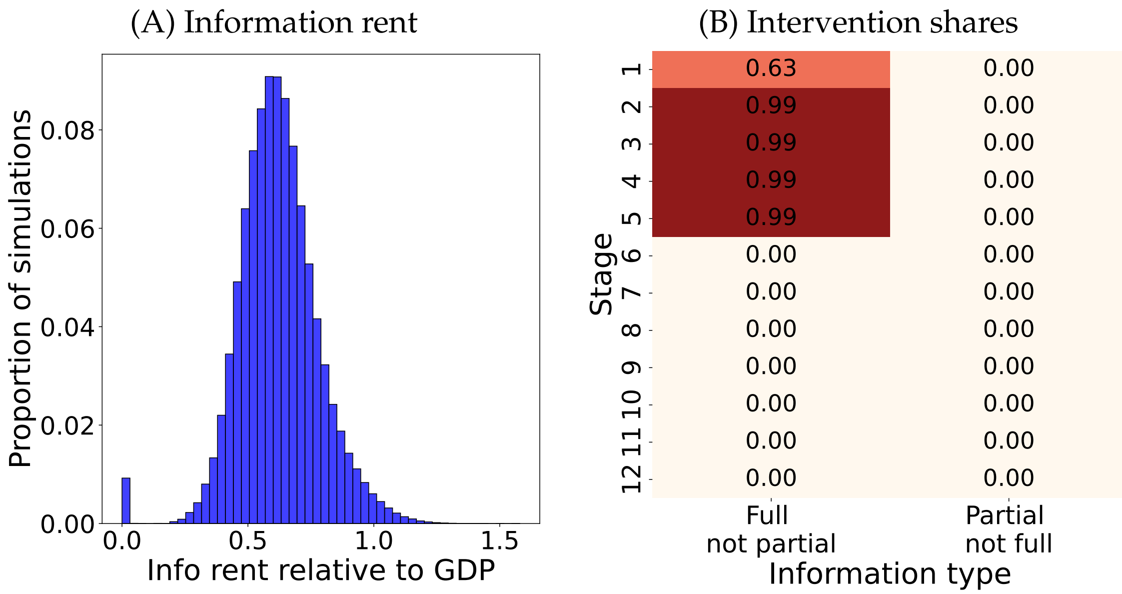

NOTES: Calculations for panel (A) follow Equation (8) and include GDP damages. The resulting information dividend by stage is reduced to a single value using the ergodic distribution weights across stages in full information without intervention. For panel (B), the left column of each heatmap shows the share of times a policymaker optimally intervenes in full information, but not partial information, across all countries and draws. The right column shows the share of times one optimally intervenes in partial information, but not in full information.

flict.

We have held the policy tools available to the policymaker constant. In the terminology of Kleinberg et al. (2015), our policy experiment is only changing the prediction quality of the policymaker, not the available treatments. This change in the information environment leads to a dramatic shift in implemented policies and large economic gains. One way to understand this is by contrasting the ergodic distributions with full and partial information. The main change is in the share of countries that are in stage 1 in the long run which increases from 32% under partial information policies to 46% under full information. Some of this change comes from other peaceful stages but the total weight on stages 1 to 7 goes up by close to 4 percentage points and the share of countries in conflict falls by a percentage point. This is despite the fact that we use a pseudo-out-of-sample forecast with realistically low performance in stages 1 to 5 and policy treatments that fail most of the time, i.e., in more than four out of five cases. The information environment has profound welfare implications.

4.3.1 Robustness and Additional Results

We run two sets of robustness checks and present these results in Table 2. First, we analyze countries above and below median GDP per capita. This shows that, in relative terms, a substantial information dividend arises regardless of the income group we focus on. In absolute terms, the return is greater for richer countries given that there is more at stake. Second, we study the sensitivity of the result to a change in the discount rate. An increase in the discount rate from 4% to 10% yields an information rent in excess of 12% of GDP on average, which is a lot lower than the benchmark rent of 62%. However, the relationship here is to a large degree mechanic where stronger discounting leads to lower present values. It is still optimal to intervene in stages 2 to 5 under full information - even with a 10% discount rate.

Table 2: Robustness checks

Country sample	Discount	Information rent (% of GDP)	
	rate (%)	mean	standard deviation
Benchmark	4	62.2	15.8
Above median GDP per capita	4	57.5	11.8
Below median GDP per capita	4	66.8	17.8
Benchmark	10	12.5	3.4

Notes: The mean and standard deviations and the information rents are computed using 10,000 simulations.

However there is the possibility that the policymaker uses extreme levels of discounting for political economy reasons.[29] A drastic increase of the discount rate to 50% would lower BCRs across the board, and even decrease them below 1 for lower risk stages given that so little value is placed on (potential) future events. The policymaker acts more according to the static damage model shown in Figure 4(A) above. Only some preventative incentives, especially in stages 4 and 5, survive even such extreme levels of discounting under full information.

[29]For recent examples see Aguiar and Amador (2011) or Besley and Persson (2011a).

Optimal policy also changes dramatically when we disregard the development incentives for donors altogether, i.e. when we don't treat damages to GDP growth as costs from conflict. The withdrawal of the main driver of damages reduces the returns to interventions dramatically. As expected, in panel (B) of Appendix Figure D2 we see a dramatic decrease in intervention incentives with BCRs even dropping below 1 for stages 1 to 5 under full information. Conflict prevention in our model is, to a large extent, motivated by developmental gains, i.e., by the economic welfare of conflict-affected countries.

5 Conclusion

This article presents a dynamic decision model to analyze the trade-offs involved in deciding whether to engage in early action (prevention) and/or late action (de-escalation) in the face of recurring risks such as armed conflicts. The Markov model has brought to light that prevention tends to encourage subsequent de-escalation actions, while de-escalation can, in certain contexts, deter preventive measures. This interdependence highlights the complexity inherent in crafting policies for conflict management and resolution.

A critical revelation of our model is the transformative role of forecasts in policymaking. Their integration fundamentally reshapes the approach to risk management. By applying our framework to a global dataset of armed conflicts spanning from 2010 to 2022, we have highlighted the existence of an information rent in armed conflict policies, currently overlooked or underutilized by governments and international organizations. Our analysis of status quo policies suggests that this rent could be realized through systematic use of forecast and the potential long-term benefits of such a strategic policy are significant and far-reaching.

But, the implications of our findings extend well beyond the sphere of armed conflicts. Similar frameworks might help shed light on other policy challenges as well, ranging from public health crises, inflation, financial crises, and climate change policy decisions. Future research could explore these avenues, examining the role of forecasts in dynamic decision-making across diverse domains. Moreover, empirical research focusing on the practical application of these models in real-world policy settings could yield valuable insights into the challenges and efficacy of forecast-based prevention and de-escalation strategies.

But there is also a long tradition within the political economy literature that underpins preventative policies in several areas. For example, institutional frameworks like the European stability pact can be traced back to academic contributions to fiscal debt.[30] Similarly, one of the arguments for central bank independence has been that short-cited politics or commitment problems would get in the way of a more preventative approach to inflation.[31] Accordingly, central bankers counter inflation not only when it arises.

Two shortcomings are that the analysis ignores uncertainty around the forecast and policies that are already being implemented. The forecast system does not explicitly take into account the implemented policies but we know from the literature on conflict that effective preventative policies have been implemented in many cases in the past. This means that our forecast model relies on data that includes attempts to counter risk. Unraveling the mix of policies and other risk factors will be crucial to understand risk better and at the same time, get a better understanding of the effectiveness of preventative policies. This is also crucial for implementing a policy system based on forecasts as systematic, effective policies will make forecasts obsolete. A systematic collection of data on policies is crucial here but, at this point, we also lack the econometric and machine learning tools to integrate forecasting and causal inference.

References

Acemoglu, Daron and James A Robinson, "A theory of political transitions," *American Economic Review*, 2001, *91* (4), 938–963.

_ **and James A. Robinson**, *Why Nations Fail: The Origins of Power, Prosperity, and Poverty*, Crown Business, 2012.

Acharya, Sushant, Julien Bengui, Keshav Dogra, and Shu Lin Wee, "Slow recoveries and unemployment traps: Monetary policy in a time of hysteresis," *The Economic Journal*, 2022, *132* (646), 2007–2047.

[30] See Alesina and Perotti (1995) for an overview.

[31] See Fraccaroli and Whitworth (2020) for a recent overview and empirical test.

Aguiar, Mark and Manuel Amador, " Growth in the Shadow of Expropriation *," *The Quarterly Journal of Economics*, 06 2011, *126* (2), 651–697.

Alesina, Alberto and Roberto Perotti, "The Political Economy of Budget Deficits," *Staff Papers (International Monetary Fund)*, 1995, *42* (1).

Aruoba, S Borağan and Thomas Drechsel, "Identifying monetary policy shocks: A natural language approach," Technical Report, National Bureau of Economic Research WP 32417 2024.

Auerbach, Alan J. and Yuriy Gorodnichenko, "Measuring the Output Responses to Fiscal Policy," *American Economic Journal: Economic Policy*, May 2012, *4* (2), 1–27.

Baker, Scott R., Nicholas Bloom, and Steven J. Davis, "Measuring Economic Policy Uncertainty*," *The Quarterly Journal of Economics*, 07 2016, *131* (4), 1593–1636.

Bazzi, Samuel, Robert A Blair, Christopher Blattman, Oeindrila Dube, Matthew Gudgeon, and Richard Peck, "The promise and pitfalls of conflict prediction: evidence from Colombia and Indonesia," *Review of Economics and Statistics*, 2022, *104* (4), 764–779.

Beaudry, Paul, Thomas J Carter, and Amartya Lahiri, "The Central Bank's Dilemma: Look Through Supply Shocks or Control Inflation Expectations?," Technical Report, National Bureau of Economic Research 2023.

Bell, Christine and Sanja Badanjak, "Introducing PA-X: A new peace agreement database and dataset," *Journal of Peace Research*, 2019, *56* (3).

Berman, Eli, Joseph H. Felter, Jacob N. Shapiro, and Erin Troland, "Modest, secure, and informed: Successful development in conflict zones," *American Economic Review*, 2013, *103* (3), 512–517.

Berman, Nicolas, Mathieu Couttenier, Dominic Rohner, and Mathias Thoenig, "This mine is mine! How minerals fuel conflicts in Africa," *American Economic Review*, 2017, *107* (6), 1564–1610.

Bernanke, Ben S and Frederic S Mishkin, "Inflation targeting: a new framework for monetary policy?," *Journal of Economic perspectives*, 1997, *11* (2), 97–116.

Besley, Timothy and Hannes Mueller, “Fiscal Capacity and State Fragility,” in Raphael Espinoza Ralph Chami and Peter Montiel, eds., *Macroeconomic Policy in Fragile States*, Oxford University Press, 2021.

_ **and Thorsten Persson**, *Pillars of Prosperity: The Political Economics of Development Clusters*, Princeton University Press, 2011.

_ **and Torsten Persson**, “Political violence and state capacity,” *Journal of the European Economic Association*, 2011, *9* (3), 551–579.

_ **and** _ , “The Logic of Political Violence*,” *The Quarterly Journal of Economics*, 08 2011, *126* (3), 1411–1445.

Bianchi, Javier, “Efficient bailouts?,” *American Economic Review*, 2016, *106* (12), 3607–3659.

Blattman, Christopher, *Why we fight: The roots of war and the paths to peace*, Penguin, 2023.

_ **and Edward Miguel**, “Civil War,” *Journal of Economic Literature*, March 2010, *48* (1), 3–57.

Blattmann, Christopher and Jeannie Annan, “Can Employment Reduce Lawlessness and Rebellion? A Field Experiment with High-Risk Men in a Fragile State,” *American Political Science Review*, 2016, *110* (1), 1–17.

Chadefaux, Thomas, “Early warning signals for war in the news,” *Journal of Peace Research*, 2014, *51* (1), 5–18.

Chalmers, Malcolm G., “Spending to save? The cost-effectiveness of conflict prevention,” *Defence and Peace Economics*, 2007, *18* (1), 1–23.

Collier, Paul et al., *Breaking the conflict trap: Civil war and development policy*, Vol. 41181, World Bank Publications, 2003.

Coulombe, Philippe Goulet, “The macroeconomy as a random forest,” *Journal of Applied Econometrics*, 2024, *39* (3), 401–421.

Couttenier, Mathieu, Julian Marcoux, Thierry Mayer, and Mathias Thoenig, “The Gravity of Violence,” Technical Report 2024.

Crost, Benjamin, Joseph Felter, and Patrick Johnston, "Aid under fire: Development projects and civil conflict," *American Economic Review*, 2014, *104* (6), 1833–1856.

_, _, and _, "Conditional cash transfers, civil conflict and insurgent influence: Experimental evidence from the Philippines," *Journal of Development Economics*, 2016, *118*, 171–182.

Davies, Shawn, Therése Pettersson, and Magnus Öberg, "Organized violence 1989–2022, and the return of conflict between states," *Journal of peace research*, 2023, *60* (4), 691–708.

Depetris-Chauvin, Emilio, Ruben Durante, and Filipe Campante, "Building Nations through Shared Experiences: Evidence from African Football," *American Economic Review*, May 2020, *110* (5), 1572–1602.

Dube, Oeindrila and Juan F Vargas, "Commodity Price Shocks and Civil Conflict: Evidence from Colombia," *The Review of Economic Studies*, 2013, *80* (4), 1384–1421.

Esteban, Joan, Laura Mayoral, and Debraj Ray, "Ethnicity and Conflict: An Empirical Study," *American Economic Review*, 2012, *102* (4), 1310–1342.

Falcone, Stefano and Michele Rosenberg, "Agricultural Modernization and Land Conflict," *BSE Working Paper*, 2022, (1314).

Fatás, Antonio and Sanjay R Singh, "Supply or demand? Policy makers' confusion in the presence of hysteresis," *European Economic Review*, 2024, *161*, 104617.

Federle, Jonathan, André Meier, Gernot Müller, Wili Mutschler, and Moritz Schularick, "The Price of War," Discussion Paper DP18834, Centre for Economic Policy Research (CEPR) 2024.

Fetzer, Thiemo, "Can Workfare Programs Moderate Conflict? Evidence from India," *Journal of the European Economic Association*, 02 2020, *18* (6), 3337–3375.

_ and Stephan Kyburz, "Cohesive Institutions and Political Violence," *The Review of Economics and Statistics*, 01 2022, pp. 1–46.

for Afghanistan Reconstruction SIGAR, Special Inspector General, "Quarterly Report to the United States Congress, Section 2: Funding," 2021. Accessed: 2024-09-25.

Francois, Patrick, Ilia Rainer, and Francesco Trebbi, "How is power shared in Africa?," *Econometrica*, 2015, *83* (2), 465–503.

Goldstone, Jack A, Robert H Bates, David L Epstein, Ted Robert Gurr, Michael B Lustik, Monty G Marshall, Jay Ulfelder, and Mark Woodward, "A global model for forecasting political instability," *American journal of political science*, 2010, *54* (1), 190–208.

Hansen, Stephen and Michael McMahon, "Shocking language: Understanding the macroeconomic effects of central bank communication," *Journal of International Economics*, 2016, *99*, S114–S133. 38th Annual NBER International Seminar on Macroeconomics.

Hassan, Tarek A., Stephan Hollander, Laurence van Lent, and Ahmed Tahoun, "Firm-Level Political Risk: Measurement and Effects," *The Quarterly Journal of Economics*, 2019, *134* (4), 2135–2202.

Hegre, Håvard, Forogh Akbari, Mihai Croicu, James Dale, Tim Gåsste, Remco Jansen, Peder Landsverk, Maxine Leis, Angelica Lindqvist-McGowan, Hannes Mueller et al., "Forecasting fatalities," 2022.

Jeanne, Olivier and Anton Korinek, "Macroprudential regulation versus mopping up after the crash," *The Review of Economic Studies*, 2020, *87* (3), 1470–1497.

Joseph, Andreas, Galina Potjagailo, Chiranjit Chakraborty, and George Kapetanios, "Forecasting UK inflation bottom up," *International Journal of Forecasting*, 2024, *40* (4), 1521–1538.

Jugl, I., F. Lösel, D. Bender, and S. King, "Psychosocial Prevention Programs against Radicalization and Extremism: A Meta-Analysis of Outcome Evaluations," *The European Journal of Psychology Applied to Legal Context*, 2020, *13* (1), 37–46.

Kleinberg, Jon, Jens Ludwig, Sendhil Mullainathan, and Ziad Obermeyer, "Prediction Policy Problems," *American Economic Review*, May 2015, *105* (5), 491–95.

Laurent-Lucchetti, Jérémy, Dominic Rohner, and Mathias Thoenig, "Ethnic conflict and the informational dividend of democracy," *Journal of the European Economic Association*, 2024, *22* (1), 73–116.

León, Gianmarco and Edward Miguel, "Risky transportation choices and the value of a statistical life," *American Economic Journal: Applied Economics*, 2017, *9* (1), 202–228.

Liu, Yang, "Mending the Crystal Ball: Enhanced Inflation Forecasts with Machine Learning," *IMF Working Papers*, 9 2024, *2024* (206), 1.

Manacorda, Marco and Andrea Tesei, "Liberation Technology: Mobile Phones and Political Mobilization in Africa," *Econometrica*, 2020, *88* (2), 533–567.

Margalef, Joan and Hannes Mueller, "Caught in a Trap: The Economic Consequences of the Conflict Trap," Technical Report 2023.

Montalvo, José G and Marta Reynal-Querol, "Ethnic polarization, potential conflict, and civil wars," *American economic review*, 2005, *95* (3), 796–816.

Moreno-Pino, Fernando, Emese Sükei, Pablo M Olmos, and Antonio Artés-Rodríguez, "PyHHMM: A Python library for heterogeneous hidden Markov models," *arXiv preprint arXiv:2201.06968*, 2022.

Mueller, Hannes and Christopher Rauh, "Reading between the lines: Prediction of political violence using newspaper text," *American Political Science Review*, 2018, *112* (2), 358–375.

_ and _ , "The hard problem of prediction for conflict prevention," *Journal of the European Economic Association*, 2022, *20* (6), 2440–2467.

_ and _ , "Building bridges to peace: A quantitative evaluation of power-sharing agreements," *Economic Policy*, April 2024, *39* (118), 411–467.

_ , _ , and Alessandro Ruggieri, "Dynamic Early Warning and Action Model," 2022.

_ , _ , and Ben Seimon, "Introducing a global dataset on conflict forecasts and news topics," *Data & Policy*, 2024, *6*, e17.

— , — , **Benjamin Seimon, and Raphael Espinoza**, "The Urgency of Conflict Prevention – A Macroeconomic Perspective," *IMF Working Paper*, 2024.

Müller, Karsten and Carlo Schwarz, "Fanning the Flames of Hate: Social Media and Hate Crime," *Journal of the European Economic Association*, 2021, *19* (4), 2131–2167.

Nicolò, Sowerbutts Rhiannon Fraccaroli and Andrew Whitworth, "Does Regulatory and Supervisory Independence Affect Financial Stability?," *Bank of England Working Paper*, 2020, *2020* (893).

Nunn, Nathan and Nancy Qian, "US Food Aid and Civil Conflict," *American Economic Review*, 2014, *104* (6), 1630–1666.

Ochs, Adrian, "A New Monetary Policy Shock with Text Analysis," Technical Report, Cambridge Working Papers in Economics 2148 2021.

— **and Christian Rörig**, "State-Dependent Macroeconomic Policy Effects: A Varying-Coefficient VAR," 2022.

OECD, "OECD.Stat: Creditor Reporting System (CRS)," `https://stats.oecd.org/Index.aspx?DataSetCode=CRS1` 2022. Accessed: June 19, 2022.

Ramey, V.A., "Chapter 2 - Macroeconomic Shocks and Their Propagation," in John B. Taylor and Harald Uhlig, eds., *John B. Taylor and Harald Uhlig, eds.*, Vol. 2 of *Handbook of Macroeconomics*, Elsevier, 2016, pp. 71–162.

Ramey, Valerie A., "Identifying Government Spending Shocks: It's all in the Timing*," *The Quarterly Journal of Economics*, 02 2011, *126* (1), 1–50.

Ree, Joppe De and Eleonora Nillesen, "Aiding violence or peace? The impact of foreign aid on the risk of civil conflict in sub-Saharan Africa," *Journal of Development Economics*, 2009, *88* (2), 301–313.

Rohner, Dominic, "Mediation, Military and Money: The Promises and Pitfalls of Outside Interventions to End Armed Conflicts," *Journal of Economic Literature*, 2024, *62*, 155–195.

— **and Mathias Thoenig**, "The Elusive Peace Dividend of Development Policy: From War Traps to Macro Complementarities," *Annual Review of Economics*, 2021, *13* (1), 111–131.

Sims, Christopher A., "Macroeconomics and Reality," *Econometrica*, 1980, *48* (1), 1–48.

Sonno, Tommaso, "Globalization and Conflicts: the Good, the Bad and the Ugly of Corporations in Africa," *Economic Journal*, 2024. Conditionally accepted.

Sundberg, Ralph and Erik Melander, "Introducing the UCDP georeferenced event dataset," *Journal of peace research*, 2013, *50* (4), 523–532.

Tan, Timothy M, Paul Spiegel, Christopher Haskew, and P Gregg Greenough, "Does spending on refugees make a difference? A cross-sectional study of the association between refugee program spending and health outcomes in 70 sites in 17 countries," *Conflict and Health*, 2016, *10*, 1–11.

Tetlock, Paul C., "Giving Content to Investor Sentiment: The Role of Media in the Stock Market," *The Journal of Finance*, 2007, *62* (3), 1139–1168.

Thoenig, Mathias, "Trade policy in the shadow of war: A quantitative toolkit for geoeconomics," 2023.

UNHCR, "UNHCR Refugee Data," https://www.unhcr.org/refugee-statistics/ 2022. Data on refugees, asylum-seekers, and internally displaced persons (IDPs).

United Nations Secretary-General, "Maintenance of International Peace and Security: Conflict Prevention and Sustaining Peace," January 2017. Remarks delivered at the Security Council Open Debate. Available online: https://www.un.org/sg/en/content/sg/statement/2017-01-10/secretary-generals-remarks-security-council-open-debate-maintenance.

Weitzman, Martin L, "A review of the Stern Review on the economics of climate change," *Journal of economic literature*, 2007, *45* (3), 703–724.

World Bank, "World Development Indicators," `https://databank.worldbank.org/source/world-development-indicators` 2022. Indicators collected: Total Population (SP.POP.TOTL), GDP at constant 2015 US dollars (NY.GDP.MKTP.KD), GDP per capita at constant 2015 US dollars (NY.GDP.PCAP.KD).

Yanagizawa-Drott, David, " Propaganda and Conflict: Evidence from the Rwandan Genocide *," *The Quarterly Journal of Economics*, 11 2014, *129* (4), 1947–1994.

_ , "Propaganda and Conflict: Evidence from the Rwandan Genocide," *The Quarterly Journal of Economics*, 2014, *129* (4), 1947–1994.

A Model Appendix

A.1 Preliminaries for Proof of Propositions 1 and 2

In this section, we solve for steady-state values in the two-stage model. Let V_1 and V_2 be the present discounted value of stages 1 and 2 when the optimal intervention policies are implemented. With a slight abuse of notation, denote $V_1(\phi_1, \phi_2)$ and $V_2(\phi_1, \phi_2)$ the present discounted values in stages 1 and 2 conditional on implementing policy $\phi_1 \in \{0, 1\}$ in stage 1 and $\phi_2 \in \{0, 1\}$ in stage 2.

Therefore, the values of stages 1 and 2 conditional on not implementing any policy in both states, $(\phi_1 = 1, \phi_2 = 0)$ are equal to:

$$\begin{aligned} V_1(\phi_1 = 0, \phi_2 = 0) &= \beta \left(p_{11} V_1 + (1 - p_{11}) V_2 \right) = \beta \frac{(1 - p_{11}) V_2}{1 - \beta p_{11}} \\ V_2(\phi_1 = 0, \phi_2 = 0) &= -D + \beta \left((1 - p_{22}) V_1 + q V_2 \right) = \frac{-D + \beta (1 - p_{22}) V_1}{1 - \beta q} \end{aligned}$$

which yields

$$\begin{aligned} V_1(\phi_1 = 0, \phi_2 = 0) &= \frac{\beta (1 - p_{11}) (-D)}{(1 - \beta) (1 + \beta - p_{11}\beta - q\beta)} \\ V_2(\phi_1 = 0, \phi_2 = 0) &= \frac{(1 - p_{11}\beta) (-D)}{(1 - \beta) (1 + \beta - p_{11}\beta - q\beta)} \end{aligned}$$

The values of stages 1 and 2 conditional on implementing only prevention policy, $(\phi_1 = 1, \phi_2 = 0)$, are equal to:

$$\begin{aligned} V_1(\phi_1 = 1, \phi_2 = 0) &= -I_1 + \beta \left(\left(p_{11} + \tau (1 - p_{11}) \right) V_1 + \left(1 - p_{11} - \tau (1 - p_{11}) \right) V_2 \right) \\ &= \frac{-I_1 + \beta \left(1 - \left(p_{11} + \tau (1 - p_{11}) \right) \right) V_2}{1 - \beta \left(p_{11} + \tau (1 - p_{11}) \right)} \\ V_2(\phi_1 = 1, \phi_2 = 0) &= \frac{-D + \beta (1 - p_{22}) V_1}{1 - \beta q} \end{aligned}$$

or

$$\begin{aligned}
V_1(\phi_1 = 1, \phi_2 = 0) &= \frac{\beta\left(1 - \left(p_{11} + \tau\left(1 - p_{11}\right)\right)\right)(-D) + \left(1 - \beta q\right)(-I_1)}{(1-\beta)\left(1 + \beta - \left(p_{11} + \tau\left(1 - p_{11}\right)\right)\beta - q\beta\right)} \\
V_2(\phi_1 = 1, \phi_2 = 0) &= \frac{\left(1 - \beta\left(p_{11} + \tau\left(1 - p_{11}\right)\right)\right)(-D) + \beta\left(1 - p_{22}\right)(-I_1)}{(1-\beta)\left(1 + \beta - \left(p_{11} + \tau\left(1 - p_{11}\right)\right)\beta - q\beta\right)}
\end{aligned}$$

The values of stages 1 and 2 conditional on implementing only de-escalation policy, $(\phi_1 = 0, \phi_2 = 1)$, are equal to:

$$\begin{aligned}
V_1(\phi_1 = 0, \phi_2 = 1) &= \beta\frac{\left(1 - p_{11}\right)V_2}{1 - \beta p_{11}} \\
V_2(\phi_1 = 0, \phi_2 = 1) &= -D - I_2 + \beta\left(\left(1 - p_{22} + \tau\left(1 - p_{22}\right)\right)V_1 + \left(q - \tau_2\right)V_2\right) \\
&= \frac{-D - I_2 + \beta\left(1 - p_{22} + \tau\left(1 - p_{22}\right)\right)V_1}{1 - \beta\left(q - \tau\left(1 - p_{22}\right)\right)}
\end{aligned}$$

so that

$$\begin{aligned}
V_1(\phi_1 = 0, \phi_2 = 1) &= \frac{\beta\left(1 - p_{11}\right)\left(-D - I_2\right)}{(1-\beta)\left(1 + \beta - p_{11}\beta - \left(q - \tau\left(1 - p_{22}\right)\right)\beta\right)} \\
V_2(\phi_1 = 0, \phi_2 = 1) &= \frac{\left(1 - \beta p_{11}\right)\left(-D - I_2\right)}{(1-\beta)\left(1 + \beta - p_{11}\beta - \left(q - \tau\left(1 - p_{22}\right)\right)\beta\right)}
\end{aligned}$$

The values of stages 1 and 2 conditional on implementing both prevention and de-escalation policies $(\phi_1 = 1, \phi_2 = 1)$, are equal to:

$$\begin{aligned}
V_1(\phi_1 = 1, \phi_2 = 1) &= \frac{\beta\left(1 - \left(p_{11} + \tau\left(1 - p_{11}\right)\right)\right)\left(-D - I_2\right) + \left(1 - \beta\left(q - \tau\left(1 - p_{22}\right)\right)\right)(-I_1)}{(1-\beta)\left(1 + \beta - \left(p_{11} + \tau_1\right)\beta - \left(q - \tau\left(1 - p_{22}\right)\right)\beta\right)} \\
V_2(\phi_1 = 1, \phi_2 = 1) &= \frac{\left(1 - \beta\left(p_{11} + \tau\left(1 - p_{11}\right)\right)\right)\left(-D - I_2\right) + \beta\left(1 - \beta\left(p_{22} - \tau\left(1 - p_{22}\right)\right)\right)(-I_1)}{(1-\beta)\left(1 + \beta - \left(p_{11} + \tau_1\right)\beta - \left(p_{22} - \tau\left(1 - p_{22}\right)\right)\beta\right)}.
\end{aligned}$$

If we insert the present values without de-escalation and without prevention into the optimal intervention conditions (1) we get

$$I_1 < \beta\tau\left(1 - p_{11}\right)\frac{D}{1 + \beta - p_{11}\beta - p_{22}\beta}$$

and the necessary condition for de-escalation without prevention is

$$I_2 < \beta\tau\left(1 - p_{22}\right)\frac{D}{1 + \beta - p_{11}\beta - p_{22}\beta}.$$

A.2 Proof of Proposition 1

Proof. Prevention with de-escalation is optimal, iff

$$
\begin{aligned}
I_1 &< \beta\tau\left(1-p_{11}\right)\left(V_1-V_2\right) \\
I_1 &< \beta\tau\left(1-p_{11}\right)\left(\frac{D+I_2}{1+\beta-p_{11}\beta-\left(p_{22}-\tau_2\right)\beta}\right)
\end{aligned}
$$

which is harder to satisfy than prevention without de-escalation if

$$
\left(\frac{D+I_2}{1+\beta-p_{11}\beta-\left(p_{22}-\tau\left(1-p_{22}\right)\right)\beta}\right)<\left(\frac{D}{1+\beta-p_{11}\beta-p_{22}\beta}\right)
$$

or

$$
I_2<\beta\tau\left(1-p_{22}\right)\frac{D}{1+\beta-p_{11}\beta-p_{22}\beta}
$$

which is the condition for de-escalation without prevention. This completes the proof that de-escalation crowds out prevention. □

A.3 Proof of Proposition 2

Proof. De-escalation without prevention if

$$
I_2<\beta\tau\left(1-p_{22}\right)\left(\frac{D}{1+\beta-p_{11}\beta-p_{22}\beta}\right)
$$

and with prevention

$$
I_2<\beta\tau\left(1-p_{22}\right)\left(\frac{D-I_1}{1+\beta-\left(p_{11}+\tau_1\right)\beta-p_{22}\beta}\right)
$$

This is easier to satisfy if

$$
\frac{D-I_1}{1+\beta-\left(p_{11}+\tau\left(1-p_{11}\right)\right)\beta-p_{22}\beta}>\frac{D}{1+\beta-p_{11}\beta-p_{22}\beta}
$$

or

$$
\beta\tau_1\frac{D}{1+\beta-p_{11}\beta-p_{22}\beta}>I_1
$$

which is the condition for prevention. This completes the proof that prevention crowds in de-escalation. □

B Hidden Markov Model

We opt to use a Hidden Markov Model (HMM) that describes five panel series through twelve latent states. We use five different panel series in order to capture different measures of conflict risk and intensity. The five panel series are two forecasts of armed conflict risk (the 'full model' uses past violence and newspaper text and the 'text model' relies on newspaper text only), the intensity forecast (based on past violence and newspaper text), observed current log(1+ fatalities per 1 million inhabitants), and an indicator of the outbreak of armed conflict regressed on a polynomial of degree three of the number of months since the last conflict.[1] The latter input reinforces the high risk in the months following armed conflict that decays exponentially.

The HMM assumes that these five observable variables are generated by a sequence of internal hidden states S which we call *risk stages*. The hidden states are not observed directly. The transitions between hidden states are assumed to be a (first-order) Markov chain. They can be specified by a start probability vector and a transition matrix $\mathbf{T}$. In our case, the emission probabilities (i.e. the probability of an observation being generated by a stage s_i) are assumed to be Gaussian. The state space of the HMMs is discrete, hence the key hyperparameter is the number of states. We estimate the model on the entire data from 2010m1 to 2024m2.

We solve the dynamic programming problem involved in backing out the states through the combination of two algorithms called the Viterbi algorithm, a Forward-Backward algorithm, and an Expectation-Maximization (EM) algorithm, known as the Baum-Welch algorithm. We use the PyHHMM package in Python. For more technical details see the PyHHMM documentation and Moreno-Pino et al. (2022).

[1] The values are derived from the following model:

$$\hat{y} = \begin{cases} max(\hat{\beta}_0 + \hat{\beta}_1 x + \hat{\beta}_2 x^2 + \hat{\beta}_3 x^3, 0), & \text{if the country is not in armed conflict,} \\ 1, & \text{if the country is in armed conflict.} \end{cases}$$

where x represents the number of months since the country last experienced armed conflict. In the OLS regression used to derive the $\hat{\beta}$, the dependent variable is a binary indicator for the onset of an armed conflict within the next twelve months. The point estimates for $\hat{\beta}_0$, $\hat{\beta}_1$, $\hat{\beta}_2$, and $\hat{\beta}_3$ are 0.2853, -0.0028, 9.401e-06 and -1.057e-08 respectively. Note that where a country is deemed to be in armed conflict in that month, we set $\hat{y}$ equal to 1.

Figure B1: Distributions of variables entering HMM

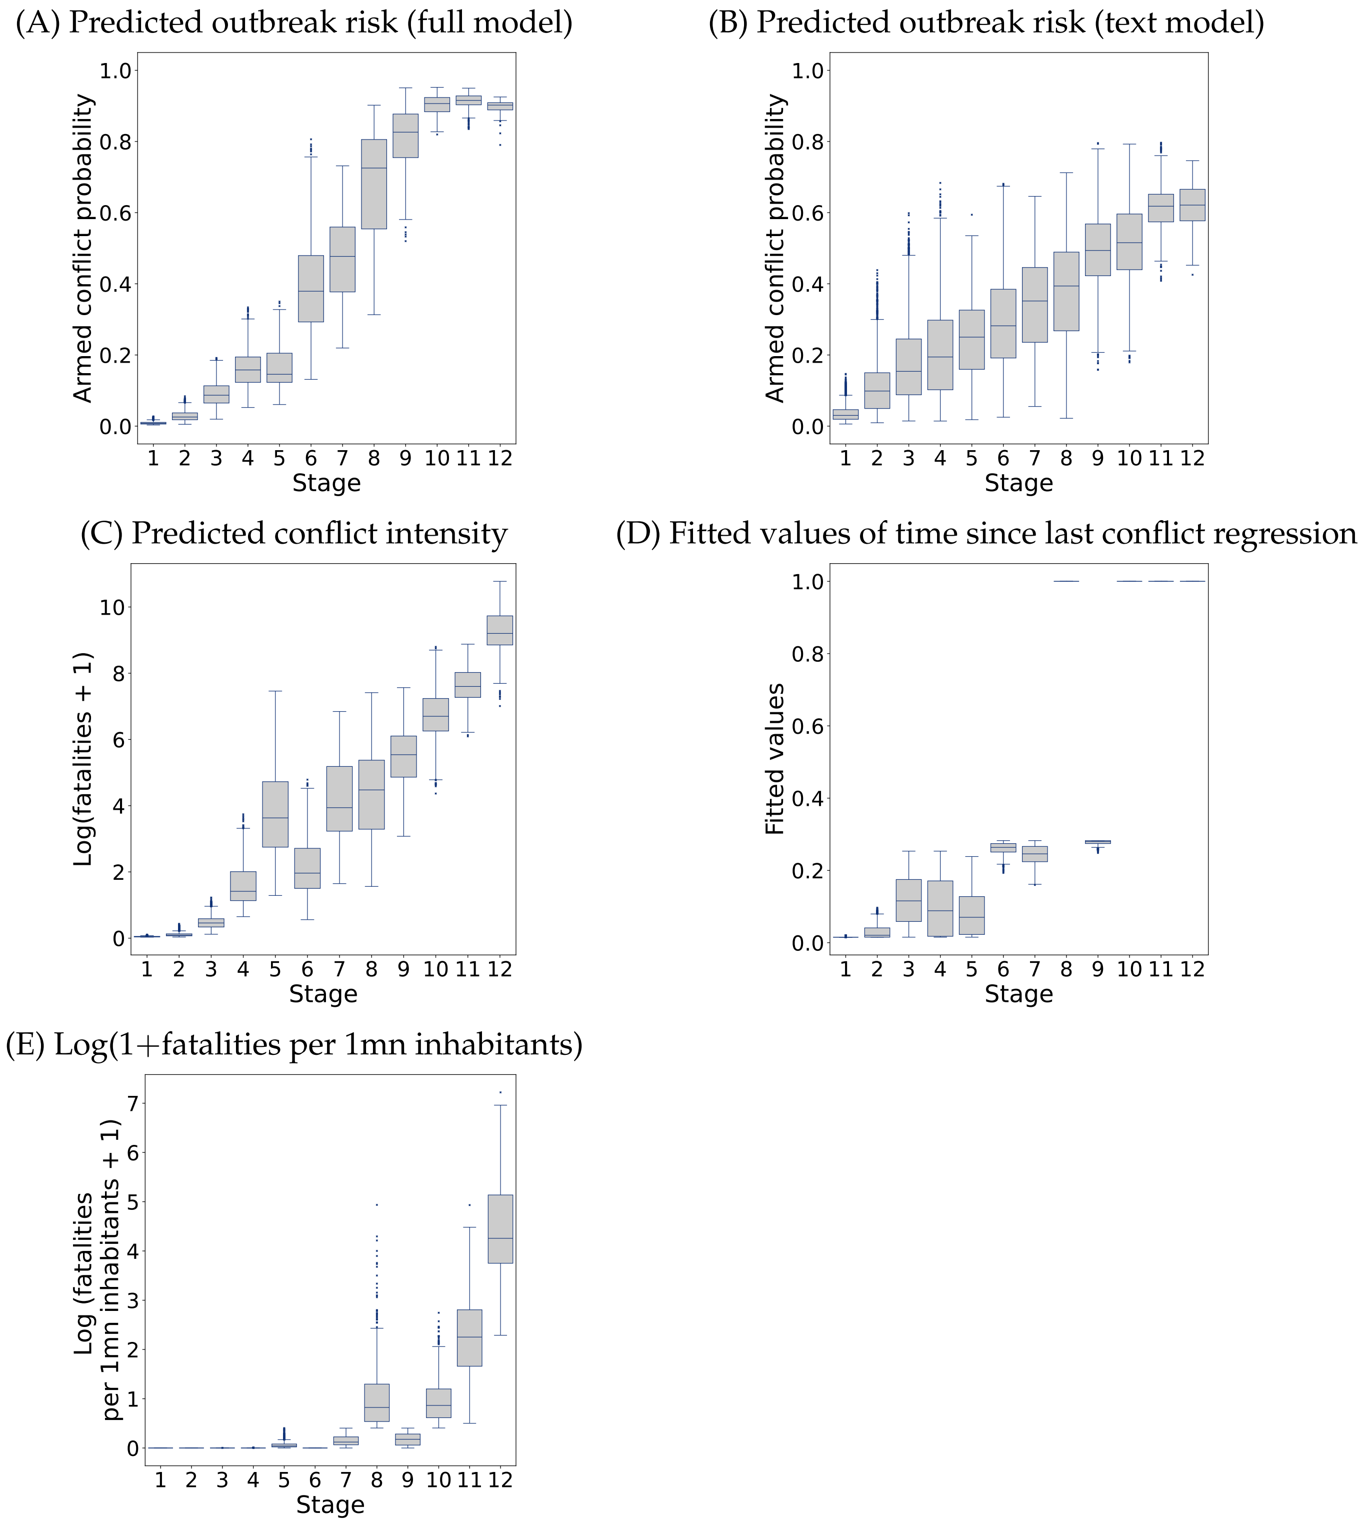

NOTES: Figure shows the distribution of the features conditional on the stage. The box plots display the minimum, first quartile, median, third quartile, and maximum and outliers are shown as individual points beyond the whiskers. Panel (D) shows fitted values OLS regressions with armed conflict in the next month as the dependent variable of the OLS. The functional form is $\hat{y} = max(\hat{\beta}_0 + \hat{\beta}_1 x + \hat{\beta}_2 x^2 + \hat{\beta}_3 x^3, 0)$ if the country is not in armed conflict, where x represents the number of months since the country last experienced armed conflict, and $\hat{y} = 1$ if the country is in armed conflict.

Figure B2: Predictive performance

(A) Receiver-operating curve **(B)** Precision-recall curve

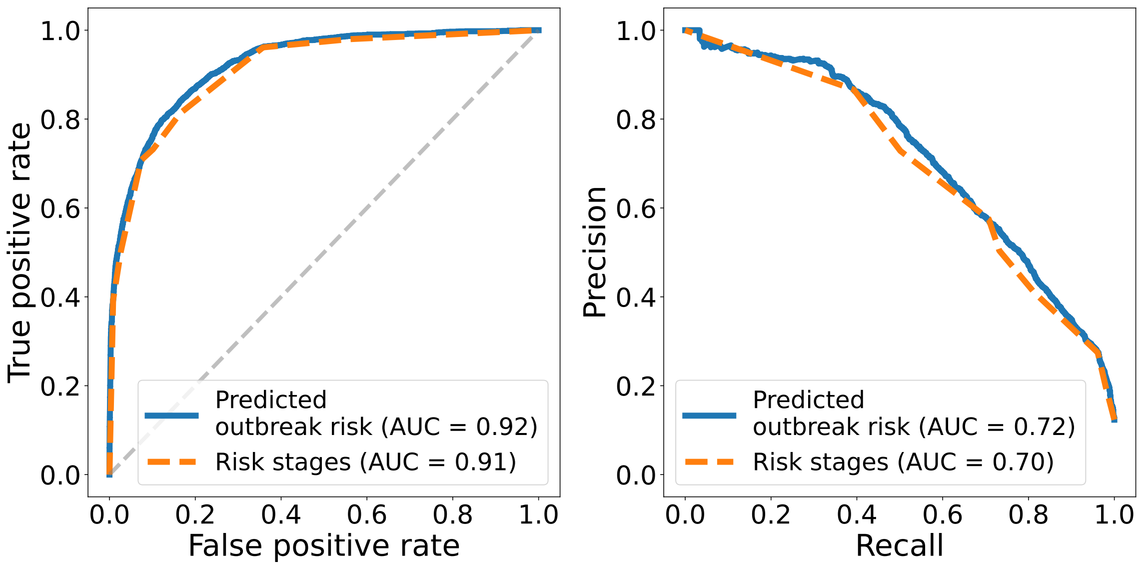

NOTES: The blue solid line shows the predictive performance of the continuous predicted outbreak probability by `conflictforecast.org`. The orange dashed line shows the predictive performance using the twelve risk stages. Predictive performance is summarized by the trade-offs between true positives (TP), false positives (FP), false negatives (FN), and true negatives (TN). Panel A summarizes predictive performance in terms of the false positive rate ($\frac{FP}{FP+TN}$) on the x-axis and true positive rate ($\frac{TP}{TP+FN}$) on the y-axis. In Panel B recall on the x-axis is the true positive rate, i.e. $\frac{TP}{TP+FN}$, and precision is defined as $\frac{TP}{TP+FP}$.

C Calibration

C.1 Damages

In Figure C1 we summarize all four damage categories. In the regression of GDP per capita on stage dummies we control for country-fixed effects and a one-year lag of GDP per capita levels. For ODA, fatalities, and refugees we do not include country-fixed effects as these are directly linked to conflict, so identification is less of an issue.

For the translation of fatalities to monetary costs we use a cost per life of 0.9 million US$ (León and Miguel 2017). Refugees per capita are indicating huge disruptions for stages 11 and 12 with 4% to 13% of the entire population leaving their home country on average in a year. We use a study from Tan et al. (2016) to translate refugee costs into a monetary equivalent. They report a cost of $623 in 2011 prices which we convert to $660 in 2015 prices. This will be an underestimate of the true costs of displacement which entail human suffering, the loss of mental and physical health, and reduced education.

Economic losses from GDP are a key dimension of our loss calculation. Note that we need to weigh the growth loss by the level of GDP to get to a total number. This means that our estimates of the total damages from conflict will be a function of the level of GDP. Furthermore, we need to translate growth numbers into static damages. As there is no significant recovery from conflict these damages need to take into account the reduced capacity of GDP in the future. We therefore calculate the discounted GDP reduction from now to up to 10 years. In this calculation, we do not assume large GDP baseline growth which would increase the total cost significantly.

Figure C1: Estimated damages across risk stages

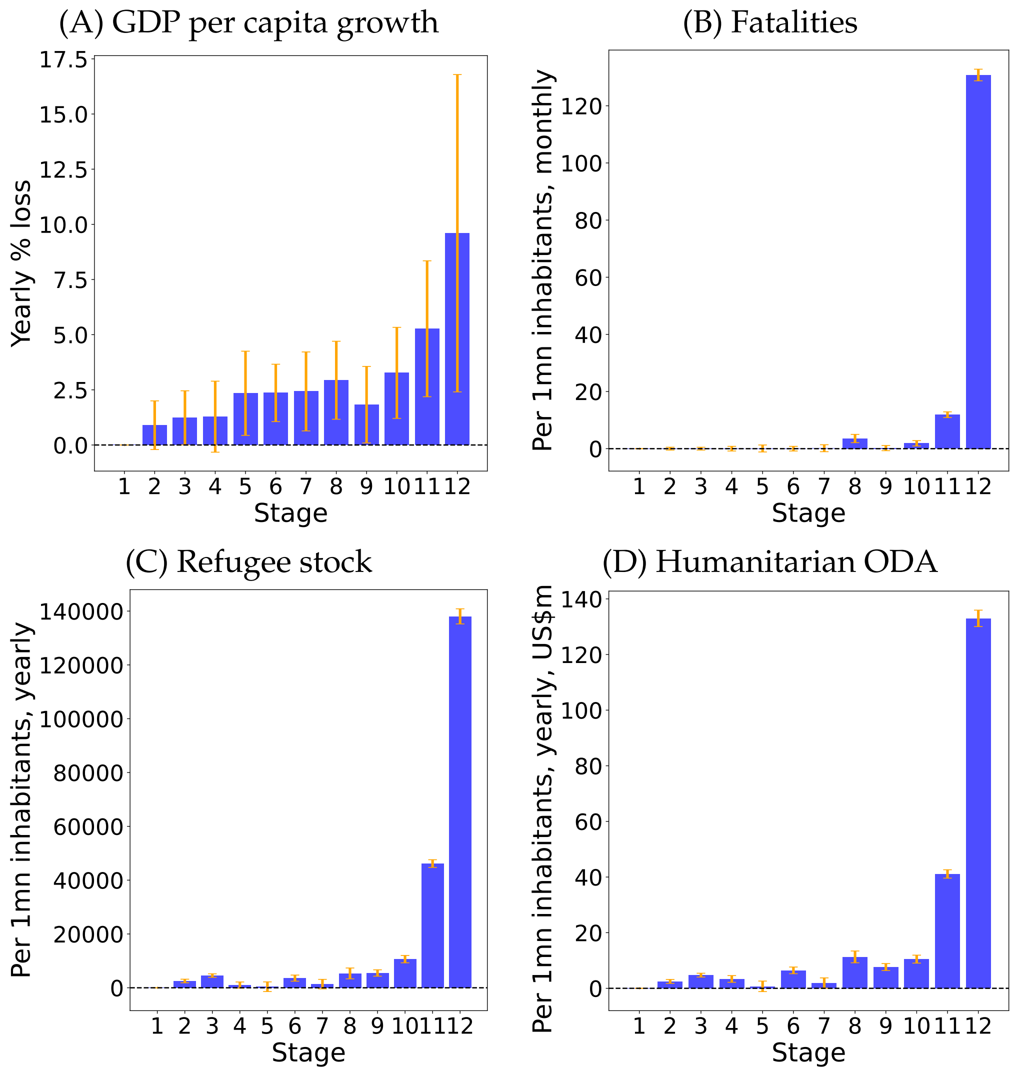

NOTES: The figure shows coefficients from regressing damages on dummies for each stage with 95% confidence intervals. GDP growth loss regressions include country-fixed effects and a one-year lag of GDP per capita levels together. Standard errors are robust and clustered at the country/year level. The data sources for fatalities are UCDP (Sundberg and Melander 2013, Davies et al. 2023), World Bank for GDP growth and population statistics (World Bank 2022), UNHCR for refugees (UNHCR 2022), and the OECD for ODA (OECD 2022).

C.2 Intervention Costs

We rely on a dataset from 2010 to the end of 2023 which contains monthly fatalities (UCDP), the average monthly peacekeeping budget (UN), and power-sharing agreements (PA-X version 8). First, we subset to all country/month observations where there is an agreement. We then compute the average peacekeeping spend in the 18 months of an agreement (so the month of the agreement and the following 17 months). We then regress this 18 month average of the peacekeeping budget on the log of fatalities per 1 million inhabitants + 1.[2] Our coefficient estimate indicates

The table shows the regression to derive our cost per fatality estimate for the intervention cost function. This cost is supposed to capture the fact that meaningful policy implementation becomes much more expensive with ongoing violence. We rely on a dataset from 2010 to the end of 2023 which contains monthly fatalities (UCDP), the average monthly peacekeeping budget (UN), and power- sharing agreements (PA-X version 8). First, we subset to all country/month observations where there is an agreement. We then compute the average peacekeeping spend in the 18 months of an agreement (the month of the agreement and the following 17 months). We then regress this 18-month average of the peacekeeping budget on the log(1+fatalities per 1mn inhabitants).

Table C1: Costs in terms of peacekeeping budget per fatality

	(1)
Log (1+fatalities per 1mn)	14.956***
	(2.583)
Time FE	No
Country FE	No
Observations	343
R^2	0.090

Notes: *Notes:* The dependent variable is average 18 month peacekeeping budget. *p<0.1; **p<0.05; ***p<0.01. Standard errors in parentheses.

[2] Appendix Figure C1 shows the estimated coefficients.

C.3 Static and Dynamic Damages by GDP per Capita

Figure C2: Static vs dynamic damages by stage, full information

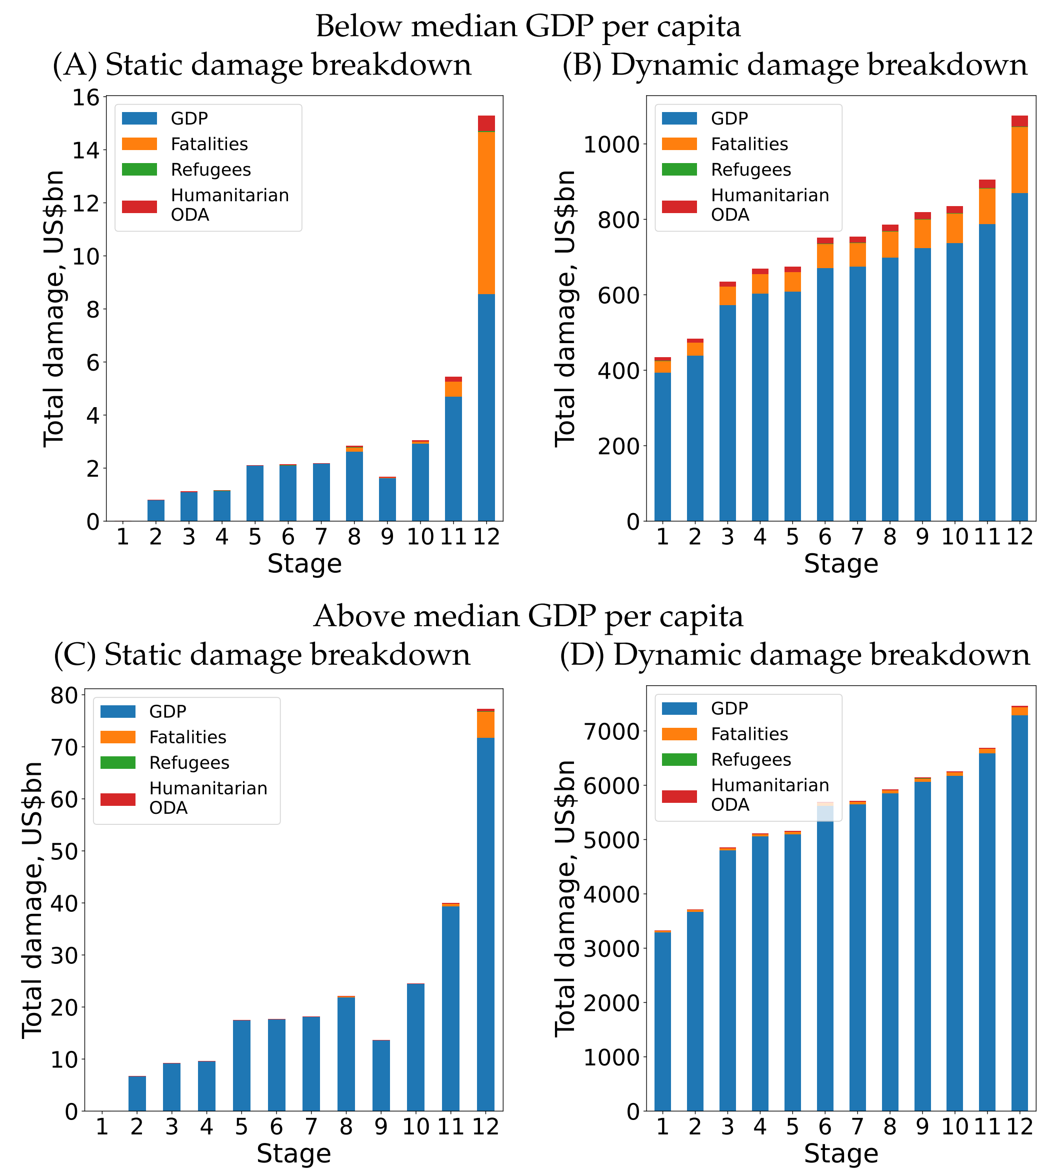

NOTES: The left panel shows the damages in a given period D_s, while the right shows D_s^{PV} defined by Equation (9), which is the present value including all future damages discounted using an annual rate of 4% and assuming a 1% annual GDP growth.

D Additional Figures

D.1 Transition Matrix

For each country/month observation, we sample a stage from the posterior distribution 10,000 times. This gives 10,000 draws of the sequence of stages for each country, from which we can compute 10,000 different transition matrices. Figure D1 shows the resulting standard deviations. Note that in the partial information, any stage draw of 1-5 is assumed to be equivalent to a draw of the super-stage 1.

Figure D1: Standard deviations of transition matrices

(A) Partial information **(B)** Full information

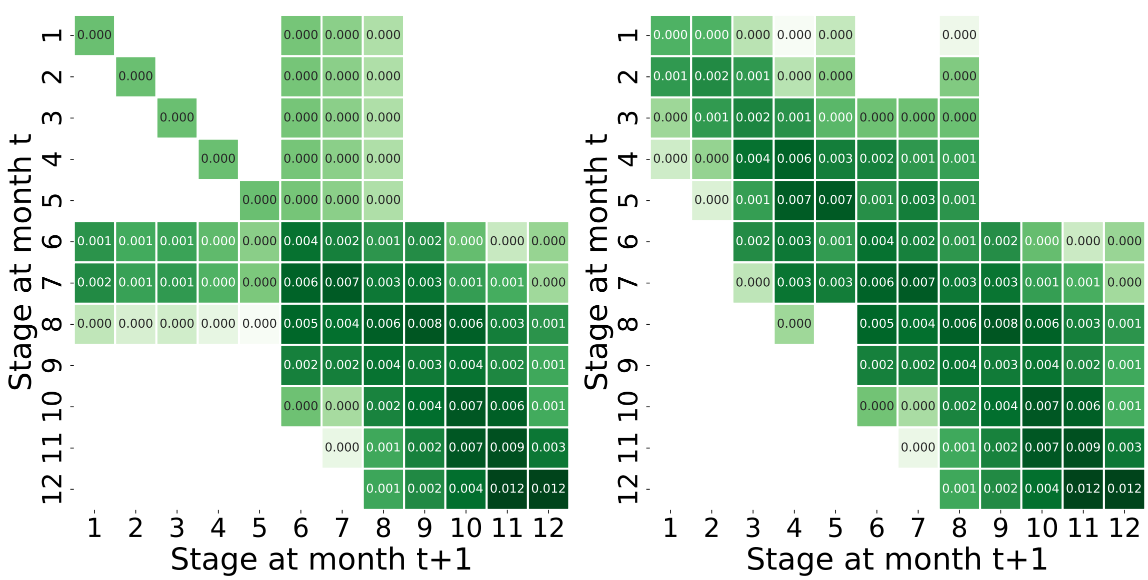

NOTES: The cells show the standard deviations around each estimated transition probabilities.

D.2 Additional Results

Figure D2: Benefit-cost ratios of interventions with full information

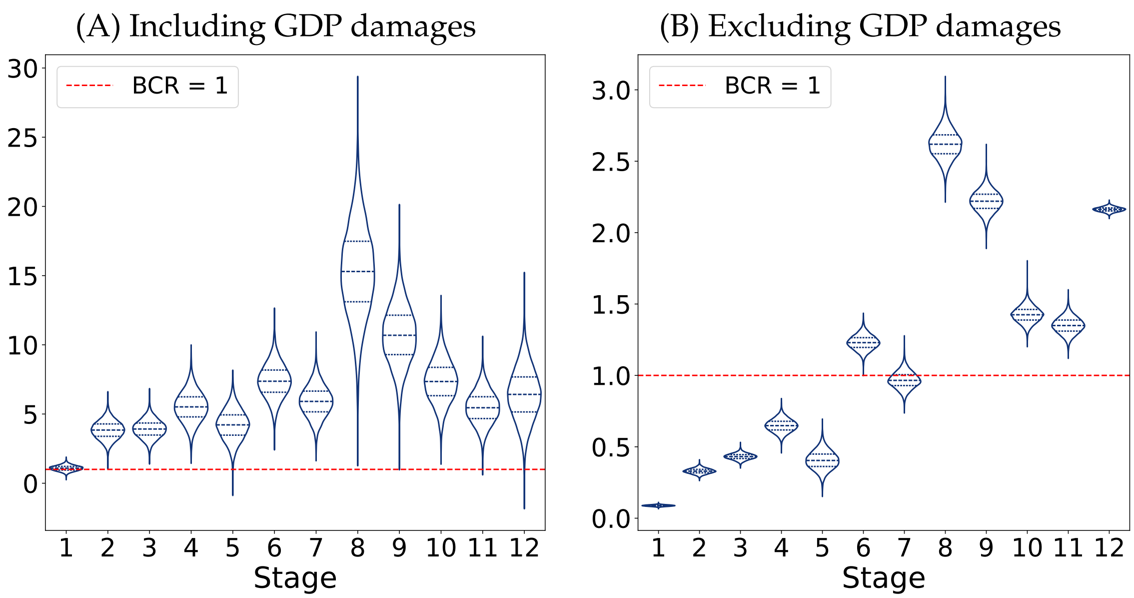

NOTES: Figure shows reports benefit cost ratios (BCRs) from 10,000 simulations across all countries by stage calculated from equation (4.2). In panel (A) GDP damages are included and in panel (B) GDP damages are excluded. For each simulation we draw a transition matrix and damage vector, calculate the BCR for each country, and average across countries. Calculations include GDP damages. Figure reports the distribution of the mean BCR's as a violin plot, where the Scott method is used to compute the kernel bandwidth. Inner lines represent the 25th, 50th, and 75th percentile respectively.

Figure D3: Intervention shares by stage and under partial and full information

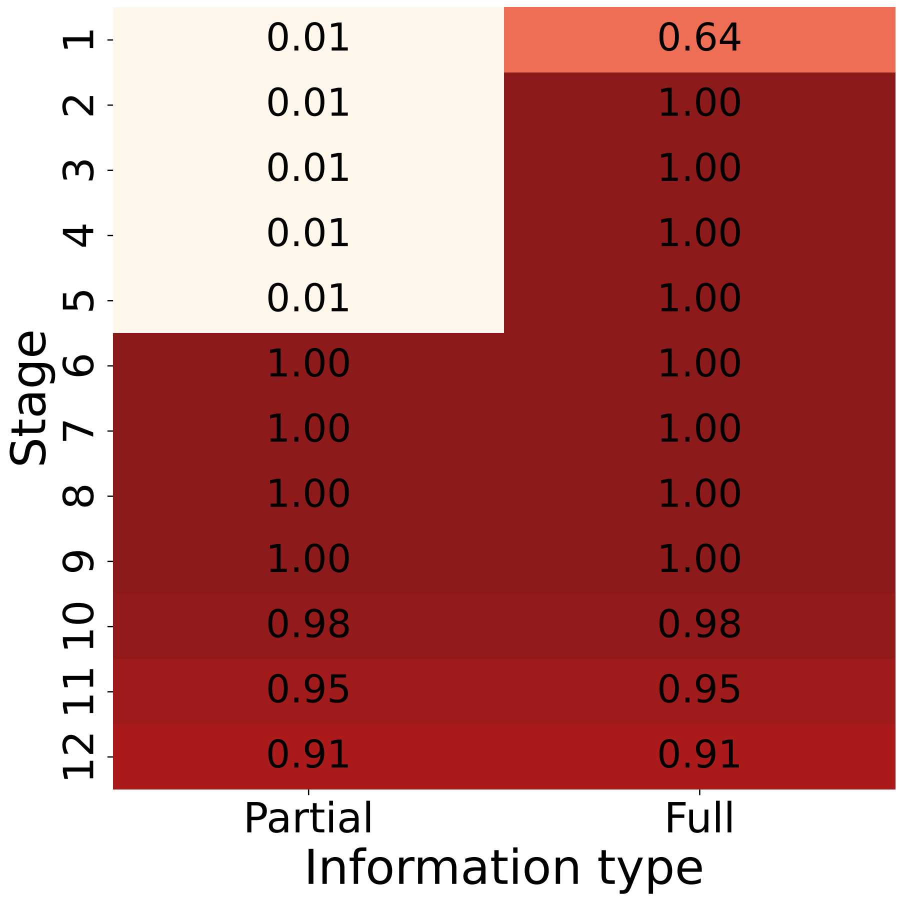

NOTES: The left column of the heatmap shows the share of times countries optimally intervene in partial information across all draws. The right columns show the share of times countries optimally intervene in full information across all draws.

.

Table D.1: Ergodic distributions across stages under full and partial information (%)

	Information	
Stage	Full	Partial
1	45.76	31.56
2	16.02	15.26
3	16.01	19.42
4	5.72	8.30
5	1.96	3.49
6	5.20	7.86
7	1.78	2.70
8	1.28	1.94
9	2.61	3.95
10	1.82	2.75
11	1.48	2.24
12	0.36	0.54

Notes: The percentages represent the ergodic distributions across stages resulting from optimal policy under the two information environments.

E Approximating Effectiveness τ Using Power-Sharing Agreements

The purpose of this appendix is to outline our methodology for calibrating our policy effectiveness parameter τ. We start with our panel of risk stages for 168 countries from Jan 2010 to Feb 2024, monthly UCDP fatalities, and the PA-X power-sharing agreements dataset version 8 (Bell and Badanjak 2019) which contains data from 1990 to the end of 2023. We consider all available power-sharing agreements in the dataset, only filtering out those that are not signed. In total, there are 514 instances of a signed power-sharing agreements in our dataset.

Simulating expected violence with and without policies in place is possible as it requires multiplying the transition matrix with the average fatalities by stage. Every additional multiplication gives the expected number of fatalities for the next month. We first simulate the model without policies in place by multiplying the transition matrix with the damage vector.

We then simulate violence reduction from policy by imposing policies of effectiveness τ in all stages for which we see power-sharing agreements in the data. Denote the transition matrix with interventions as $\mathbf{T}_i$ and the transition matrix without interventions as $\mathbf{T}$ and take the distribution across stages at the moment of an agreement as $\mathbf{S}_0$, and the damage vector (fatalities in different stages) to be $\mathbf{D}$.

Then we calculate at each time step the distributions across stages in the post-treatment period $t > 0$ as

$$\mathbf{S}_{i,t} = \mathbf{S}_{i,t-1}(\mathbf{T}_i)^t$$

and without treatment $\mathbf{S}_{n,t}$:

$$\mathbf{S}_{n,t} = \mathbf{S}_{n,t-1}(\mathbf{T})^t$$

If we multiply these two changing distributions by the damage vector we arrive at the simulated average of fatalities with and without treatment. To calibrate τ we target the value of

$$(\mathbf{S}_{i,t} * \mathbf{D} - \mathbf{S}_{i,n} * \mathbf{D}) / (\mathbf{S}_{n,t} * \mathbf{D})$$

which is the relative reduction in violence with treatment.

Figure E4: Approximating effectiveness

(A) Stages in which power-sharing agreements take place

(B) Violence reduction

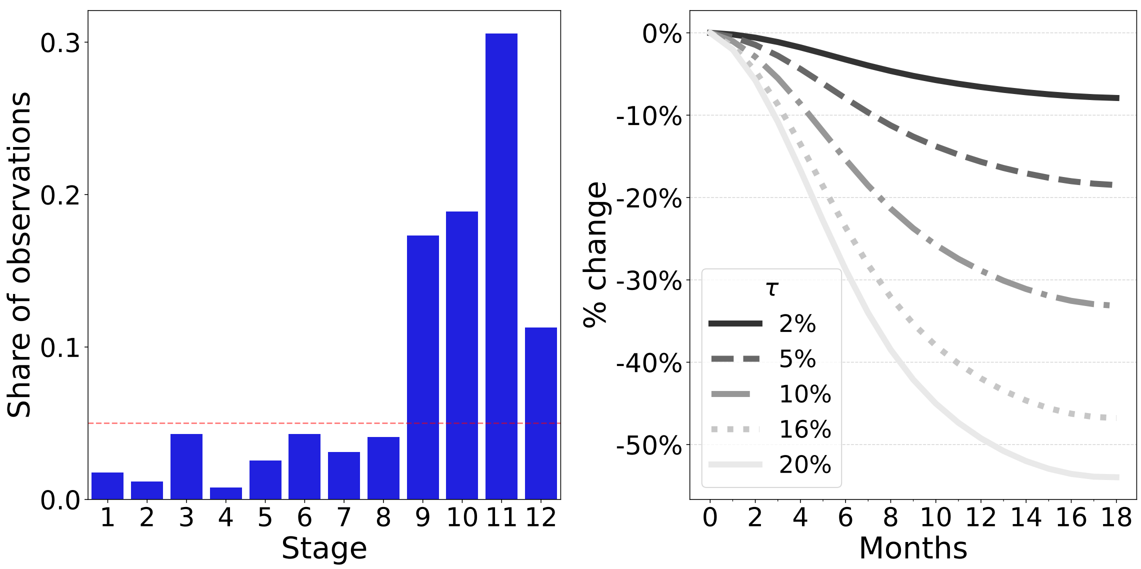

NOTES: Panel A shows the distribution across stages at the moment of a power-sharing agreement. For example, we see that 30% of power-sharing agreements take place in stage 11. Panel B shows the percentage reduction in violence over 18 months, assuming a τ of 2, 5, 10, 16, and 20%. The optimal policy vector is assumed to take a value of 1 where the distribution across stages at the moment of a power-sharing agreement exceeds 5%. In practice, this means the simulation is conducted under the assumption that a policymaker intervenes in stages 9, 10, 11, and 12.

F Prevention in Practice

F.1 Conflict Prevention Policies

The framework we propose offers a dynamic model of both targeting of policies and their effect. A central pillar for this is the huge body of work on conflict prevention and de-escalation policies and their effectiveness. However, we will argue in this section that data and, hence, quantitative research on conflict prevention policies, as we define them, are essentially missing.

An important policy document published on conflict prevention was the joint UN/World Bank report *Pathways to Peace*. The report stresses that affecting conflict outcomes means addressing both short-term triggers of conflict and more structural factors. Examples are changing incentives for violence through institutional reform or by directly changing incentives for specific actors, power-sharing agreements, reforms that equalize spending patterns across regions, mediation and promoting of peaceful narratives. Countries stay peaceful if they ensure broad access to political power and representation, natural resources (in particular, land and extractives), security and justice, and basic services. A lot of this is backed by the academic evidence.[3].

However, when searching both policy reports and the academic literature it becomes clear that there is very little work on policies targeting armed conflict prevention in situations without previous violence. Measured policies related to conflict are typically implemented in the context of ongoing or recent violence. At the same time, there are few well-identified studies in the conflict literature on preventative policies like diplomacy or civil society work. There is, for example, little quantitative work on prevention programs against radicalization despite its relevance (Jugl et al. 2020) or studies that analyze how leaders could be prevented from *playing the ethnic card* despite the fact that this is an extremely well-researched mechanism for internal armed conflict (Montalvo and Reynal-Querol 2005, Esteban et al. 2012, Yanagizawa-Drott 2014a). At the cross-country level there is no systematic data on prevention policies of this sort.

This is not to say that there are no researched ways to make societies less susceptible to armed conflict. We know, for example, that economic development and high state capacity are associated with more stability across countries (Besley and Persson 2011a, Acemoglu and Robinson 2012). This suggests that the simple idea that helps on these two fronts might be an effective conflict prevention strategy. And indeed, there is a significant level of foreign intervention in prevention stages with an economic focus like development aid or IMF programs. However, foreign aid has very mixed support as a preventative measure and IMF programs target sound fiscal policies, not violence prevention. The evidence for development aid preventing conflict is mixed at best.[4] However, some of the contradictory effects can be explained by theory. For example, aid will contribute to instability if it leads

[3]See, for example, Depetris-Chauvin et al. (2020), Fetzer and Kyburz (2022), Mueller and Rauh (2024), Rohner (2024)

[4]See De Ree and Nillesen (2009), Berman et al. (2013), Crost et al. (2014), Nunn and Qian (2014), and Crost et al. (2016) for some encouraging and discouraging examples.

to the generation of rapacity effects or displaces the local population through high demands on land.[5]

Targeting the human capital channel seems more promising. Studies have found that productivity-boosting policies such as employment programs, vocational training, school construction or health-promotion programs have led to a drop in fighting.[6] Given the role attributed to political institutions and state capacity it could well be that aid support for building non-exclusive institutions with strong checks and balances and fiscal capacity have a positive effect on robustness but it is hard to establish clean treatment effects here.[7]

However, many of these policies cannot be implemented fast enough. What would be needed is an approach that builds a long-term prevention strategy with short-term safeguards for emergencies which can be targeted based on forecasts. Given the latest research on the importance of communication for protest and conflict outcomes[8] ethically justifiable policy experiments on this front seem an essential pillar of the fast deployment strategy. The introduction of special financial instruments by the IMF and World Bank and rethinking of how to measure *fragility* is very promising in this regard. We discuss these initiatives in more detail in a companion paper.

Many of these policies have prevention as a side product. We do not have sufficient quantitative research on fast-to-deploy conflict prevention. But we argue in Appendix F.2 that some of the evidence is simply not there because situations for prevention are not systematically identified. This highlights the problem posed by the absence of targeting methods - the formulation of systematic prevention policies is hindered significantly by the absence of an effective targeting mechanisms. This motivates our approach of calibrating our model under the partial information and then using this to explore the viability of a more preventative approach even the detailed formulation of this approach is lacking.

[5] See Dube and Vargas (2013), Berman et al. (2017), Falcone and Rosenberg (2022), Sonno (2024)

[6] See, for example, Blattmann and Annan (2016) and Fetzer (2020). See Rohner (2024) for a theory-led review.

[7] See Besley and Persson (2011c) for a cross-country study and for a detailed discussion of fiscal policies see Besley and Mueller (2021).

[8] See, for example, Yanagizawa-Drott (2014b), Manacorda and Tesei (2020), Müller and Schwarz (2021).

F.2 The Role of Forecasting in Prevention Policy

The size and robustness of prevention benefits and the information rent derived suggest that significant rents could be generated by basing policy on systematic quantitative forecasting, as, for example, in Central Banking. But for this argument to have any relevance, we need to argue that these preventative policies are not already implemented.

The international organization with the most direct mandate to engage in preventive action is the Department of Political and Peacebuilding Affairs (DPPA). Its Mid-year report for 2022 states:

With US$9.4 million received from 11 donors, DPPA supported peace processes, mediation, and elections, throughout the world. From Ethiopia, Libya, and Myanmar to Syria and Yemen, we continued to advance political solutions in some of the most challenging crises, provided mediation between parties at odds, and encouraged preventive action for lasting peace.

This is a revealing statement for two reasons. First, the five countries mentioned spent the entirety of 2022 in one of stages 9, 10, 11, or 12. Syria and Myanmar spent the entire year in stage 11, Yemen spent 9 months in stage 11 and 3 months in stage 12, while Ethiopia spent 6 months in stage 11 and 4 months in stage 12. Second, the budget is extremely small for an organization that deals with conflict worldwide. For comparison, the UNHCR budget in 2022 was US$10.714 billion. The small DPPA budget seems to be completely absorbed by attempts to de-escalate the most violent ongoing conflicts.

The World Bank has a much larger budget and it is making fragile and conflict-affected situations (FCS) one of its main priorities. The Bank has an FCS country list and special funding and policy instruments in place for countries on the list. This suggests a big focus on prevention as *fragility* is a concept that is intimately connected to forecasting - the word is capturing an uncertain, risky future. However, to enter the FCS list, countries need to have a) a Country Policy and Institutional Assessment (CPIA) rating of 3.2 or less or b) the presence of a United Nations and/or regional peace-keeping or peace-building mission. This second aspect very obviously targets the aftermath of armed conflict - not its prevention. The first element has preventive elements but the data underlying the CPIA is released with significant delay and is heavily affected by armed conflict. The result is largely a list of countries with recent or ongoing conflicts and has very little predictive power beyond the, admittedly strong, persistence of conflict. This means that the FCS policy response by the World Bank and other actors using the list will be largely directed away from prevention and towards the trap stages of 9 to 12.

Most organizations working in this area are separated into development and conflict departments and the conflict specialists are only allocated to countries when a country lands on one of the lists, i.e. when there are substantial levels of violence. In addition, conflict-specific programs are withdrawn relatively quickly when violence stops. This is clearly visible in the data if we look at average responses across stages. We have conducted a systematic analysis of existing policy variables and everything points towards a late policy response in high stages 9 to 12. Sanctions and mediation for power-sharing agreements both clearly react to violence and therefore accumulate in the highest stages - so do peacekeeping missions and UN resolutions.

A summary of the cross-sectional variation of the allocation of different resources is presented in Figure F1. Power-sharing agreements are shown in panel A. We see that there are relatively few agreements below stage 10 and almost none below stage 8 - stages which the model identified as crucial. A similar pattern emerges in panel B for peace and security spending. Government and civil society ODA spending in panel C shows slightly more attention to lower stages but the dominance of stages with actual conflict persists. Only economic ODA in panel D paints a more nuanced picture. Economic aid, typically not explicitly associated with conflict prevention, is relatively strong in preventative stages. In Mueller et al. (2024b) we show the same pattern for IMF macro programs.

The documented allocations of resources reinforce the impression that there are little to no systematic conflict prevention policies in place at the global level. Generally, there is substantial engagement in countries at high risk but the engagement is focused on economic development and fiscal aspects.[9] The intuition could be that what is good for economic development is automatically reducing conflict. However, the literature documents some counter-examples, e.g. Dube and Vargas (2013), Berman et al. (2017), Sonno (2024).

In summary, our analysis suggests that there are attainable information gains.

[9] It should be mentioned that there is country heterogeneity by donor country here with the UK more heavily engaged with a different policy mix in preventative stages. Our conversations with the FCDO clearly indicate an awareness of the problems pointed out here.

Status quo policies are reacting to violence or not directly addressing conflict risk. This is despite a large literature in economics analyzing causal factors for risks and showing that there are interventions that can help address conflict risks.[10]

[10] For a recent review see Rohner (2024).

Figure F1: Allocation of resources across stages

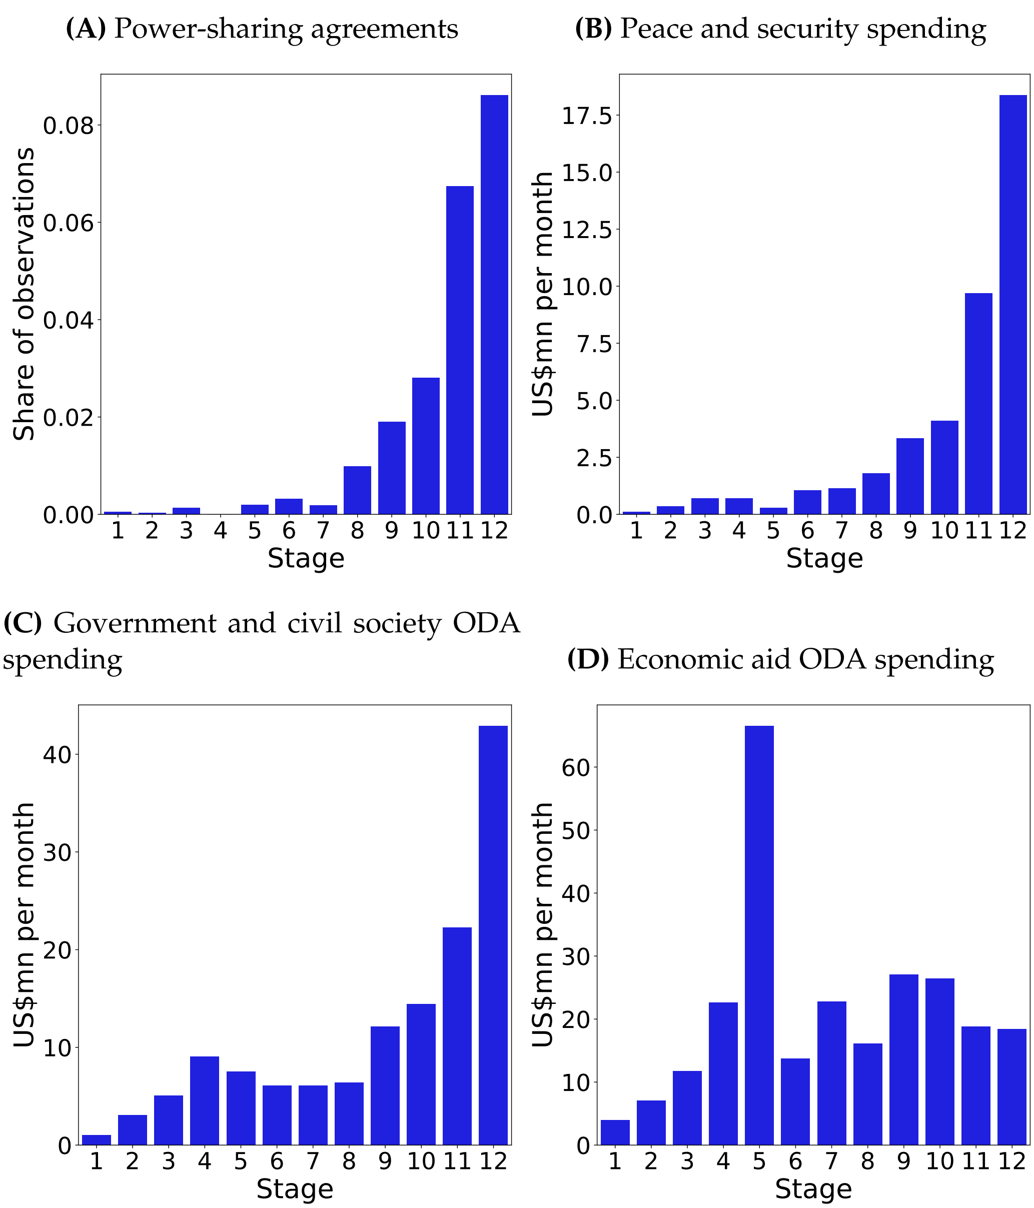

(A) Power-sharing agreements

(B) Peace and security spending

(C) Government and civil society ODA spending

(D) Economic aid ODA spending

NOTES: In Panel A an observation is defined as a country/month. The figure shows the share of observations, conditional on the stage, in which power-sharing agreements took place in the period 2010-2020. Panels B, C, and D show average monthly expenditure (US$mn) on specific ODA categories by stage, sourced from the OECD. Panel B shows the ODA spending on peace and security, panel C shows ODA spending on government and civil society, and panel D shows ODA spending on economic aid.