UAB
Universitat Autònoma de Barcelona

3 Essays in Disaster Risk, Conflict, and Tax Evasion

Joan Margalef España

3 Essays in
Disaster Risk, Conflict, and Tax Evasion

by Joan Margalef España

Supervisors: Pau Milán & Hannes Mueller

A dissertation submitted in fulfillment of the requirements for the degree of
Doctor of Philosophy in the International Doctorate in Economic Analysis,
Departament d'Economia i Història Econòmica

DOCTORAL THESIS
2025

UAB
Universitat Autònoma
de Barcelona

BSE
Barcelona School of Economics

A Marisol i Carlos.

Acknowledgements

This thesis is the result of many years of work and countless interactions. While I cannot mention everyone I am thankful to, I would like to acknowledge a group of people to whom I owe a special debt of gratitude.

First, I want to thank my family and friends. Though distant from the academic content of this work, they have always believed in me and supported me unconditionally. I thank my parents, Marisol and Carlos, and my siblings Maria, Ferran, and Carles for their love and encouragement.

Second, I am grateful to those within the IDEA program who have supported me professionally and personally. Thanks to Àngels and Mercè for their tireless efforts in keeping everything running. Thanks to Inés, Joan Maria Esteban, Luis, Massó, Caballer, Joan Llull, Caminal, Xavier, and Amedeo for their valuable guidance. Special thanks to David and Laura Mayoral for their time writing letters of recommendation.

Third, a warm thanks to my peers, who walked this journey with me, sharing feedback, frustrations, and resilience. To Pau's "minions", Hannes' "minions", and especially my cohort: Tino, Abhinav, Marta, Conghan, Alessandra, and Clemente.

Finally, and most importantly, I thank my two parents in research, Pau and Hannes. Pau took me under his wing during my Master's and helped me take my baby steps in research with enormous care and patience. I will never forget our long afternoons drawing equations on the blackboard or our eternal Zoom meetings during Covid. I deeply admire his passion for research and his ability to empower others.

Hannes first introduced me to topics that resonated deeply with me and later became both a co-author and a supervisor. Our conversations—whether over coffee in Plaça Catalunya or in airport calls, often in between laughs—shaped much of my thinking and professional path. I truly admire his ability to work from anywhere, connect with people, and bring them along with him. What he is building with EconAI is truly unique.

To both Pau and Hannes—whom I truly consider exceptionally talented, hardworking, and generous, and whom I look up to and feel honored to have had as my supervisors—thank you for believing in this work, and for believing in me, especially during the moments I was not able to.

Thanks.

Preface

This thesis examines three topics of academic and policy relevance—disaster risk, conflict, and tax evasion—using a combination of theoretical models and empirical techniques. It contributes to deepening our understanding of complex challenges affecting economic governance.

The first chapter explores the idea that investors' beliefs can be used to anticipate economic disasters. The main challenge is that these beliefs are not directly observable and must be inferred from asset prices. I develop a method to estimate investors' perceived probability of disasters—such as sovereign defaults or wars— from yield curve data, and assess their anticipatory value. By combining a theoretical asset pricing model with high-frequency yield curve data, I produce daily estimates of one-year-ahead disaster probabilities for over 60 countries from 2000 to 2023. These probabilities show signs of anticipation of major events, including debt restructurings and the Russian invasion of Ukraine, although the timing and intensity of responses vary across cases. They also show meaningful predictive power. The findings suggest that bond markets can provide a valuable high-frequency signal for assessing tail risks.

The second chapter, co-authored with Hannes Mueller, assesses the long-term economic cost of conflict. Armed conflict is a phenomenon that creates trap dynamics. Once a country falls into it, the likelihood of future outbreaks remains high. A key challenge in estimating its long-term impact is then accounting for the subsequent conflicts triggered by the first. We propose a model that treats conflict as a Markov process with states of conflict and peace. We estimate the transition matrix and the GDP per capita growth distribution associated with each state. We then perform simulations to estimate the distribution of costs over a long horizon. On average, countries that enter into conflict experience a 20% loss in GDP per capita after 30 years, with losses exceeding 40% in worst cases. These results highlight the economic value of effective prevention and peacebuilding strategies, which can yield substantial gains.

The third chapter examines why international efforts to curb tax evasion through tax havens have fallen short. I developed a coordination game, grounded in the global games literature, to analyze how investors respond to different enforcement strategies, and how these responses shape overall levels of evasion. The model shows that when one tax haven becomes relatively more attractive, evaders tend to concentrate there, facilitating offshore activity. As a result, policies that differentiate between havens can even backfire and increase evasion. This framework helps explain the limited effectiveness of past OECD-G20 initiatives and underscores the importance of harmonized approaches, such as the Global Minimum Tax.

Contents

List of Figures

List of Tables

Chapter 1

Disaster Risk through Investors' Eyes: a Yield Curve Analysis

1.1 Introduction

There are many reasons to care about investors' beliefs. A central one is that they can help anticipate economic outcomes, as investors have strong incentives to be accurate in their predictions.[1] This is the basis of the Efficient Market Hypothesis, which posits that asset prices reflect all available information and adjust quickly to new developments.[2] This is especially relevant in the context of economic disasters, such as defaults, interstate wars, or depressions, given their profound economic consequences. Thus, understanding and monitoring investors' perceived probability of disaster can be crucial for anticipating disasters and informing policy actions. The challenge is that investors' beliefs are not directly observable. However, since asset prices are directly influenced by these beliefs,[3] they can be used to reveal the probability of disaster as perceived by investors.

This paper provides a model to extract investors' perceived probability of disaster from yield curve data. The yield curve, which plots government bond yields against their maturities, consolidates investors' beliefs over different time horizons. By integrating an asset pricing model with yield curve data from Datastream, I provide daily estimates of the one-year-ahead disaster probability as perceived by investors for approximately 60 countries from 2000 to 2023. Then, I study the predictive value of these probabilities through several exercises.

To estimate the disaster probabilities, I use a classic asset pricing model, based on Rietz (1988) and Barro (2006), that incorporates time-varying disaster probabilities. A representative consumer maximizes expected consumption in a closed economy where she can invest in government bonds. The equilibrium conditions imply that prices depend on the expectation of consumption growth, inflation, and sovereign default. The occurrence of a disaster induces significant shifts in these variables, which I refer to as "jumps". Thus, the probability of such disasters shapes investors' expectations regarding these variables, and then prices. The nature of these "jumps" varies depending on the type of disaster being analyzed: sovereign default or interstate war. The model implies that observed bond prices can be decomposed into a theoretical non-disaster price and a *disaster wedge*. The non-disaster price reflects the price determined by current business cycle conditions, while the disaster wedge captures the component of the price driven by disaster-related elements. The theoretical model is calibrated to many countries, allowing us to compute the non-disaster theoretical prices for each of them over time. I bring the model to the data by regressing observed bond prices on the computed theoretical non-disaster prices using a fixed effects regression. By exploiting variation across countries, time periods, and maturities, this approach accounts for potential model misspecifications and isolates key fixed effects essential for estimating the disaster wedge. Finally, by specifying a type of disaster based on each country's context, I estimate the disaster probability.

To assess the predictive value of the estimated disaster probabilities, I first analyze how they evolve before disasters through case studies. This reveals how investors anticipate them and respond to new

1. Another reason is that investors' beliefs can have real effects. They may trigger self-fulfilling crises (Lorenzoni and Werning 2019; De Grauwe and Ji 2013), increase the debt burden, and ultimately hinder economic growth (Reinhart and Rogoff 2010). This is why central banks aim to shape expectations through effective communication and policy interventions (Blinder et al. 2008).

2. This notion originates from Fama (1970). See Malkiel (2003) for a review. Prediction markets exemplify this principle by aggregating beliefs through trading contracts on future events, such as elections or sports outcomes, often demonstrating remarkable predictive accuracy (Wolfers and Zitzewitz 2004; Arrow et al. 2008).

3. Ross (2015) refers to disaster risk as *dark matter*: "It is unseen and not directly observable but exerts a force that can change over time and profoundly influence markets."

information as the disaster approaches. Second, I examine their relationship with credit ratings data. Finally, I conduct a series of forecasting exercises using machine learning techniques to evaluate the value of the estimated disaster probabilities. Specifically, I compare the Area Under the Curve (AUC) of the Receiver Operating Characteristic (ROC) curve and the Precision-Recall (PR) curve across three models: one using only disaster probabilities, another using only credit rating variables, and a third combining both sources.

Probabilities rise before disaster events, such as the debt restructurings of Greece, Sri Lanka, and Ghana, as well as the onset of the Russia-Ukraine war. In some cases, such as Greece and Ghana, the probabilities reached 100% months before the event, while in others, smaller increases were observed weeks ahead, followed by sharp spikes shortly before or on the day of the event. This suggests that investors respond quickly to new information and adjust probabilities as the disaster approaches, but the timing and intensity of these responses might vary across events. Furthermore, disaster probabilities are strongly associated with higher-risk credit ratings. Finally, the estimated probabilities have meaningful predictive power for forecasting disasters (ROC-AUC 0.79, PR-AUC 0.24) and also improve the predictive performance of credit ratings, raising their ROC-AUC from 0.94 to 0.97 and their PR-AUC from 0.76 to 0.92. Overall, while investors' beliefs are not perfect, the findings show they provide useful information and respond quickly to rising risks.

This article relates to at least two strands of literature. The first is the macroeconomic literature on "(rare) disasters" or "tail events". The early disaster literature was theoretical, addressing asset pricing puzzles—such as the risk-free rate premium—by introducing the concept of a low-probability of a "consumption" disaster (Rietz 1988; Barro 2006; Gabaix 2008; Backus et al. 2011; Gourio 2012; Gabaix 2012; Wachter 2013; Farhi and Gabaix 2016).[4] A consumption disaster unifies extreme events like wars and depressions into a single concept, representing any significant decline in consumption observed in U.S. history. More recently, research has focused on empirically estimating the probability of these consumption disasters. This body of research is largely based on reduced-form models (Berkman et al. 2011; Schreindorfer 2020). A notable utility-based model in Barro and Liao (2021) uses option prices and fixed-effect regressions to estimate consumption disaster probabilities across major economies.[5] This paper contributes both methodologically and through its practical applications. First, it introduces a structural model to estimate disaster probabilities from high-frequency yield curve data. The yield curve data I use is particularly insightful due to its panel data structure, including a maturity dimension, and its well-established role as a key financial indicator.[6] The model is flexible enough to be calibrated for a wide range of countries, offering daily updates, and it can account for different types of disasters, not just consumption disasters. This is important because consumption disasters are rare and may not be the most salient risk. Second, the paper assesses the predictive value of the computed disas-

4. Julliard and Ghosh (2012) argue that rare events alone cannot adequately explain asset pricing puzzles like the equity premium.

5. A related contribution is Ross (2015), which introduces the Recovery Theorem—a method to disentangle investors' beliefs about future outcomes from their risk preferences, allowing one to infer the perceived probability of catastrophic events.

6. Substantial empirical evidence suggests that the yield curve is one of the most informative indicators, particularly for forecasting economic downturns (Estrella and Hardouvelis 1991; Estrella and Mishkin 1998; Ang et al. 2006). Even the Federal Reserve Bank of New York has a webpage dedicated to the yield curve and its predictive power for recessions. See https://www.newyorkfed.org/research/capital_markets/ycfaq.html and https://www.newyorkfed.org/medialibrary/media/research/capital_markets/Prob_Rec.pdf. Furthermore, other studies show that the yield curve responds to economic policy uncertainty (Leippold and Matthys 2022), political uncertainty (Pástor and Veronesi 2013; Smales 2016) and international political risk (Huang et al. 2015).

ter probabilities, paving the way for their use in future applications. The computed disaster probabilities will be available on my GitHub repository[7], enabling researchers and policymakers to incorporate them into their own analyses.

Secondly, financial literature has extensively examined the predictive power of asset prices, including the government yield curve, in forecasting economic outcomes. A widely used measure is the sovereign yield spread, which compares yields from "safe" countries like Germany or the U.S. to those from riskier assets, such as foreign bonds.[8] Yields and their spreads have been frequently applied in reduced-form models, capturing sovereign default and political risks (Clark 1997; Remolona et al. 2007; Bekaert et al. 2016), financial crises (Bluwstein et al. 2023), and wars (Chadefaux 2017). This paper adds to this body of research by leveraging the yield curve in a structural model. Computing the theoretical non-disaster price that accounts for current business cycle conditions, I "control" for factors that also influence yields in a theoretically consistent way. Furthermore, fixed effects allow for a further decomposition of yields, isolating the variation driven by disaster risk. Finally, it's important to differentiate between the predictive power of asset prices and whether the beliefs inferred from these prices are ultimately accurate.

This paper is structured as follows. The next section outlines the asset pricing model. Section 1.3 describes the data sources. In Section 1.4, I present the methodology for estimating investors' perceived probability of disaster. In Section 1.5, I discuss the results, followed by the conclusions in the final section.

1.2 Model setup

The model follows Rietz (1988) and Barro (2006), which I extend by including time-varying probabilities of disasters. It will later be calibrated separately for different countries, but for clarity and simplicity, the country-specific indices are omitted in this section.

The representative consumer maximizes a time-additive utility function:

$$\mathbb{E}_t \sum_{t=0}^{\infty} \beta^t U\left(C_t\right), \tag{1.1}$$

where β is the time discount factor, and the utility function takes the CRRA form

$$U\left(C_t\right) = \frac{C_t^{1-\theta}}{1-\theta} \tag{1.2}$$

θ is the coefficient of relative risk aversion. In each period, agents can invest in government nominal zero-coupon bonds, each of which will pay out one unit of currency at maturity. Q_{Nt} is the price at t of a bond that matures in N periods, and X_{Nt} is the amount bought. The government can default on its obligations and pay a fraction F_{Nt} of the bond's face value. F_{Nt} represents the recovery rate. $F_{Nt} = 1$ indicates full payment with no default, and $F_{Nt} = 0.8$ implies that the government pays 80% of the

7. https://github.com/joanmargalef

8. The spread between corporate and government bonds is also commonly used as an indicator of economic activity (Gilchrist and Zakrajšek 2012; Gilchrist et al. 2016), and corporate default risk (Duffee 1999; Dionne et al. 2010). Another method for corporate default risk is the Merton Distance to Default model (Merton 1974; Bharath and Shumway 2008).

bond's face value. The per-period budget constraint of the agents is given by

$$P_t C_t = W_t - \sum_{N=1}^{H} Q_{Nt} X_{Nt} \tag{1.3}$$

where P_t is the price of consumption and W_t corresponds to the wealth if no bond is bought, which includes the payments from previously purchased bonds.[9] H represents the maximum maturity. Using the usual first-order conditions, I derive the fundamental asset pricing equation:

$$Q_{Nt} = \beta^N \mathbb{E}_t \left[\frac{U'(C_{t+N}) P_t}{U'(C_t) P_{t+N}} \right] \tag{1.4}$$

The relationship between bond prices and bond yields is given by

$$Y_{Nt} = \left(\frac{1}{Q_{Nt}} \right)^{\frac{1}{N}} - 1 \tag{1.5}$$

where Y_{Nt} is the yield of a bond that matures in N periods at time t. The yield curve is the graph that plots Y_{Nt} against N. This equation allows us to translate bond prices to yields and vice versa.

Substituting in the functional form of the marginal utilities of consumption, Equation 1.4 can be rewritten as

$$Q_{Nt} = \beta^N \mathbb{E}_t \left[\frac{F_{Nt}}{\prod_{j=1}^{N} G_{t+j}^{\theta} \Pi_{t+j}} \right] \tag{1.6}$$

with $G_{t+1} = C_{t+1}/C_t$ being consumption growth and $\Pi_{t+1} = P_{t+1}/P_t$ being inflation. Note that bond prices decrease in expected consumption growth and inflation. Since the bond is a mechanism to transfer consumption to the future, there are fewer incentives to buy the bond if consumption is expected to be high. Higher expected inflation diminishes the real value of the bond. The price also decreases as the expected recovery rate decreases.

Following the standard approach in asset pricing, I will analyze this equilibrium price equation using exogenous processes for consumption growth, inflation, and the recovery rate.[10]

In each period, a disaster may or may not occur. For simplicity, disasters are assumed to be independent of one another.[11] $\delta_{\tau,t}$ denotes the probability at t of a disaster happening in τ periods s.t.

$$\delta_{\tau+1,t} = \phi_\delta \delta_{\tau,t} \tag{1.7}$$

with $\phi_\delta \in [0, 1]$ being the persistence parameter of the disaster probability. This allows us to express all disaster probabilities in terms of $\delta_{1,t}$ since $\delta_{\tau,t} = \phi_\delta^{\tau-1} \delta_{1,t}$.

The law of motion of consumption growth is

$$G_{t+1} = \alpha_G G_t^{\phi_G} \varepsilon_{t+1} V_{t+1} \tag{1.8}$$

where α_G is a constant term, ϕ_G represents a persistence parameter, $\varepsilon_{t+1} \overset{\text{iid}}{\sim} \log \mathcal{N}(0, \sigma_\varepsilon^2)$ is white noise,

9. The model shows a closed economy, where all that is produced is consumed. The BIS report Fang et al. (2022) shows that the majority of government bonds are held by domestic investors, especially during crises.

10. See Cochrane (2009).

11. While independence is assumed for tractability, the realization of disasters may exhibit persistence or avoidance effects. Wars may trigger war traps, while major depressions are less likely to recur immediately.

and $V_{t+\tau}$ is the "disaster impact factor on consumption growth" s.t.

$$V_{t+\tau} = \begin{cases} 1 & \text{if no disaster at } t+\tau \\ J_G & \text{if disaster at } t+\tau \end{cases}$$

Therefore, the disaster affects consumption growth through V_{t+1}. When the disaster does not occur, the log of consumption growth follows an AR(1) process. $J_G > 0$ represents the "jump" in consumption growth induced by the disaster. A value of $J_G = 0.98$ implies that the disaster reduces consumption growth by 2%. Note that the disaster directly impacts consumption growth in the same period it occurs and indirectly in future periods. If the disaster occurs at $t+1$, it will directly impact G_{t+1} through V_{t+1}. Additionally, it will indirectly affect $G_{t+2}, G_{t+3}, \ldots$ through their dependence on G_{t+1}.

Analogously, the process of inflation is

$$\Pi_{t+1} = \alpha_\Pi \Pi_t^{\phi_\Pi} \eta_{t+1} W_{t+1} \tag{1.9}$$

where α_Π is a constant term, ϕ_Π is the persistence parameter, $\eta_{t+1} \overset{\text{iid}}{\sim} \log\mathcal{N}(0, \sigma_\eta^2)$ is white noise, and $W_{t+\tau}$ is the "disaster impact factor on inflation" s.t.

$$W_{t+\tau} = \begin{cases} 1 & \text{if no disaster at } t+\tau \\ J_\Pi & \text{if disaster at } t+\tau \end{cases}$$

$J_\Pi > 0$ represents the "jump" in inflation induced by the disaster. A $J_\Pi = 1.05$ means the disaster increases inflation by 5%. As with consumption growth, the disaster directly impacts inflation in the same period it occurs and indirectly in future periods.

When a disaster occurs, there is a probability γ that it will lead to a sovereign default, which I model as an equal haircut across all bonds. When there is no disaster, the probability of default is zero. Then, the recovery rate is given by

$$F_{Nt} = 1 \cdot \prod_{\tau=1}^{N} Z_{t+\tau} \tag{1.10}$$

with $Z_{t+\tau}$ being the "disaster impact factor on the recovery rate" s.t.

$$Z_{t+\tau} = \begin{cases} 1 & \text{if no disaster at } t+\tau \\ 1 & \text{if disaster but no partial default at } t+\tau \\ 1 - J_F & \text{if disaster and partial default at } t+\tau \end{cases}$$

$J_F \in [0, 1]$ denotes the size of the haircut. A $J_F = 0.2$ means that the government does not pay 20% of the face value of the bond. A $J_F = 1$ is full default. The product of $Z_{t+\tau}$ over all periods until maturity implies that haircuts are cumulative, making long-term bonds riskier since they can suffer several haircuts.

Note that independence between disasters implies that the disaster impact factors are independent across periods, i.e., $V_t \perp V_{t'}, W_{t'}, Z_{t'}$ for $t' \neq t$. However, V_t, W_t, and Z_t are perfectly correlated through the disaster event.

Given this, the bond price from Equation 1.6 can be expressed as

$$Q_{Nt} = \underbrace{Q_{Nt}^{ND}}_{\text{Non-disaster price}} \underbrace{\prod_{\tau=1}^{N} \left(1 + \phi_{\delta}^{\tau-1} \delta_{1,t} \left(J_{\tau,N} - 1\right)\right)}_{\text{Disaster wedge}} \tag{1.11}$$

Q_{Nt}^{ND} represents the bond price in the absence of disasters, and $J_{\tau,N}$ synthesizes all the jump effects of a disaster happening in τ periods to a bond that matures in N periods. I refer to $J_{\tau,N}$ as the "overall jump". Remember that $\delta_{1,t}$ is the probability at t of a disaster happening in 1 period. Thus, the price of the bond consists of the non-disaster price, Q_{Nt}^{ND}, multiplied by a "disaster wedge" that accounts for the risks of all disasters that may occur before the bond reaches maturity. This wedge depends on the disaster probabilities for all periods before maturity, $\delta_{\tau,t} = \phi_{\delta}^{\tau-1} \delta_{1,t}$ for $\tau \in [1, N]$, and the potential impact of each, summarized in $J_{\tau,N}$. For example, a 2-period bond is affected by the risk of a disaster happening in 1 and 2 periods, but not after, as it will have already matured. The term $1 + \phi_{\delta}^{\tau-1} \delta_{1,t} \left(J_{\tau,N} - 1\right)$ is the specific disaster wedge induced by the disaster in τ periods. Long-term bonds accumulate more elements in the product, as they are exposed to more periods where disasters can occur.

The functional form of the non-disaster price is

$$Q_{Nt}^{ND} = \beta^{N} \frac{e^{\frac{1}{2}\left(\sum_{i=1}^{N} (\sum_{j=0}^{i-1} \phi_G^j)^2 \theta^2 \sigma_{\varepsilon}^2 + \sum_{i=1}^{N} (\sum_{j=0}^{i-1} \phi_{\Pi}^j)^2 \sigma_{\eta}^2\right)}}{\left(\alpha_G^{\sum_{i=1}^{N} i \phi_G^{N-i}} G_t^{\sum_{i=1}^{N} \phi_G^i}\right)^{\theta} \alpha_{\Pi}^{\sum_{i=1}^{N} i \phi_{\Pi}^{N-i}} \Pi_t^{\sum_{i=1}^{N} \phi_{\Pi}^i}} \tag{1.12}$$

This incorporates the expectations based on current business cycle conditions since it contains the effect of current consumption growth (G_t) and inflation (Π_t).

Finally, the functional form of $J_{\tau,N}$ is

$$J_{\tau,N} = \frac{1 - \gamma J_F}{J_G^{\sum_{j=1}^{N+1-\tau} \theta \phi_G^{j-1}} J_{\Pi}^{\sum_{j=1}^{N+1-\tau} \phi_{\Pi}^{j-1}}} \tag{1.13}$$

This illustrates that the effect of a disaster depends on the interplay between jump effects in consumption growth (J_G), inflation (J_{Π}), and default risk (γ and J_F), which may offset each other. As the gap between the disaster's occurrence (τ) and bond maturity (N) increases, the summations in the exponents include more components. This reflects that long-term bonds have more indirect effects by a single disaster due to the persistence of the underlying variables' processes (ϕ_G and ϕ_{Π}). As a result, short- and long-term bonds may behave very differently, even in opposite directions. For example, if a disaster causes a sharp drop in consumption growth and moderate inflation, but consumption growth is less persistent, the recessionary impact will be strong initially, leading to an increase in short-term bond prices. However, as the effect fades quickly and inflation persists, long-term bond prices will eventually decrease as inflation outweighs the recessionary impact.

This model offers tractable solutions for decomposing bond prices and allows us to analyze how disaster probabilities affect them.

Proposition 1.1 *The bond price with maturity N at time t, Q_{Nt}, decreases with the probability of a*

disaster occurring in the next period, $\delta_{1,t}$, if and only if

$$\sum_{\tau=1}^{N} \frac{\phi_{\delta}^{\tau-1}\left(J_{\tau,N}-1\right)}{1+\phi_{\delta}^{\tau-1}\left(J_{\tau,N}-1\right)} < 0 \tag{1.14}$$

A sufficient condition for this to hold is that $J_{\tau,N} < 1$ for all τ.

The proof is in Appendix A.1. When the probability of a disaster in the next period ($\delta_{1,t}$) increases, the probability of disasters in all future periods before maturity also increases, since $\delta_{\tau,t} = \phi_{\delta}^{\tau-1}\delta_{1,t}$ with $\phi_{\delta} \in [0,1]$. The overall effect on the bond price is ambiguous because the impact of a disaster in different periods ($J_{\tau,N}$) can have opposing effects. A sufficient condition for the bond price to fall as disaster risk rises is that the impact of a disaster is negative in every period before maturity, i.e., $J_{\tau,N} < 1$ for all τ. In that case, an increase in $\delta_{1,t}$ will unambiguously reduce Q_{Nt}.

1.3 Data

I use yield curve and credit ratings data from Datastream, macroeconomic data from the International Monetary Fund's International Financial Statistics (IMF/IFS) and the World Bank's World Development Indicators (WB/WDI), and conflict data from the Uppsala Conflict Data Program's Georeferenced Event Dataset (UCDP/GED).

The yield curve data is used to derive the bond prices (Q_{Nt}). IMF/IFS data is used to calibrate current inflation (Π_t) and consumption growth (G_t), which are input into the non-disaster price formula. WB data is used to estimate the parameters governing the laws of motion for consumption growth and inflation (α_G, ϕ_G, σ_ε and α_Π, ϕ_Π, σ_η). This is later combined with UCDP data to estimate the effects of conflict, particularly the size of the jumps associated with disaster events (J_G and J_Π). The credit ratings from Datastream are used for comparison in the forecasting exercises.

1.3.1 Yield curve data

Refinitiv's Datastream provides daily government bond yields for a wide range of countries, including both developed and developing economies. The availability of bond data varies by country; more developed countries typically offer a greater variety of bonds and longer maturity horizons. The analysis includes 64 countries over various time horizons.[12]

I retrieved the daily "benchmark" yield curve, which is based on "benchmark" bonds.[13] These are the most liquid government bonds, which are particularly relevant for analyzing investor expectations, as they capture actively traded securities that swiftly respond to market developments.[14] These cover standard government bonds with fixed rates and fixed maturity dates while excluding bonds with variable

12. For a detailed list of all countries, including their respective time spans and maturity coverage, see Table A.1 in Appendix A.3.

13. These are based on Refinitiv Government Bond Indices, which are calculated using methodologies recommended by the European Federation of Financial Analysts Societies (EFFAS).

14. The Refinitiv Government Bond Indices include three main types: All Traded Index, which includes all eligible bonds, providing comprehensive market coverage; Tracker Index, a sample of bonds that closely tracks overall market performance; and Benchmark Index, focusing on the most liquid bonds.

rates and other features that distort predictability.[15] All the bonds are denominated in the local currency of the issuing country. I use the yield curve data provided directly by Refinitiv without any time lags.[16]

Finally, I restrict the sample to bonds with maturities between 1 and 10 years for two main reasons. First, this range aligns with the year-over-year growth rates of the macroeconomic variables. Second, these maturities are more frequently available in the dataset, ensuring adequate data coverage and consistency in the analysis.

1.3.2 Economic and conflict data

The economic variables of interest are consumption growth and inflation, which were obtained at a quarterly frequency from the IMF/IFS and an annual frequency from the WB/WDI. In both datasets, consumption growth is proxied by GDP growth in constant local currency units. Inflation is measured using the Consumer Price Index (CPI).

The IMF data, which provides quarterly updates, allows me to run the model at a quarterly frequency by inputting per-period consumption growth (G_t) and inflation (Π_t). The annual WB data offers a longer time span, which is especially useful for estimating the parameters of the laws of motion for consumption growth and inflation, as well as the disaster parameters.

Finally, to link economic effects to interstate wars, I utilize battle-related fatality data from the UCDP/GED. I aggregate this data to the country-year level.

1.3.3 Credit ratings data

Refinitiv's Datastream includes ratings from multiple credit rating agencies, such as Moody's, Fitch, Dominion Bond Rating Service (DBRS), and Rating & Investment (R&I). Additionally, it provides an equivalence mapping to the Standard & Poor's (S&P) rating scale, which consists of over 20 categories (e.g., AAA, AA+, AA, AA-, etc.).

To ensure consistency and comparability, I standardize all relevant ratings to their equivalent S&P categories. This transformation creates a unified scale for evaluating and aggregating ratings from different agencies. For each country and time period, the average rating is calculated as the mean of the S&P-equivalent ratings across all agencies.

1.4 Estimating investors' perceived probability of disaster

To estimate disaster probabilities, I first calibrate the model for every country to compute theoretical non-disaster prices. Then, I bring the model to the data by running a fixed effects regression to attribute part of the difference between the observed prices and the computed theoretical ones to the disaster wedge. Finally, based on a specific disaster type, I estimate the probability of the disaster for each country and day.

15. Excluded bonds include those with inflation-linked, floating rate, convertible, and bonds with embedded options or warrants.

16. Refinitiv also offers computed yield curves for third parties, which may have pricing lags.

1.4.1 Calibration

Calibrating the model for all countries requires setting parameters for the utility function, the laws of motion for consumption growth and inflation, and disaster-related parameters.

1.4.1.1 Utility function and laws of motion

I derive the utility function parameters from established literature. Following the methodology posited by Barro (2006), I set the discount factor, β, to 0.97 per year, and the coefficient of relative risk aversion, θ, to 4, which are common to all the countries.

Table 1.1: Summary of calibration: utility function and laws of motion

Variable	Value	Source
Time preference (β)	0.97	Barro (2006)
Risk aversion (θ)	4	Barro (2006)
Consumption growth (G_{ct})	Country-specific	IMF/IFS
Inflation (Π_{ct})	Country-specific	IMF/IFS
Constant of consumption growth ($\alpha_{G,c}$)	Country-specific	Estimated from WB/WDI
Constant of inflation ($\alpha_{\Pi,c}$)	Country-specific	Estimated from WB/WDI
Persistence of consumption growth ($\phi_{G,c}$)	Country-specific	Estimated from WB/WDI
Persistence of inflation ($\phi_{\Pi,c}$)	Country-specific	Estimated from WB/WDI
S.d. of consumption growth ($\sigma_{\varepsilon,c}$)	Country-specific	Estimated from WB/WDI
S.d. of inflation ($\sigma_{\eta,c}$)	Country-specific	Estimated from WB/WDI

Notes: The table presents the calibration of the utility function and the laws of motion for consumption growth and inflation. The utility parameters (β, θ) are taken from the rare disasters literature. Consumption growth (G_{ct}) and inflation (Π_{ct}) are input from the IMF/IFS database, measured as quarterly GDP growth in constant local currency units and quarterly changes in CPI, respectively. The constants ($\alpha_{G,c}$, $\alpha_{\Pi,c}$), persistence parameters ($\phi_{G,c}$, $\phi_{\Pi,c}$), and standard deviations ($\sigma_{\varepsilon,c}$, $\sigma_{\eta,c}$) are estimated using OLS regressions on annual country-level data from WB/WDI.
Sources: Barro (2006), IMF/IFS, WB/WDI, and author's calculations.

The laws of motion for consumption growth and inflation are represented by Equation 1.8 and 1.9. Taking logs transforms the laws of motion into a linear form, which, in the absence of disaster shocks, follows an AR(1) process. For each country c, I estimate the constant parameters ($\alpha_{G,c}$ and $\alpha_{\Pi,c}$), the persistence parameters ($\phi_{G,c}$ and $\phi_{\Pi,c}$), and the standard deviations ($\sigma_{\varepsilon,c}$ and $\sigma_{\eta,c}$) using OLS on WB/WDI time series from 1989 to 2023. The distribution of the estimated parameters shows that the constant parameters are around 1.02 for both variables. Both log consumption growth and log inflation exhibit mean reversion. The inflationary process is more persistent, with an average persistence of 0.5, compared to 0.2 for consumption growth. See Figure A.1 in Appendix A.2 for kernel density plots of these estimates. For the period-specific consumption growth and inflation, G_{ct} and Π_{ct}, I use quarterly year-over-year data from IMF/IFS.

With all these parameters, I can compute the theoretical non-disaster prices, $\hat{Q}^{ND}_{Nct}$, for each maturity N, country c, and period t. As only G_{ct} and Π_{ct} vary over time, the non-disaster price is updated quarterly. See Table 1.1 for a summary of the calibration of the utility function and the laws of motion.

1.4.1.2 Disaster parameters

The disaster parameters to be calibrated include the jumps in consumption growth (J_G) and inflation (J_Π), the probability of default during a disaster (γ), and the haircut size (J_F). These parameters are calibrated for each type of disaster and are the same for all countries. I define two types of disasters: interstate war and sovereign default.

For interstate war, I conducted a two-way fixed effects analysis using WDI/WB data.[17] The results show that a year in war reduces consumption growth by 2% and increases inflation by 2%. Therefore, I set $J_G = 0.98$ and $J_\Pi = 1.02$. The regression results are presented in Table A.2 in Appendix A.3. I set J_F, to 0.56, based on the haircut analysis from Luckner et al. (2023), which uses historical data on sovereign defaults triggered by geopolitical disasters. Given that they recorded 45 defaults resulting from 95 interstate wars, I set the probability of default γ to 0.5.

For sovereign default, the objective is to capture the probability of default using $\delta_{1,t}$. Thus, the conditional probability of default given a disaster, γ, becomes redundant, so it is set to 1. The parameter J_F is set to 0.44, corresponding to the average sovereign debt haircut reported in Meyer et al. (2022). The jumps in consumption growth and inflation are set to 1.

Table 1.2: Summary of calibration: disaster parameters by disaster type

Disaster type	J_G	J_Π	γ	J_F	Source
Interstate war	0.98	1.02	0.5	0.56	Von Laer & Bartels (2023), author's calculations
Sovereign default	1	1	1	0.44	Meyer et al. (2022)

Notes: The table presents the calibration of disaster parameters by disaster type. J_G and J_Π are the jumps in consumption growth and inflation, respectively. γ is the probability of default when a disaster occurs, and J_F is the haircut size. $\phi_\delta = 0.5$ for all disaster types.
Source: Barro (2006), Luckner et al. (2023), and Meyer et al. (2022) and author's calculations on WB/WDI and UCDP/GED data.

Based on the calibrated disaster parameters and laws of motion, I can compute the overall jump, $\hat{J}_{\tau,cN}$, using Equation 1.13. The persistence parameter of the disaster probability, ϕ_δ, is set to 0.5 for all disaster types. Table 1.2 summarizes the calibration for each disaster.

Figure 1.1 shows the simulated impact of varying $\delta_{1,ct}$ values for each disaster type on the price curve, based on the U.S. calibration. The likelihood of either interstate war or default leads to a price drop across all maturities. For interstate war, the inflationary and default risks outweigh the recessionary effects, causing a decline in prices. Default has a stronger negative impact.

17. To match the conflict size with Luckner et al. (2023), I define war as having more than 1,000 deaths per year using UCDP data.

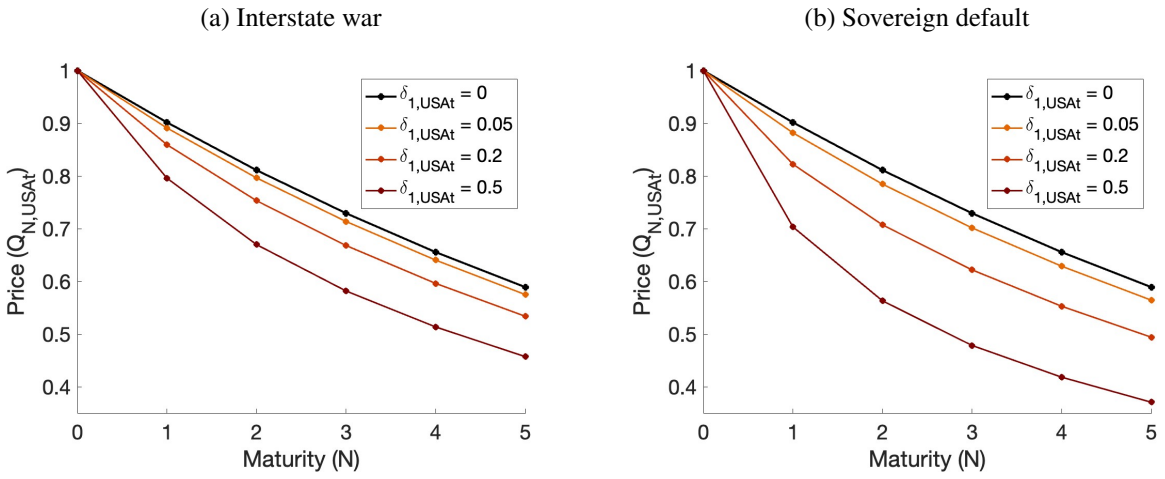

Figure 1.1: Comparative statics of disaster probabilities

Notes: The figure presents simulated bond price curves for each disaster type—interstate war (left) and sovereign default (right)—using the U.S. calibration. The simulations vary the one-period-ahead disaster probability, $\delta_{1,ct}$, while holding all other parameters constant. The persistence of disaster risk is set to $\phi_\delta = 0.5$, and all jump parameters are taken from Table 1.2. The U.S. calibration is based on the estimation of the laws of motion for GDP growth and inflation using WB/WDI data. For GDP growth in constant local currency, the constant is $\alpha_G = 1.02$, the persistence $\phi_G = -0.05$, and the residual standard deviation $\sigma_\varepsilon = 0.01$. For inflation, the constant is $\alpha_\Pi = 1.01$, the persistence $\phi_\Pi = 0.517$, and the residual standard deviation $\sigma_\eta = 0.01$. Prices are computed using Equation (1.11) with $G_t = \Pi_t = 1.02$.
Source: Author's calculations.

1.4.2 Bringing the model to the data

Building on the calibration, I bring the theoretical model to the data to estimate investors' perceived probability of disaster. Incorporating the panel structure using the country index c and taking logs, Equation 1.11 transforms into

$$q_{Nct} = q_{Nct}^{ND} + \sum_{\tau=1}^{N} \log\left(1 + \phi_\delta^{\tau-1} \delta_{1,ct} \left(J_{\tau,cN} - 1\right)\right) \tag{1.15}$$

with $q_{Nct} = \log(Q_{Nct})$ and $q_{Nct}^{ND} = \log(Q_{Nct}^{ND})$. The difference between the observed log price and the theoretical log price captures the disaster wedge. I bring this equation to the data by employing a fixed effects regression specified as:

$$q_{Nct} = \beta \hat{q}_{Nct}^{ND} + \kappa_{Nc} + \kappa_{Nt} + \kappa_{ct} + u_{Nct} \tag{1.16}$$

where κ_{Nc}, κ_{Nt} and κ_{ct} represent fixed effects for country-maturity, maturity-time, and country-time interactions, respectively, and u_{Nct} is the error term. q_{Nct} is the observed log price of a bond sourced from Datastream. $\hat{q}_{Nct}^{ND}$ is the computed non-disaster theoretical log price derived from the model's calibration. While observed bond prices are available daily, the computed theoretical prices are updated quarterly based on economic data. To align frequencies, I interpolate the quarterly theoretical prices to a daily level. The results are robust when using the quarterly model, which uses quarterly averages of

observed prices.[18]

This equation suggests that observed bond prices can be explained by the non-disaster theoretical price, which reflects expectations based on the current business cycle, plus a set of unobserved factors varying at different levels. Fixed effects regression offers several advantages. First, it corrects for potential model misspecifications. Comparing Equation 1.15 with the regression equation, if the model perfectly captures the bond price data-generating process, $\hat{\beta}$ would approximate 1. However, allowing it to deviate provides a more accurate reflection of the relationship and serves as a measure of model fit. Second, because $\delta_{1,ct}$ varies at the ct level, the model's country-time interaction term, κ_{ct}, isolates variations in country-specific factors over time, which is essential for estimating disaster probabilities. Finally, the additional fixed effects address structural and temporal influences, which enhance identification. The country-maturity interaction term, κ_{Nc}, captures structural yield curve differences across countries, reflecting time-stable variations potentially due to regulatory or market-specific conditions. The maturity-time interaction term, κ_{Nt}, controls for maturity-specific factors impacting all countries in a given period, such as global shifts in demand for certain maturities or adjustments in term premiums.

Table 1.3: Fixed effect regression

	Observed price (q_{Nct})				
	(1)	(2)	(3)	(4)	(5)
Non-disaster price ($\hat{q}^{ND}_{Nct}$)	0.109***	0.206***	0.218***	0.312***	0.292***
	(0.001)	(0.001)	(0.0004)	(0.0002)	(0.0002)
Country-time FE	✓	✓	✓	✓	
Maturity-country FE	✓	✓			
Maturity-time FE	✓		✓		
Observations	1,765,539	1,765,539	1,765,539	1,765,539	1,766,001
Adjusted R^2	0.973	0.946	0.850	0.831	0.350

Note: The table presents a fixed effects regression of the observed log bond price ($q_{Nct} = \log(Q_{Nct})$) on the log of the theoretical non-disaster price ($q^{ND}_{Nct} = \log(Q^{ND}_{Nct})$), estimated using the methodology described above. Observed prices are derived from Refinitiv Datastream yield curve data, based on benchmark government bond yields. The sample includes annual bond prices with maturities from 1 to 10 years, across 64 countries from 2000 to 2023. Fixed effects are κ_{Nc} (Maturity-country), κ_{Nt} (Maturity-time), and κ_{ct} (Country-time). Models differ by their inclusion of these fixed effects. Robust standard errors are reported in parentheses, with $^{*}p < 0.1$, $^{**}p < 0.05$, and $^{***}p < 0.01$ indicating significance levels.
Source: Datastream data for observed prices, and theoretical prices are calculated based on WB/WDI and IMF/IFS data.

Table 1.3 presents the regression results from different specifications, which vary in the fixed effects included. The favorite specification, from which I estimate the disaster probabilities, includes all fixed effects. Across all specifications, $\hat{\beta}$ is positive and significant, showing that the theoretical price moves in the same direction as observed prices. However, the values are below 1 in all specifications and generally decrease as more fixed effects are added—from 0.316 (only κ_{ct}) to 0.108 (with all fixed effects). The low $\hat{\beta}$ suggests that the non-disaster prices deviate from the actual observed prices—the model tends to underpredict prices.[19] The R^2 in the baseline model is 0.35 and rises to nearly 1 as all fixed effects are added. This pattern indicates that the model captures a moderate share of the variation, while the fixed

18. Figure A.2 in Appendix A.2 compares the daily and quarterly models for estimating $\hat{\kappa}_{ct}$.

19. This finding reflects the classic asset pricing puzzle, where standard models fail to match the low observed returns on risk-free government bonds. This puzzle motivated the inclusion of rare disasters as a potential explanation.

effects capture most of the remaining unexplained variation. Almost identical results are found when using the quarterly frequency model, see Table A.3 in Appendix A.3.

Given that $\delta_{1,ct}$ varies at the country-time level, it is captured in $\hat{\kappa}_{ct}$ as the common unobserved factor at the ct level. Its interpretation is as follows: if $\hat{\kappa}_{ct}$ is significantly positive, it indicates that there is an unobserved factor at the country-time level causing bond prices to be higher than what the current business cycle, and the other factors controlled by the other fixed effects, would suggest. Conversely, a significantly negative $\hat{\kappa}_{ct}$ implies that this unobserved factor is reducing bond prices. Since disaster risk reduces bond prices across maturities, as shown in Figure 1.1, a negative $\hat{\kappa}_{ct}$ may indicate that the unobserved factor is disaster risk.

Figure 1.2 illustrates the evolution of $\hat{\kappa}_{ct}$ for six representative countries selected for their distinct levels of disaster risk. Germany and the United States represent stable countries with no significant disaster risk. Ireland and Spain exemplify relatively stable countries that experienced periods of disaster risk (default risk during the European debt crisis) that ultimately did not materialize. In contrast, Greece and Ukraine represent countries where disaster risk materialized, including Greece's default in 2012 and Ukraine's interstate conflict in 2023. The horizontal lines at 0 and -0.5 serve as reference thresholds to facilitate comparison of $\hat{\kappa}_{ct}$ values across countries. For the evolution of $\hat{\kappa}_{ct}$ across the full set of countries, see Figures A.3, A.4, A.5, and A.6 in Appendix A.2.

Stable countries with no significant disaster risk, such as Germany and the USA, generally exhibit $\hat{\kappa}_{ct}$ values that remain positive or only briefly touch the zero baseline. This pattern extends to other economically robust regions, including Australia, the United Kingdom, Scandinavian countries (e.g., Sweden, Norway, Denmark), and stable European economies like Austria, the Netherlands, and Switzerland. Likewise, leading Asian economies such as Japan, South Korea, and Singapore also display this pattern.

Relatively stable countries that experienced periods of known disaster risk, such as Ireland and Spain, typically display positive $\hat{\kappa}_{ct}$ values with only shallow dips into negative territory during times of financial instability. Ireland's lowest $\hat{\kappa}_{ct}$ point occurs in mid-2011, aligning with austerity measures and bailout negotiations during the peak of the European debt crisis. Spain similarly reaches a minimum in mid-2012, reflecting peak financial strain in this period. Comparable patterns are seen in other southern European countries, including Portugal, Italy, and Cyprus. Other notable examples of brief negative dips without crossing the -0.5 threshold include Israel during the Second Intifada, amid severe conflict and economic disruption, and Poland in 2002, during economic adjustments following rapid liberalization and structural reforms, which led to rising unemployment, social discontent, and fiscal strain. Additionally, countries like Mexico and India experienced multiple brief dips into negative territory, indicating episodic financial pressures without prolonged instability.

In contrast, Greece and Ukraine exhibit severe declines in $\hat{\kappa}_{ct}$ leading up to their respective disasters, with values reaching their lowest points and crossing well below the -0.5 threshold as the disasters unfolded. Greece's significant drop aligns with its 2012 default during the European debt crisis, while Ukraine's plunge reflects the escalation of interstate conflict in 2023. Other severe declines are observed in Sri Lanka and Ghana in 2022, coinciding with their defaults during severe economic and political instability.

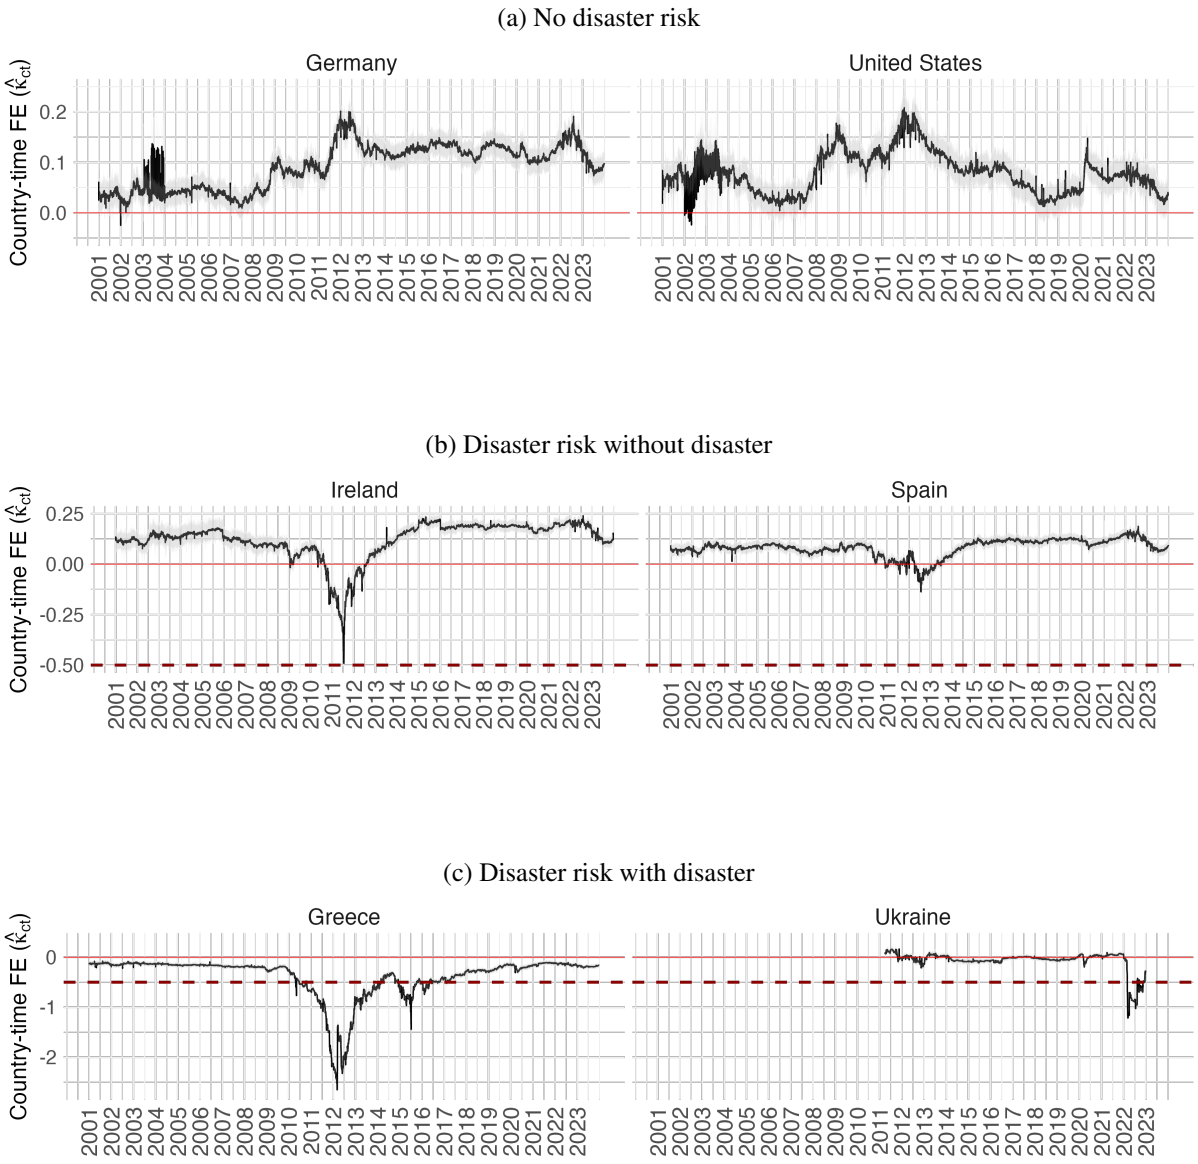

Figure 1.2: Evolution of $\hat{\kappa}_{ct}$ for selected countries

Notes: The figure shows the estimated evolution of $\hat{\kappa}_{ct}$, the country-time fixed effects from the regression in Table 1.3, interpreted as a residual component capturing the disaster wedge. Shaded grey bands represent 95% confidence intervals. The countries shown include Germany, the United States, Ireland, Spain, Greece, and Ukraine. Red lines at 0 and –0.5 serve as visual benchmarks.
Source: Author's calculations.

To derive disaster probabilities, I assume $\hat{\kappa}_{ct}$ corresponds to the disaster wedge. Specifically, $\hat{\kappa}_{ct}$ represents the average effect of disaster wedges across all maturities, i.e.,

$$\hat{\kappa}_{ct} = \frac{\sum_{N \in \mathbb{N}(c,t)} q_{Nct} - \hat{\beta}\hat{q}^{ND}_{Nct} - \hat{\chi}_N - \hat{\kappa}_{Nc} - \hat{\kappa}_{Nt}}{|\mathbb{N}(c,t)|} \tag{1.17}$$

where $\mathbb{N}(c,t)$ is the set of maturities available for country c at time t, and $|\mathbb{N}(c,t)|$ is the number of

them. The theoretical model then implies

$$\hat{\kappa}_{ct} = \frac{\sum_{N \in \mathbb{N}(c,t)} \sum_{\tau=1}^{N} \log\left(1 + \phi_{\delta}^{\tau-1} \delta_{1,ct} \left(J_{\tau,cN} - 1\right)\right)}{|\mathbb{N}(c,t)|}$$

Finally, I specify the type of disaster by inputting the estimated persistence parameter ($\hat{\phi}_{\delta}$) and the overall jump effect of the disaster ($\hat{J}_{\tau,cN}$) into the previous equation, leaving $\delta_{1,ct}$ as the only variable to be determined. The type of disaster is tailored to each country's context; for example, an interstate war is specified for Ukraine, while a sovereign default is specified for Greece. Given that the shortest maturity bond available is a one-year bond, $\delta_{1,ct}$ corresponds to the probability of a disaster occurring within one year.

I estimate $\delta_{1,ct}$ by minimizing the squared difference between $\hat{\kappa}_{ct}$ and the theoretical form of the disaster wedge:

$$\hat{\delta}_{1,ct} = \underset{\delta_{1,ct}}{\operatorname{argmin}} \left(\hat{\kappa}_{ct} - \frac{\sum_{N \in \mathbb{N}(c,t)} \sum_{\tau=1}^{N} \log\left(1 + \hat{\phi}_{\delta}^{\tau-1} \delta_{1,ct} \left(\hat{J}_{\tau,cN} - 1\right)\right)}{|\mathbb{N}(c,t)|} \right)^2$$

1.5 Results

Using the process explained in the previous section, I estimate investors' perceived probability of disaster for each country and day. To assess their predictive value, I first analyze how these probabilities evolve leading up to the disaster events. Second, I explore their relationship with credit ratings. Finally, I conduct a series of forecasting exercises using machine learning techniques to evaluate the predictive value of the estimated disaster probabilities.

1.5.1 Disaster probabilities before disasters

1.5.1.1 Sovereign defaults: Greece, Ghana, and Sri Lanka

Figure 1.3 shows the evolution of Greece's sovereign default probability during the European debt crisis, spanning from the deficit revelation in late 2009 to the execution of the Private Sector Involvement (PSI) agreement in March 2012. Key events are marked by dashed vertical lines. The probability rose sharply from 25% after late 2009 when the Greek government admitted its budget deficit was far higher than previously reported. Austerity measures were implemented, but at least an immediate impact is not observed. In April 2010, the default probability spiked dramatically from 40% to 60%. It then declined to 50% in May 2010, coinciding with the announcement of a €110 billion EU-IMF bailout package. However, this relief was short-lived, as the probability rose again above 70%, before falling back to 50%, aligning with the establishment of the European Financial Stability Facility (EFSF). The EFSF, a temporary crisis resolution mechanism, aimed to support Eurozone countries in distress by issuing bonds backed by guarantees from member states. Protests against austerity measures in May 2010 marked the beginning of a new, slow but steady increase in default probability, rising from 50% to 100% by June 2011. In July 2011, the EU announced a second bailout package that included a plan for voluntary debt restructuring. On the day of the announcement, the probability dropped by 20%. However, it quickly returned to 100% and remained elevated until the PSI agreement was finalized. The PSI included a 50%

haircut for private bondholders, but it was not until March 2012 that it was fully executed.

This case highlights the accuracy of financial markets in predicting sovereign default and their rapid reaction to new information and policy measures. The figure suggests that bailouts and international mechanisms have an immediate impact on default probabilities, while austerity measures alone do not.

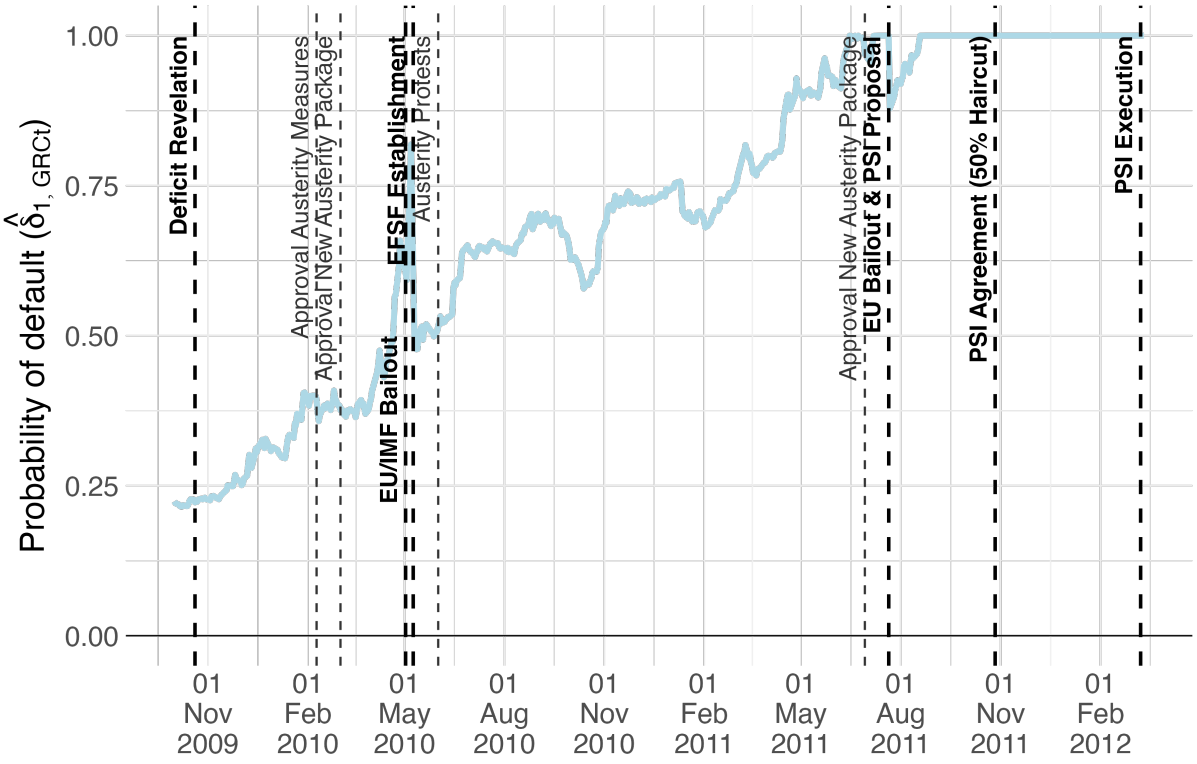

Figure 1.3: Evolution of probability of default for Greece

Notes: The figure shows the evolution of the probability of a sovereign default for Greece from October 2009 to March 2012. Bolded vertical lines mark key events: the 2009 deficit revision, the April 2010 yield spike, the May 2010 EU–IMF bailout, the creation of the European Financial Stability Facility (EFSF), the July 2011 restructuring announcement, and the final Private Sector Involvement (PSI) deal in March 2012. Other vertical lines mark additional developments such as the approval of austerity measures and major protest episodes.
Source: Author's calculations.

Figure 1.4 tracks Ghana's rising default probability from May 2018 to its official default in December 2022. The trajectory resembles Greece's case, as the probability of default steadily rises from 30% to 100% over several years. However, Ghana's experience includes fewer notable interventions. In May 2018, the probability of default stood at 30% and increased steadily to around 60% by the end of 2020. The COVID-19 lockdown caused a temporary spike in default probability, but it stabilized at approximately 60% thereafter. It was not until May 2022, when the Ghanaian government dismissed the possibility of seeking assistance from the IMF, that the probability sharply increased, eventually reaching 100%. Once the probability peaked at 100%, the government reversed its stance and began discussions with the IMF. Despite this shift, no immediate agreement was reached, and the probability remained at the 100% level. Toward the end of 2022, the government announced plans for a debt restructuring, followed shortly by a declaration of default on foreign debt.

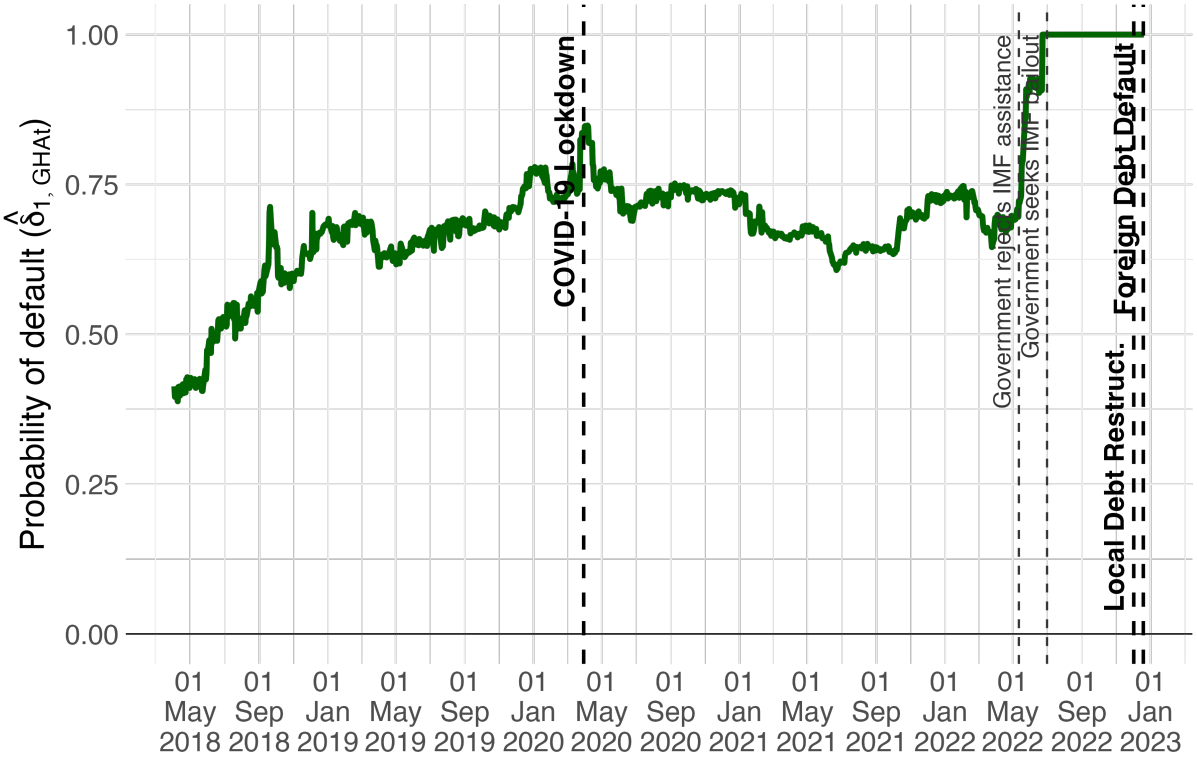

Figure 1.4: Evolution of probability of default for Ghana

Notes: The figure shows the evolution of the probability of a sovereign default for Ghana from May 2018 to December 2022. Bolded vertical lines mark key events: the COVID-19 lockdown, the announcement of debt restructuring, and the formal default on external debt. Other vertical lines indicate additional developments such as the May 2022 rejection of IMF assistance and the subsequent decision to seek an IMF bailout.
Source: Author's calculations.

Figure 1.5 depicts the probability of default for Sri Lanka from January 2022 to the declaration of default on foreign debt in April 2022. In February 2022, investors assigned a near-zero probability of default to Sri Lanka. This began to change in March, following a public statement by the IMF declaring Sri Lanka's debt unsustainable. Probability rose to 20% thereafter. Mass protests erupted, and a state of emergency was declared, keeping the probability stable around 20%. In the days leading up to the default, the probability surged to approximately 40%, and on the day of the default, it spiked further to 60%.

Although the probability did not reach 100%, the rapid increases leading up to the default provide clear evidence of investors' responsiveness of financial markets to developments even in short time frames.

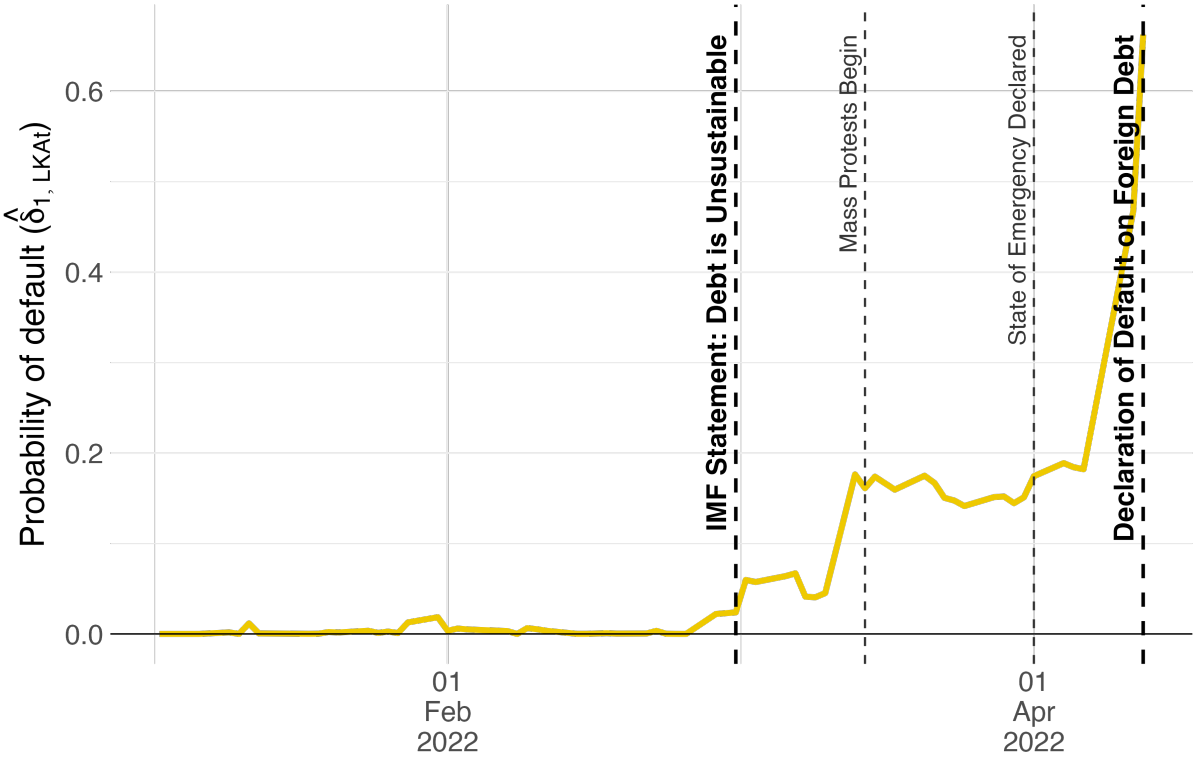

Figure 1.5: Evolution of probability of default for Sri Lanka

Notes: The figure shows the evolution of the probability of a sovereign default for Sri Lanka from January to April 2022. Bolded vertical lines mark key events: the IMF's March 2022 statement declaring debt unsustainable and the formal default on external debt in April. Other vertical lines indicate additional developments such as the outbreak of mass protests and the declaration of a state of emergency.
Source: Author's calculations.

1.5.1.2 Interstate war: Ukraine and Russia

Figure 1.6 illustrates the evolution of the estimated probability of an interstate war in Ukraine and Russia, from December 2021 to March 2022. Investors assigned virtually no probability to an interstate war until approximately two months before the conflict. Starting in January, their perception of war risk began to shift, with the probability of conflict in the two countries rising gradually to around 20%. Notably, even after Belarus's military drills on February 10 and Russia's recognition of the independence of Donetsk and Luhansk on February 21, investors' probability of a conflict remained relatively steady. It was only after the invasion commenced on February 24 that the estimated probability of interstate jumps to over 90%.

This finding suggests that, compared to sovereign defaults, such as the Greek and Ghanaian cases, interstate war appears harder to predict. Markets did begin to price in rising geopolitical risk ahead of the invasion, but not to levels commensurate with the immediacy of the conflict, even one week before it occurred. This aligns with the findings of Chadefaux (2017), which show that markets often react rather late to the outbreak of war.

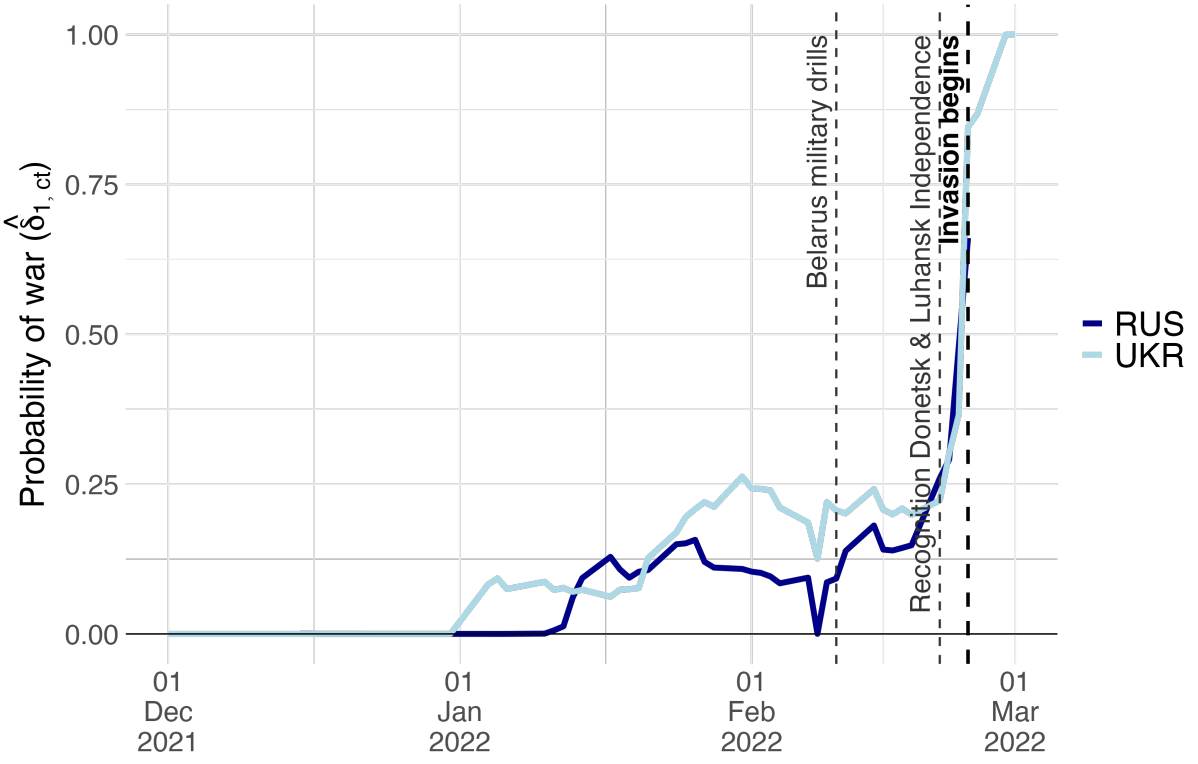

Figure 1.6: Evolution of probability of interstate war for Russia and Ukraine

Notes: The figure shows the evolution of the estimated probability of an interstate war for Russia and Ukraine from December 2021 to March 2022. Bolded vertical lines mark key events like the invasion of Ukraine. Other vertical lines represent additional developments such as Belarus's joint military drills with Russia and Russia's recognition of the Donetsk and Luhansk regions.
Source: Author's calculations.

1.5.2 Disaster probabilities and credit ratings

I analyze the relationship between estimated disaster probabilities and an established measure of default risk: credit ratings. Credit ratings are qualitative measures of a borrower's creditworthiness, reflecting the likelihood of default. Ratings incorporate a broad range of factors, including public debt levels, fiscal deficits, economic growth, political stability, and governance quality. Given that the disaster probabilities involve sovereign defaults and interstate wars—where default risk is the dominant factor—the comparison with credit ratings is appropriate. For this exercise, I use monthly averages of the estimated probabilities.

Figure 1.7 presents a boxplot of the estimated disaster probabilities across credit rating categories: AAA to B- (low risk), CCC+ to CC- (medium to high risk), and C+ to RD (very high risk or in default).[20] The boxplot illustrates the distribution of the probabilities, with the black line representing the median, the blue box indicating the interquartile range (25th to 75th percentiles), and the bars extending

20. The grouping of ratings into these categories is based on the S&P study (Rossi et al. 2023), which calculates sovereign default probabilities for each group based on observed transition frequencies. For ratings ranging from AAA to BB, the one-year default probabilities are near zero, with AAA at 0% and B at 4%. A sharp increase to 50% is observed when ratings fall within the CCC to CC category, while no data is reported for the C category.

to the 10th and 90th percentiles. For ratings from AAA to B-, the median of the default probabilities remains at 0%, with a compressed distribution reflecting the concentration of developed countries with minimal default risk. In contrast, the median of the probabilities rises to 30% for CCC+ to CC- ratings, showing significant variation in the distribution, and exceeds 90% for ratings from C+ to RD, where the distribution is concentrated at the upper end. The relationship between ratings and default probabilities is positive and notably non-linear. The AAA to B category encompasses the majority of rating values—approximately 80%—indicating a disproportionate concentration of low-risk ratings. This shows that the model's estimated probabilities follow a similar pattern to credit ratings.

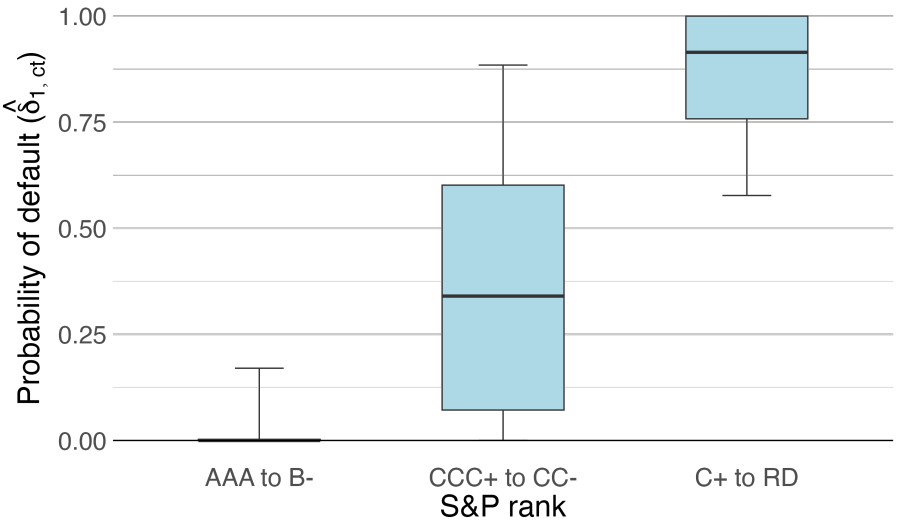

Figure 1.7: Relationship between disaster probabilities and credit ratings

Notes: The figure shows a boxplot of the estimated disaster probabilities by credit rating category. Credit ratings are obtained from Refinitiv Datastream, which reports long-term sovereign ratings from several agencies, including Moody's, Fitch, DBRS, and R&I. All ratings are converted to the S&P scale using Datastream's equivalence mapping. The boxplot illustrates the distribution of the probabilities, with the black line representing the median, the blue box indicating the interquartile range (25th to 75th percentiles), and the bars extending to the 10th and 90th percentiles.
Source: Author's calculations.

1.5.3 Predictive value of the estimated disaster probabilities

To assess the predictive value of the estimated disaster probabilities, I conduct a forecasting exercise to predict the occurrence of any disaster in the next month. Three specifications are used: one using only disaster probabilities, another using only credit rating variables, and a third combining both sources. The credit rating variables include the current rating, the magnitude of the last downgrade and upgrade, and the time since the last downgrade and upgrade.

For models using credit ratings, I employ a Random Forest, a machine learning technique that builds multiple decision trees and aggregates their predictions, allowing it to capture non-linear relationships and complex interactions.[21] When using disaster probabilities alone, they are applied directly, as they already represent one-year-ahead probabilities. The combined model includes both data sources and uses the same Random Forest specification.

21. The Random Forest model uses 500 trees, with 6 variables randomly selected at each split.

To evaluate the predictive performance of the models, I use two standard metrics: the ROC curve and the PR curve. The ROC curve plots the true positive rate (TPR) against the false positive rate (FPR) across different classification thresholds.[22] The TPR represents the proportion of actual positive events correctly identified relative to the total number of actual positive events, while the FPR indicates the proportion of negative events incorrectly classified as positive, relative to the total number of actual negative events. The ROC-AUC summarizes this curve into a single value, where 1.0 represents perfect discrimination and 0.5 indicates performance equivalent to random guessing. A higher ROC-AUC reflects the model's ability to effectively distinguish between positive and negative outcomes across all thresholds. In contrast, the PR curve evaluates the balance between precision (the proportion of correctly predicted positive events out of all predicted positives) and recall (the proportion of actual positive events correctly identified out of actual positives).[23] The ROC curve evaluates performance across all thresholds and is generally preferred for assessing overall discrimination. The PR curve focuses on how well the model captures the positive cases by comparing the share of correctly predicted positives among total predicted positives (precision) and among actual positives (recall), making it especially useful when the focus is on the positive class and class imbalance is present.

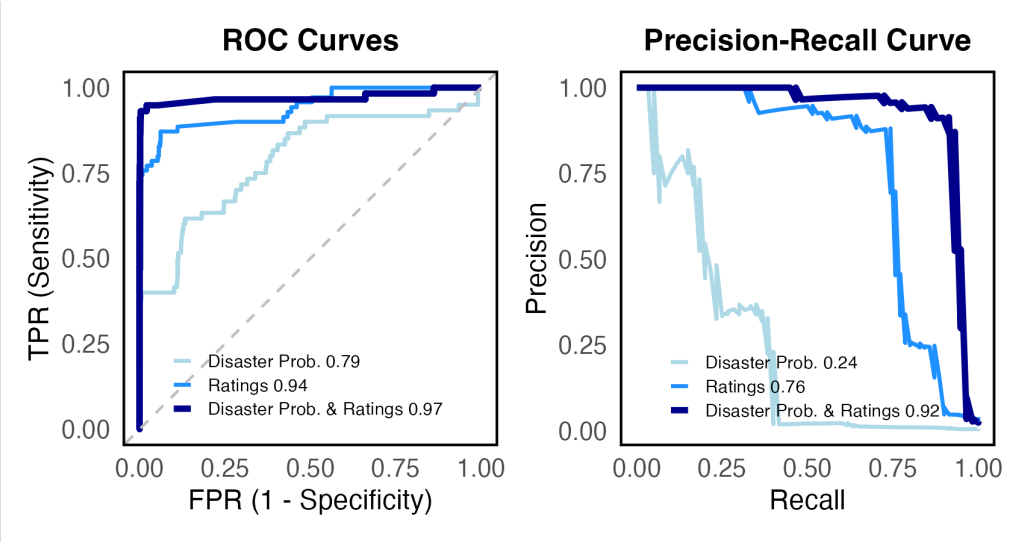

Figure 1.8: Predictive performance of disaster probabilities and credit ratings for disasters

Notes: The figure shows the predictive performance of three different models for forecasting disasters: using only disaster probabilities, using only credit rating variables (including the current rating, the magnitude of the last downgrade and upgrade, and the time since the last downgrade and upgrade), and using both sources. Disaster probabilities are used directly, while the other models rely on a Random Forest model. All models are implemented using an expanding window forecasting method.
Source: Author's calculations.

Figure 1.8 shows the predictive performance of the three models for forecasting disasters. Disaster probabilities perform reasonably well, achieving an AUC of 0.79 for the ROC curve and 0.24 for the PR curve. The Random Forest model utilizing credit ratings performs significantly better, with a ROC-AUC

22. TPR, also known as sensitivity, is defined as $\text{TPR} = \frac{\text{True Positives}}{\text{True Positives} + \text{False Negatives}}$. FPR is defined as $\text{FPR} = \frac{\text{False Positives}}{\text{False Positives} + \text{True Negatives}} = 1 - \text{Specificity}$.

23. $\text{Precision} = \frac{\text{True Positives}}{\text{True Positives} + \text{False Positives}}$. Recall is equivalent to TPR.

of 0.94 and a PR-AUC of 0.76. Notably, integrating disaster probabilities into the credit ratings model, further improves the ROC-AUC to 0.97 and the PR-AUC to 0.92. The exceptionally high performance of these models can be justified by the nature of defaults, which rarely occur as sudden surprises. Instead, they are often the result of prolonged processes and involve a self-fulfilling component as investor pessimism increases borrowing costs.

These results show that credit ratings remain the stronger standalone predictor with a very high performance. Yet, disaster probabilities improve the model further, highlighting the informational value of bond market data in forecasting events.

1.6 Conclusions

This paper introduces a novel approach to estimating investors' perceived probability of disaster from yield curve data. I provide daily estimates of the one-year-ahead disaster probability as perceived by investors for around 60 countries from 2000 to 2023. Then, I assess the predictive value of these probabilities through several exercises.

Disaster probabilities spike before major events, such as the debt restructurings in Greece and Sri Lanka, where probabilities reached 100% months in advance. In other cases like Ghana, the increase was more abrupt and closer to the default date. A similar pattern is observed for the onset of the Russia–Ukraine war, with a rise in the weeks preceding the invasion but not before. Furthermore, the estimated probabilities enhance the predictive power of credit ratings for predicting disaster events. Overall, while investors' beliefs are not perfect in all cases, the evidence shows that markets can still perform reasonably well and respond quickly to rising risks.

This model adopts a specific approach to estimate investors' beliefs, and the inferences are tied to this structure. However, future research can address its limitations and explore alternative frameworks. First, there is room for improvement in the performance of the theoretical model. Expanding the model to incorporate additional instruments, such as equities and corporate bonds, and developing a general equilibrium framework could more accurately capture the data-generating process. This would improve theoretical price estimations and enhance the identification of disaster risk. Another issue is the inability to distinguish between the probability of a disaster and its effect, as the model relies on general assumptions about the magnitude of disaster-related jumps. Calibrating these jumps to be country-specific or dependent on other variables could address this limitation. Additionally, the current approach requires specifying the type of disaster based on the country's context, limiting the model's use for long-term analysis as these contexts are likely to change. Integrating an NLP model to analyze news reports and text data could help identify not only the type of disaster but also its severity, allowing for a more precise estimation of disaster risk probabilities.

Finally, estimated disaster probabilities have broader uses beyond prediction, including identifying policies that reduce investor fears, assessing their welfare effects, and analyzing disaster spillovers.

Bibliography

Ang, Andrew, Monika Piazzesi, and Min Wei. 2006. "What does the yield curve tell us about GDP growth?" *Journal of Econometrics* 131 (1-2): 359–403.

Arrow, Kenneth J, Robert Forsythe, Michael Gorham, Robert Hahn, Robin Hanson, John O Ledyard, Saul Levmore, Robert Litan, Paul Milgrom, Forrest D Nelson, et al. 2008. "The promise of prediction markets." *Science* 320 (5878): 877–878.

Backus, David, Mikhail Chernov, and Ian Martin. 2011. "Disasters implied by equity index options." *The Journal of Finance* 66 (6): 1969–2012.

Barro, Robert J. 2006. "Rare disasters and asset markets in the twentieth century." *The Quarterly Journal of Economics* 121 (3): 823–866.

Barro, Robert J, and Gordon Y Liao. 2021. "Rare disaster probability and options pricing." *Journal of Financial Economics* 139 (3): 750–769.

Bekaert, Geert, Campbell R Harvey, Christian T Lundblad, and Stephan Siegel. 2016. "Political risk and international valuation." *Journal of Corporate Finance* 37:1–23.

Berkman, Henk, Ben Jacobsen, and John B Lee. 2011. "Time-varying rare disaster risk and stock returns." *Journal of Financial Economics* 101 (2): 313–332.

Bharath, Sreedhar T, and Tyler Shumway. 2008. "Forecasting default with the Merton distance to default model." *The Review of Financial Studies* 21 (3): 1339–1369.

Blinder, Alan S, Michael Ehrmann, Marcel Fratzscher, Jakob De Haan, and David-Jan Jansen. 2008. "Central bank communication and monetary policy: A survey of theory and evidence." *Journal of Economic Literature* 46 (4): 910–945.

Bluwstein, Kristina, Marcus Buckmann, Andreas Joseph, Sujit Kapadia, and Özgür Şimşek. 2023. "Credit growth, the yield curve and financial crisis prediction: Evidence from a machine learning approach." *Journal of International Economics* 145:103773.

Chadefaux, Thomas. 2017. "Market anticipations of conflict onsets." *Journal of Peace Research* 54 (2): 313–327.

Clark, Ephraim. 1997. "Valuing political risk." *Journal of International Money and Finance* 16 (3): 477–490.

Cochrane, John. 2009. *Asset pricing: Revised edition.* Princeton university press.

De Grauwe, Paul, and Yuemei Ji. 2013. "Self-fulfilling crises in the Eurozone: An empirical test." *Journal of International Money and finance* 34:15–36.

Dionne, Georges, Genevieve Gauthier, Khemais Hammami, Mathieu Maurice, and Jean-Guy Simonato. 2010. "Default risk in corporate yield spreads." *Financial Management* 39 (2): 707–731.

Duffee, Gregory R. 1999. "Estimating the price of default risk." *The Review of financial studies* 12 (1): 197–226.

Estrella, Arturo, and Gikas A Hardouvelis. 1991. "The term structure as a predictor of real economic activity." *The Journal of Finance* 46 (2): 555–576.

Estrella, Arturo, and Frederic S Mishkin. 1998. "Predicting US recessions: Financial variables as leading indicators." *Review of Economics and Statistics* 80 (1): 45–61.

Fama, Eugene F. 1970. "Efficient capital markets." *Journal of finance* 25 (2): 383–417.

Fang, Xiang, Bryan Hardy, and Karen K. Lewis. 2022. *Who Holds Sovereign Debt and Why It Matters.* Working Paper, Working Paper Series 30087. National Bureau of Economic Research, May.

Farhi, Emmanuel, and Xavier Gabaix. 2016. "Rare disasters and exchange rates." *The Quarterly Journal of Economics* 131 (1): 1–52.

Gabaix, Xavier. 2008. "Variable rare disasters: A tractable theory of ten puzzles in macro-finance." *American Economic Review* 98 (2): 64–67.

———. 2012. "Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance." *The Quarterly Journal of Economics* 127 (2): 645–700.

Gilchrist, S, E Zakrajsek, G Favara, and K Lewis. 2016. "Recession risk and the excess bond premium." *Fed Notes,* no. 8, 1–3.

Gilchrist, Simon, and Egon Zakrajšek. 2012. "Credit spreads and business cycle fluctuations." *American Economic Review* 102 (4): 1692–1720.

Gourio, Francois. 2012. "Disaster risk and business cycles." *American Economic Review* 102 (6): 2734–2766.

Huang, Tao, Fei Wu, Jing Yu, and Bohui Zhang. 2015. "International political risk and government bond pricing." *Journal of Banking & Finance* 55:393–405.

Julliard, Christian, and Anisha Ghosh. 2012. "Can rare events explain the equity premium puzzle?" *The Review of Financial Studies* 25 (10): 3037–3076.

Leippold, Markus, and Felix Matthys. 2022. "Economic policy uncertainty and the yield curve." *Review of Finance* 26 (4): 751–797.

Lorenzoni, Guido, and Ivan Werning. 2019. "Slow moving debt crises." *American Economic Review* 109 (9): 3229–3263.

Luckner, Clemens Graf von, Josefin Meyer, Carmen M Reinhart, and Christoph Trebesch. 2023. *Sovereign Debt: 200 years of creditor losses.* Working Paper, NBER Working Paper Series 32599. National Bureau of Economic Research.

Malkiel, Burton G. 2003. "The efficient market hypothesis and its critics." *Journal of Economic Perspectives* 17 (1): 59–82.

Merton, Robert C. 1974. "On the pricing of corporate debt: The risk structure of interest rates." *The Journal of Finance* 29 (2): 449–470.

Meyer, Josefin, Carmen M Reinhart, and Christoph Trebesch. 2022. "Sovereign bonds since Waterloo." *The Quarterly Journal of Economics* 137 (3): 1615–1680.

Pástor, L'uboš, and Pietro Veronesi. 2013. "Political uncertainty and risk premia." *Journal of Financial Economics* 110 (3): 520–545.

Reinhart, Carmen M, and Kenneth S Rogoff. 2010. "Growth in a Time of Debt." *American Economic Review* 100 (2): 573–578.

Remolona, Eli M, Michela Scatigna, and Eliza Wu. 2007. "Interpreting sovereign spreads." *BIS Quarterly Review, March.*

Rietz, Thomas A. 1988. "The equity risk premium a solution." *Journal of Monetary Economics* 22 (1): 117–131.

Ross, Steve. 2015. "The recovery theorem." *The Journal of Finance* 70 (2): 615–648.

Rossi, Luca, Nick W. Kraemer, and Vaishali Singh. 2023. *Default, Transition, and Recovery: 2023 Annual Global Sovereign Default and Rating Transition Study.* Technical report. Six sovereigns defaulted in 2023, reflecting challenges in emerging and frontier markets. S&P Global Ratings, Credit Research & Insights. https://www.spglobal.com/ratings/en/research-insights.

Schreindorfer, David. 2020. "Macroeconomic tail risks and asset prices." *The Review of Financial Studies* 33 (8): 3541–3582.

Smales, Lee A. 2016. "The role of political uncertainty in Australian financial markets." *Accounting & Finance* 56 (2): 545–575.

Wachter, Jessica A. 2013. "Can time-varying risk of rare disasters explain aggregate stock market volatility?" *The Journal of Finance* 68 (3): 987–1035.

Wolfers, Justin, and Eric Zitzewitz. 2004. "Prediction markets." *Journal of Economic Perspectives* 18 (2): 107–126.

Chapter 2

Caught in a Trap: Simulating the Economic Consequences of Internal Armed Conflict

Joint with Hannes Mueller, Institut d'Anàlisi Econòmica (CSIC)

2.1 Introduction

Internal armed conflict is recognized as a major and growing concern for development.[1] Beyond its immediate economic consequences, conflict further degrades material conditions and erodes social cohesion, setting the stage for more conflict. This leads to vicious cycles of violence—a phenomenon known as the *conflict trap*. Scholars have recognized the significant influence of the conflict trap on macroeconomic development (Collier et al. 2003; Rohner et al. 2013; Rohner and Thoenig 2021).[2] Yet, there is no analysis quantifying its effect on the long term. Such estimates would support the assessment of conflict's economic cost, clarify its developmental toll, and help quantify the value of investing in prevention strategies. A reason for this gap is that the conflict trap is hard to define as it persists from the conflict into the post-conflict period, requiring a structural model to evaluate its impact.

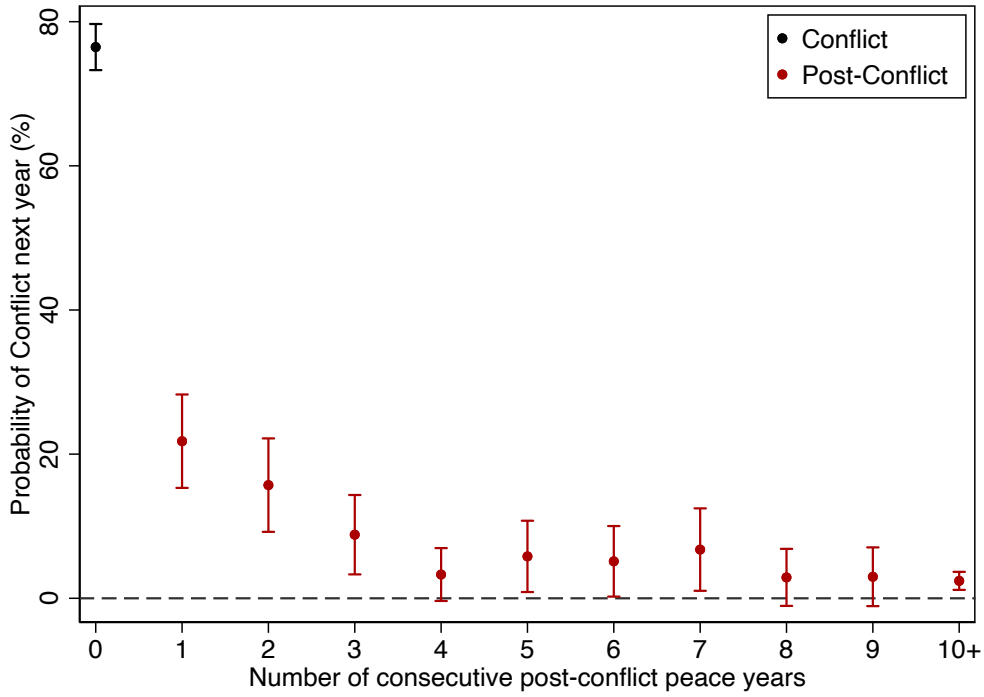

Figure 2.1: The conflict trap: high risk of conflict during conflict and post-conflict peace

Notes: The figure shows the risk of renewed outbreak for the different number of consecutive years of peace after conflict. The risk is estimated using the proportion of transitions observed in the data. Point estimates are displayed as circles where bars indicate significance at 5% using the Wald method. Conflict is defined as surpassing the threshold of violence intensity that significantly harms the economy, which corresponds to having more than 9.35 battle-related deaths per million inhabitants during a year.
Source: Authors' calculations based on population data from the World Bank's World Development Indicators, and number of battle-related fatalities from Uppsala Conflict Data Program's Georeferenced Event Dataset.

This article proposes a simple version of such a model. We follow the conflict literature and assume that conflict dynamics can be captured by a Markov process (Besley and Mueller 2012; Hegre et al. 2013; Besley et al. 2024) , which we extend to capture the conflict trap. Figure 2.1 shows the conflict dynamics we incorporate into the model. The y-axis shows the likelihood of internal armed conflict in the next year

1. Over half of the world's extreme poor is predicted to reside in countries marked by fragility, conflict, and violence (World Bank 2020). For an analysis, refer to Corral et al. (2020).
2. The conflict trap has proven useful for forecasting conflict risk; see Hegre et al. (2017) and Mueller and Rauh (2022).

during and after conflict.[3] The likelihood of remaining in conflict is very high, nearly 80%. In the first year of peace after conflict, the likelihood of a renewed outbreak is over 20%. It then falls monotonically to 3% after 4 years and stabilizes around that level. After 10 consecutive years of peace, the likelihood of resurgence is close to 2%. This means not only that conflict is self-perpetuating, but also that the post-conflict period is extremely risky - around half of the countries that escape from conflict will experience a resurgence before reaching 8 years of peace. This pattern of falling risk is extremely robust across time, sets of countries, and conflict definitions.

Our model of the conflict trap is a discrete-time Markov process where there is a state of conflict, multiple states of post-conflict peace, and a state of stable peace. Modeling the post-conflict period as a sequence of different stages allows us to model the conflict trap in a tractable manner without sacrificing the Markov assumption. We estimate the transition matrix from the history of armed violence and the GDP per capita growth distribution of each state through a country fixed effects regression. We simulate dynamics between conflict and peace by drawing transition paths from the estimated transition matrix and drawing growth rates from the corresponding distribution of the realized states. This allows us to estimate the distribution of developmental effects of the conflict trap.

We also explore the role of the severity of the conflict trap. We use a Machine Learning (ML) model with cross-validation to predict the extent of the trap for each country using fixed characteristics. We use this prediction to partition our dataset into samples that are more and less prone to trap dynamics.

We find that for the aggregate sample using the full dataset, entering a conflict results in an average GDP per capita loss of about 20% over 30 years. This loss escalates to 30% at the 75th percentile and nearly 45% at the 90th percentile. When partitioning the data, we find that within the conflict-prone sample, the average GDP per capita loss approximates 30%, with losses at the 90th percentile exceeding 50%. Even in the less conflict-prone sample, the economic impact remains substantial, with an average loss of about 15% and losses exceeding 35% at the 90th percentile.

This underscores the importance of aiding countries in both exiting the conflict trap and, more importantly, preventing them from falling into it, as key strategies for development. We also show that while the extent of the trap cannot be predicted perfectly, there is a predictable element to its dynamics. This makes our findings actionable in the sense that macroeconomic development in the most at-risk countries can be promoted through interventions designed to counter challenges posed by conflict.

Our statistical model provides a way to integrate post-war growth dynamics into the empirical discussion of the long-term effects of internal conflict. Long-run macroeconomic growth and the extent of violent conflict have often been linked theoretically (Rodrik 1999; Collier 2007; North et al. 2009; Besley and Persson 2011a). However, causal identification of the link running from conflict to long-run outcomes has been plagued by concerns of reverse causality and omitted variable bias. Empirical macro studies therefore tend to focus on the contemporaneous growth effects of war episodes (Collier 1999; Cerra and Saxena 2008; Mueller 2016; De Groot et al. 2022). Moreover, standard macroeconomic theory suggests that growth is particularly strong after civil wars as capital stocks re-adjust to higher productivity (Collier 1999). This would suggest that the long-term economic costs of conflict are relatively low. Although there is some empirical support for this hypothesis for external wars (Organski and Kugler 1977, Davis and Weinstein 2002, Steven Brakman and Schramm 2004, Miguel and Roland 2011), in the macro data there are no visible growth spurts at the end of internal conflicts (Cerra and Sax-

3. Our conflict definition is discussed below. Mueller and Rauh (2022) exploit this same pattern to forecast conflict risk.

ena 2008, Mueller 2012). In their review of the latest literature, Rohner and Thoenig (2021) emphasize the presence of macro-level complementarities that perpetuate "war traps" with devastating effects on long-term development.[4] A rationale for the absence of readjustments after internal wars is that human capital is heavily affected and cannot readjust easily (Barro and Sala-i-Martin 2004).[5] Pinotti (2015) shows that the rise of mafia activity in southern Italy led to a lasting slowdown in economic development.[6] Recent studies also show that conflict can cause broader and persistent economic disruptions by damaging supply chains and production networks (Couttenier et al. 2022; Federle et al. 2024; Korovkin et al. 2024). We quantify the long-term effects of averse trap dynamics to demonstrate that a crucial aspect is the post-war risk of re-emerging conflict. Importantly, we use fixed effects growth estimates in our simulations and find relatively minor differences in the contemporaneous growth damages of conflict between the samples. Fixed country characteristics therefore drive long-term economic outcomes through the conflict trap dynamics they are associated with.

Our conflict trap model generates averse cycles of growth. These growth dynamics can be linked to macroeconomic studies on growth and volatility (Easterly et al. 1993; Ramey and Ramey 1995; Berg et al. 2012), particularly those identifying the cyclical growth patterns of emerging countries (Aguiar and Gopinath 2007; Garcia-Cicco et al. 2010). Aguiar and Gopinath (2007) note that emerging market growth is characterized by shocks to trend growth rather than transitory fluctuations around a stable trend. In other words, growth histories in emerging markets are characterized by long-lasting episodes in which economic growth is consistently low which are then followed by sudden growth spurts with a different trend growth. This is the kind of growth behavior that conflict traps generate but the two have not been linked in the macroeconomic literature.

Finally, our methodology for endogenizing trap dynamics is based on the literature that analyzes the causes of conflict. We show issues like geographic features (Nunn and Puga 2012), political institutions (North et al. 2009; Besley and Persson 2011b; Robinson and Acemoglu 2012), natural resources (Dube and Vargas 2013; Bazzi and Blattman 2014; Berman et al. 2017) and ethnic and religious composition (Esteban and Ray 1994; Montalvo and Reynal-Querol 2005; Michalopoulos and Papaioannou 2016) can predict some of the variation in the extent of conflict and the conflict trap. We build on this literature by predicting both the extent of violence and time spent in the conflict trap in the period 1989-2021 using these factors and other factors like GDP per capita. We find that predictable conflict risk is associated with substantial long-term economic costs through a worsening of conflict trap dynamics.

The following section outlines the conceptual model of the conflict trap. Section 2.3 is dedicated to the model's estimation. In Section 2.4, we present and discuss the results, followed by our conclusions in the final section.

2.2 A model of the conflict trap

This section presents our empirical model of the conflict trap. Our central assumption is that a country's transition between conflict and peace can be described by a discrete-time Markov process.

Let s_t represent the state at period t. It can take values from the state space defined as $S =$

4. See also Mueller and Tobias (2016).

5. There is growing evidence that the costs of internal wars for health and human capital formation are severe (Ichino and Winter-Ebmer 2004; Blattman and Annan 2010; Leon 2012; Akresh et al. 2012; Tapsoba 2023).

6. The effect is driven by reduced private investment and a shift toward less productive public capital.

$\{0, 1, \ldots, i, \ldots, \tau, \tau+1\}$. When $s_t = 0$, the country is in the state of *conflict*. When $s_t = k$ s.t. $k \in [1, \tau]$, the country is in the kth consecutive year of post-conflict peace. States 1 to τ capture the post-conflict period in a strict chronological order. We refer to this post-conflict peace phase as *unstable peace*. Finally, when $s_t = \tau + 1$, the country is in *stable peace*, i.e. it is outside of the conflict trap. Stable peace represents a state that is not conditioned by the conflict history of a country – once a country escapes the conflict trap, transitions to conflict are as likely as in a country that never had a conflict.

When a country is in conflict ($s_t = 0$), it can either stay in conflict or transition to the first year of post-conflict peace ($s_t = 1$). When a country is in its first year of peace, it can either return to conflict or transition to the second year of post-conflict peace ($s_t = 2$). This pattern repeats until the country reaches τ consecutive years of post-conflict peace ($s_t = \tau$). At this point, the country can either go back to conflict or transition to stable peace ($s_t = \tau + 1$). The parameter τ determines the threshold of consecutive years of peace needed to escape from the conflict trap. Finally, a country in stable peace can either return to conflict or stay in stable peace.

By the structure of the Markov process we assume, a country can only transition in two directions, to conflict or to one more year of peace. This allows us to simplify the notation of the transition probabilities, making π_i to be the probability of transition to conflict from state i. Thus, $1 - \pi_i$ is the probability of adding a year of peace from state i. The transition matrix is

$$\mathbf{\Pi} = \begin{pmatrix} \pi_0 & 1-\pi_0 & 0 & \ldots & 0 & 0 \\ \pi_1 & 0 & 1-\pi_1 & \ldots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ \pi_\tau & 0 & 0 & \ldots & 0 & 1-\pi_\tau \\ \pi_{\tau+1} & 0 & 0 & \ldots & 0 & 1-\pi_{\tau+1} \end{pmatrix}$$

Note that the assumption of a Markov process implies that the dynamics in all countries can be described by a single transition matrix. Importantly, this also means that the likelihood of re-entering the conflict during the stabilization process or from stable peace is not a function of the country's longer history. This is a stark simplification as much of the conflict literature links armed conflict outcomes to fixed characteristics like the colonial history of a country or its geographic features. However, we show in Section 2.3.2 that predicting which countries will get stuck in the trap based on pre-determined features is surprisingly difficult.

To capture the impact of conflict on development we focus on real GDP per capita. Denote GDP_t as the real GDP per capita when the Markov process is at period t. The realization of the state affects GDP_{t+1} by determining the distribution from which growth, u_{t+1}, is drawn, i.e., $u_{t+1}|s_{t+1} \sim f_{i=s_{t+1}}$. Then, at each period, GDP_{t+1} updates according to

$$GDP_{t+1} = GDP_t(1 + u_{t+1}) \tag{2.1}$$

For example, if a country is in conflict at period 2 ($s_2 = 0$), u_2 is drawn from f_0. Then, if $u_2 = -0.01$, it implies that GDP per capita is reduced by 1% from period 1 to 2 in the state of conflict.

We consolidate all the growth distributions into a single vector $\boldsymbol{f} = \begin{pmatrix} f_0 & f_1 & \ldots & f_{\tau+1} \end{pmatrix}$ which we refer to as *growth vector*. We expect to draw lower growth from f_0 compared to the other elements of the growth vector. This kind of model will therefore lead to dynamics in which countries cycle back

and forth between high and low growth episodes.

2.3 Estimation

2.3.1 Defining conflict and conflict trap's length

Our definition of conflict is based on surpassing the threshold of violence intensity that significantly harms the economy. We measure violence intensity by calculating the number of battle-related deaths per capita using the Uppsala Conflict Data Program's Georeferenced Event Dataset (UCDP/GED)[7] together with population data from the World Bank's World Development Indicators (WB/WDI).[8] To find this threshold, we identify country-years with violence, arrange them in order of violence intensity, and bin them into deciles. We then run a country fixed effects regression of growth on these deciles, controlling for time fixed effects. We find that an appropriate threshold to define conflict is having as much violence as our seventh decile (top 40% most violent years), which corresponds to having more than 9.35 battle-related deaths per million inhabitants during a year.[9] We also show robustness checks for more restrictive thresholds to define conflict.[10]

From Figure 2.1 we know that the risk of a renewed outbreak next year stabilizes for higher states. We choose $\tau = 7$ so be sure to have enough observations to have meaningful transition likelihoods.[11] This leads to a total of 9 states. We also conducted a robustness check with a shorter conflict trap.[12]

2.3.2 Predicting the extent of the conflict trap

To study the role of the conflict trap at the intensive margin, we employ a ML model with cross-validation to predict its extent for each country. This prediction enables us to partition our dataset into a more and less conflict-prone sample. By re-running our analysis separately on these subsets, we quantify the effects of more and less pervasive types of conflict traps. This approach also enables us to identify which countries are at higher risk of being trapped.

A comprehensive dataset of country characteristics is constructed from various sources: Nunn and Puga (2012) (legal origin France, Africa dummy, ruggedness, % fertile soil, % desert, distance to coast, longitude and latitude), Montalvo and Reynal-Querol (2005) (ethnic polarization/fractionalization and religious polarization/fractionalization), Polity5 dataset (polity2 score, executive constraints, executive openness, and executive competitiveness) and the WB/WDI (natural resource dependence of GDP, GDP, and population). We use only the values of these variables that are pre-determined in our sample period to minimize concerns of reverse causality. This set reflects structural and highly persistent characteristics associated with conflict, as emphasized in the conflict literature. While some measurement error is

7. UCDP defines an event as: "An incident where armed force was used by an organized actor against another organized actor, or against civilians, resulting in at least 1 direct death at a specific location and a specific date". We are using the "best" estimate of fatalities summing all types of violence and aggregating all fatalities for each country-year.

8. See Mueller (2016) for a more detailed analysis along these lines using subnational data.

9. The results are displayed in Figure B.1 in Appendix B.2. A clear pattern emerges in which the most intense conflicts are also associated with the largest contraction of growth. The coefficients at the 7th, 8th, 9th, and 10th deciles are significantly negative at the 10% level, and among these deciles, only the coefficient at the 8th decile is not significant at the 5% level.

10. They lead to even larger estimates of the impact of internal armed conflict. See Figure B.6 in Appendix B.2.

11. Mueller and Rauh (2022) show that conflict history loses its predictive power for renewed outbreaks between 4 and 10 years.

12. We tested $\tau = 4$ and the results are similar. See Figure B.4 in Appendix B.2.

inevitable, we rely on standard sources widely used in the empirical literature. For a more detailed explanation of how missing data was addressed, see Appendix B.1.

We employ two ML models: a linear Lasso regression and a Random Forest. The predicted score is the share of years that the country is in the conflict trap in the period 1989-2021. We then use cross-validation to tune hyperparameters[13] and to calculate the R-squared statistics. For the Random Forest, the cross-validated R-squared score is 0.128, for the Lasso regression this is 0.149, and when we combine both through an average we get 0.173. We therefore always stick to the average of the two models (ensemble). Predicting the extent of the conflict trap in this way is relatively hard ex-ante with an R-squared of less than 20 percent. The out-of-sample prediction ensures that we only use predictable variation of which countries are more conflict-affected in our analysis. We find that GDP per capita levels are a main predictor of trap risk.[14]

Using the fitted value from this exercise we generate two samples. Countries are binned into deciles based on their prediction scores. Figure 2.2 displays the model fit, with cross-fitted values on the x-axis and actual variation in conflict trap exposure on the y-axis. We split the sample so that each sub-sample contains a similar total number of years of conflict. The conflict-affected group defined as the countries in the 9th and 10th deciles would go on to spend half their years in the conflict trap. We will return to this point after presenting our main results in Section 2.4.

We refer to the specification of the whole data set as *aggregate sample*, to the less conflict-prone sample as *peaceful sample*, and to the more conflict-prone as *conflict sample*.

2.3.3 Transition matrix and growth vector

The transition matrix is estimated using the proportion of transitions observed in the data. The corresponding estimates for each specification are shown in Table 2.1. When using the whole dataset, the likelihood of staying in an additional year of conflict is 76%. This leads to an expected duration of uninterrupted conflict of 4.17 years.[15] Once conflict ends, the likelihood of going back to conflict falls dramatically to 22% in the first year after conflict, 16% in the second year, and 7% after seven years in peace. Finally, the baseline likelihood of conflict is just 2% in stable peace. On average, it takes 18.46 years, i.e. almost two decades, to escape from the trap.[16]

Comparing the other two specifications, the transition probabilities show that the likelihood of transitioning to conflict is significantly higher in the conflict sample for most of the states, indicating a higher tendency to remain in the conflict trap. Specifically, the expected duration of uninterrupted conflict is 3.45 years for the peaceful sample and 5.56 years for the conflict one. Meanwhile, the expected number of periods in the conflict trap is 15.72 for the peaceful sample and 24.49 for the conflict one. Note that the pattern of high persistence in conflict and falling risks in post-conflict is robust across samples. However, the conflict sample suffers a significantly higher baseline risk of conflict outbreaks from stable peace.

To estimate the growth vector, we regress GDP per capita growth on a set of dummy variables d_{ijt}

13. Everything is implemented in Python using the `sklearn` package. We find $Lasso - alpha : 0.0001, RandomForest - max\_depth : 4, min\_samples\_leaf : 20, n\_estimators : 500$.

14. See features' importances in Figure B.2 in Appendix B.2.

15. We calculate it using the geometric distribution formula: $1/(1-\hat{\pi}_0)$.

16. This statistic is obtained using the fundamental matrix of $\hat{\mathbf{\Pi}}$, treating stable peace as an absorbing state. The fundamental matrix is defined as $\boldsymbol{F} = (\boldsymbol{I} - \boldsymbol{Q})^{-1}$, where $\boldsymbol{Q}$ is the submatrix of $\hat{\mathbf{\Pi}}$ excluding the absorbing states.

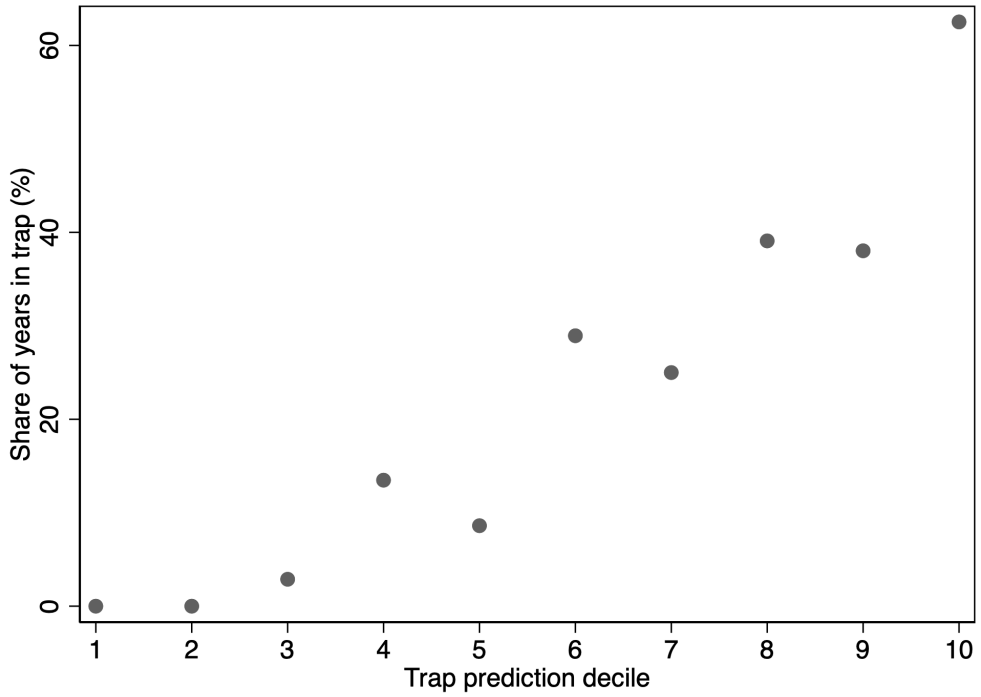

Figure 2.2: Risk deciles and extent of conflict trap

Notes: The figure shows the ensemble prediction value binned by decile together with the average extent of the conflict trap in each of the bins. The 20 countries in the top trap risk decile spent more than 60% in the conflict trap in the period 1989-2021. *Source:* Authors' calculations based on replication dataset from Nunn and Puga (2012) (legal origin France, Africa dummy, ruggedness, % fertile soil, % desert, distance to coast, longitude and latitude), Montalvo and Reynal-Querol (2005) (ethnic polarization/fractionalization and religious polarization/fractionalization), Polity5 dataset (polity2 score, executive constraints, executive openness, and executive competitiveness) and the WB/WDI (natural resource dependence of GDP, GDP, and population).

that correspond to the state of country j at year t. We use GDP (in constant 2015 US$) covering the period from 1989 to 2021 across over 190 countries. The regression equation is

$$Growth_{jt} = \sum_{i=0}^{\tau} \beta_i d_{ijt} + \mu_j + \gamma_t + \varepsilon_{jt}, \tag{2.2}$$

where μ_j and γ_t represent country and time fixed effects respectively. The omitted category is stable peace. The results for each specification are shown in Table 2.2. In all specifications, a year in conflict lowers growth by more than 3 percentage points, which lies inside the bounds of the literature. Importantly, the different conflict samples produce very similar findings regarding the growth effects of conflict. The coefficients of the rest of the states are small and insignificant in all specifications.

Both concerns of omitted variable bias and reverse causality apply to cross-country regressions. There is now, however, a large literature that shows that conflict has large economic causal effects at the micro level,[17] and Rohner and Thoenig (2021) concludes that these probably still constitute a lower bound. The causal interpretation is in line with the finding that more intense violence is associated with

17. Guiso et al. (2009), Cassar et al. (2011), and Rohner (2011) provide evidence that conflict diminishes economic activity by eroding trust, cooperation, and trade. Conversely, findings on the reverse channel at the country level have given contradictory or weak results Bazzi and Blattman (2014) and Berman and Couttenier (2015).

Table 2.1: Estimated transition probabilities

	Stay in Conflict	Unstable Peace to Conflict							Stable Peace to Conflict
	$\hat{\pi}_0$	$\hat{\pi}_1$	$\hat{\pi}_2$	$\hat{\pi}_3$	$\hat{\pi}_4$	$\hat{\pi}_5$	$\hat{\pi}_6$	$\hat{\pi}_7$	$\hat{\pi}_8$
AS	0.76	0.22	0.16	0.09	0.03	0.06	0.05	0.07	0.02
PS	0.71	0.20	0.12	0.07	0.03	0.07	0.04	0.06	0.01
CS	0.82	0.24	0.23	0.12	0.03	0.03	0.08	0.08	0.05

Notes: The table shows the estimated transition probabilities for each specification: the aggregate sample (AS), the peaceful sample (PS), and the conflict sample (CS). They are estimated using the proportion of transitions observed in the data. $\hat{\pi}_i$ denotes the probability of transitioning to conflict from state i, where $i = 0$ is conflict, $i = k$ s.t. $k \in [1, 7]$ is the kth number of consecutive years in post-conflict peace, and $i = 8$ is stable peace. Conflict is defined as surpassing the threshold of violence intensity that significantly harms the economy, which corresponds to having more than 9.35 battle-related deaths per million inhabitants during a year.
Source: Authors' calculations based on population data from WB/WDI, and number of battle-related fatalities from UCDP/GED. The methodology for sample partitioning is described in Section 2.3.2

Table 2.2: Estimation of the growth vector

	Aggregate Sample	Peaceful Sample	Conflict Sample
	GDP per capita Growth		
Conflict	-0.0320***	-0.0335**	-0.0320**
	(0.00823)	(0.0112)	(0.0111)
1st Year Post-conflict Peace	0.00481	0.00419	0.00568
	(0.00761)	(0.0108)	(0.00894)
2nd Year Post-conflict Peace	0.00176	0.00494	-0.00276
	(0.00713)	(0.0102)	(0.00710)
3rd Year Post-conflict Peace	0.000921	0.000639	-0.00111
	(0.00633)	(0.00790)	(0.0106)
4th Year Post-conflict Peace	0.00608	0.00495	0.00355
	(0.00562)	(0.00700)	(0.00998)
5th Year Post-conflict Peace	0.000295	0.000306	-0.000776
	(0.00503)	(0.00570)	(0.00977)
6th Year Post-conflict Peace	0.00217	0.000923	0.00112
	(0.00461)	(0.00541)	(0.00950)
7th Year Post-conflict Peace	0.00155	-0.00387	0.0122
	(0.00417)	(0.00455)	(0.00840)
Observations	5730	4676	1054
Country FE	Yes	Yes	Yes
Time FE	Yes	Yes	Yes
(Within country) R^2	0.122	0.175	0.0815

Notes: The table shows the regression results from equation 2 used to estimate the growth vector for each specification: the aggregate sample, the peaceful sample, and the conflict sample. A country fixed effects model is employed where GDP per capita growth is regressed on a set of dummies representing the states of the model. The base category is stable peace, defined as having more than 7 consecutive years of post-conflict peace. Year fixed effects were also included as control variables. Conflict is defined as surpassing the threshold of violence intensity that significantly harms the economy, which corresponds to having more than 9.35 battle-related deaths per million inhabitants during a year. Robust standard errors are in parentheses. $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.
Source: Authors' calculations based on GDP (constant 2015 US$) and population data from WB/WDI, and number of battle-related fatalities from UCDP/GED. The methodology for sample partitioning is described in Section 2.3.2.

more dramatic declines in growth.[18] In any case, we are using reliable estimates with a reasonable claim to identification and outside validity. These results are also robust to a more demanding specification with country time trends.[19] Lastly, considering recent debates on the problems of two-way fixed effects in cases of potentially heterogeneous treatment (De Chaisemartin and d'Haultfoeuille 2020; Callaway and Sant'Anna 2021; Sun and Abraham 2021),[20] we can show that these are not an issue in our case.[21]

Since stable peace is the omitted category, we normalize growth for this state to 0, i.e. $\hat{f}_{\tau+1} = 0$. For the rest of the states, the growth distribution follows a normal distribution with the estimated coefficient as the mean and the standard error as the standard deviation. Formally, for $h \in S$ s.t. $h \leq \tau$, $\hat{f}_h = \mathcal{N}(\hat{\beta}_h, \hat{\sigma}_h^{SE})$ with $\hat{\sigma}_h^{SE}$ being the standard error of β_h. Therefore, the estimated growth vector is given by the growth estimates from Table 2.2 with a 0 added, i.e., $\hat{\boldsymbol{f}} = \begin{pmatrix} \mathcal{N}(\hat{\beta}_0, \hat{\sigma}_0^{SE}) & \mathcal{N}(\hat{\beta}_1, \hat{\sigma}_1^{SE}) & \dots & 0 \end{pmatrix}$. This means that GDP per capita growth is stochastic for all states but not for stable peace due to the 0 growth normalization that we impose. Assumptions regarding the baseline growth do not affect the relative GDP losses we show in the results.

Note that the different specifications have very similar growth vectors. This suggests that any variation between them in the simulation results can be mainly attributed to differences in conflict dynamics, rather than to variations in the impact of conflict itself. Furthermore, a reduction of 3 percentage points is significant, but it is unclear whether such a number can explain large cross-country differences.

2.3.4 Simulation

Given our estimated Markov model composed by the estimated transition matrix ($\hat{\boldsymbol{\Pi}}$) and growth vector ($\hat{\boldsymbol{f}}$), we can simulate growth paths that countries experience as they move through the state space.

We make a thought experiment in which a control group of countries is always in stable peace which always generates 0 growth. The treated group of countries suffers an outbreak of conflict. Treated countries start with $GDP_0 = 100$ and in the state of conflict.[22] Importantly, we make stable peace absorbing ($\hat{\pi}_8 = 0$) to capture the net aggregate effect of falling into the conflict trap once. We draw transition paths from the estimated transition matrix and the growth from the corresponding distribution in the growth vector. The simulation has $T = 30$ periods and it is repeated $N = 100000$ times to get a good sense of the distribution of the GDP loss (compared to $GDP_t = 100$) over time.

Note that we only impose a starting year in conflict, along with its associated immediate economic impact. The long-term losses stem from the repeated outbreaks of violence a typical country will suffer after this initial onset. In this way, we are able to study the effect of the conflict trap in isolation. This provides a nuanced interpretation of the results, which differs from the conflict literature where aggregate losses are typically reported conditional on the specific duration of the conflict episode.

18. In our robustness analysis with stricter conflict definitions, we observe that higher violence intensity results in greater economic losses. See Figure B.1 in Appendix B.2.

19. See Table B.10 in Appendix B.3.

20. See De Chaisemartin and d'Haultfoeuille (2023) for a related survey.

21. Using the command `twowayfeweights` we find that the proportion of negative weights is less than 5% and their sum is -0.0018.

22. The distribution of states in period 0 is $\boldsymbol{p_0} = \begin{pmatrix} 1 & 0 & \dots & 0 \end{pmatrix}'$.

2.4 Results

The results of the simulations for each specification are shown in Figure 2.3. The x-axis of the figure counts the years after an outbreak of conflict in the treated sample. Since all treated countries start in conflict, the loss increases sharply. Then, as time goes by, more countries first leave ongoing conflict and escape the conflict trap. When a large part of the treated samples reaches absorbing stable peace growth converges back to the benchmark's growth rate.

In the aggregate sample (shown on the left), the average loss in GDP per capita after 30 years is almost 20%. This is a large effect. For comparison, the median growth of GDP per capita in the 30 years between 1990 to 2020 was about 50%. Importantly, there is a large heterogeneity across simulations with the 75th percentile experiencing a decline of 30% while the 90th percentile declines by almost 45%. In other words, doing better or worse inside the conflict trap can explain substantial changes in the long run.[23]

When comparing the two other samples, the higher conflict tendency of the conflict sample leads to significantly greater losses, the average loss in the conflict sample is almost twice as high than in the peaceful sample. More strikingly, the 90th percentile of the conflict sample reaches losses above 50%. Losses in the 90th percentile in the peaceful sample are 35% which makes the conflict trap important even here. Overall, we get a good understanding that both the extensive and intensive margins matter.

As discussed in the previous section the larger losses in the conflict sample are the result of a more severe conflict trap dynamics with more persistent conflict and a higher likelihood of re-surging conflict in the years after conflict. Keep in mind that our classification was based on cross-validated predictions which suggest that the conflict sample would spend around half of its time in the conflict trap. Our results suggest that if such a sample can be identified today—perhaps through existing conflict histories or other structural factors—we can expect severe macroeconomic effects of the conflict trap for this sample.

2.5 Conclusion

This study aims to enhance our comprehension of the long-term consequences of the conflict trap. To achieve this, we propose a simple framework that combines a probabilistic model of conflict dynamics with within-country estimates of economic costs. Together, we can simulate the growth trajectories that countries follow as they navigate through periods of conflict and peace. We use ML techniques to predict conflict trap's extent for each country, allowing us to classify our dataset into more and less conflict-prone samples and conduct separate analyses on these groups.

The simulation results show that entering conflict will induce an average loss in GDP per capita after 30 years by close to 20%, a loss of 30% in the 75th percentile, and a decline of nearly 45% in the 90th percentile. The loss in the 90th percentile for the conflict sample is more than 50% and for the peaceful one is still 35%. We only impose one outbreak of conflict, and the remainder of the losses are due to the effect of the estimated conflict dynamics. This underscores that aiding countries in exiting the conflict trap, and more importantly, preventing them from falling into it in the first place, are key to development.

23. We checked whether the random growth element coming from the growth regression alone can explain some of this variation. We find that the long-run level changes that could be explained by this part are small. For more information, see Figure B.5 in Appendix B.2.

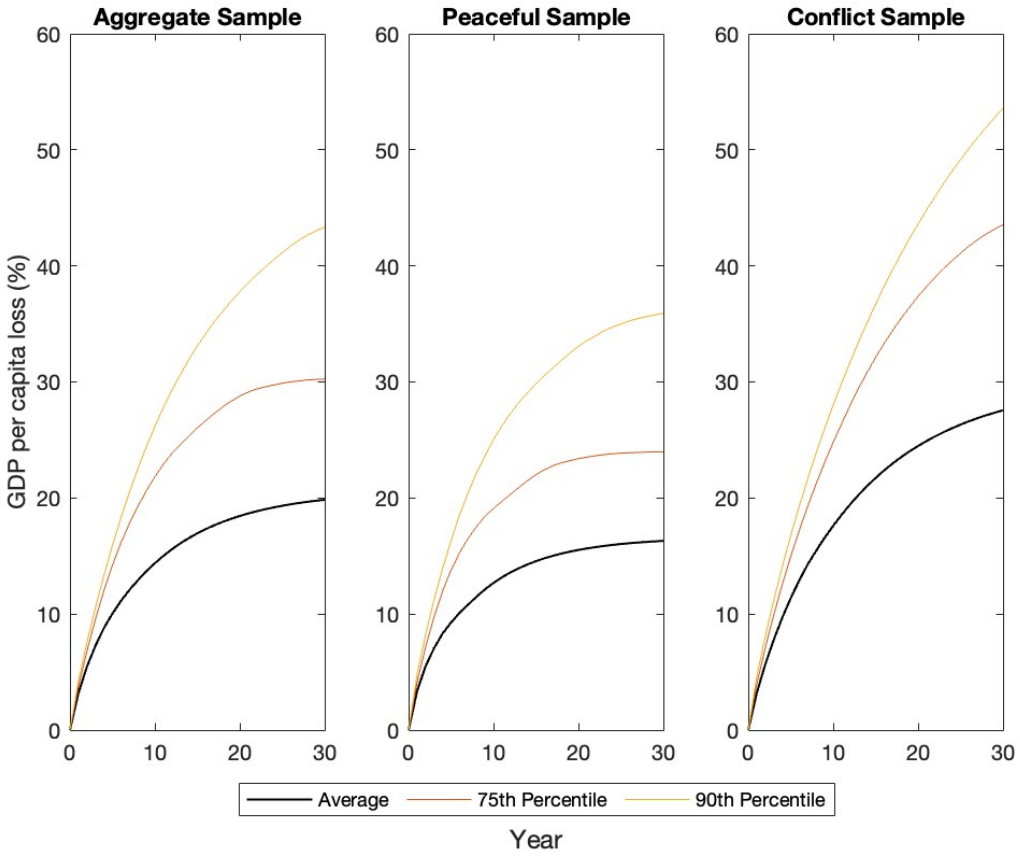

Figure 2.3: Evolution of GDP per capita loss

Notes: The figure shows the evolution of GDP per capita loss due to entering into conflict for each specification: the aggregate sample, the peaceful sample, and the conflict sample. Growth paths are simulated for countries as they transition through the state space, utilizing the respective estimated transition matrix from Table 2.1 and growth vector from Table 2.2. Countries start in conflict and stable peace is absorbing. Conflict is defined as surpassing the threshold of violence intensity that significantly harms the economy, which corresponds to having more than 9.35 battle-related deaths per million inhabitants during a year. The distribution at each period is described by the average, 75th percentile, and 90th percentile. The horizon for simulation is $T = 30$, and the number of simulations conducted is $N = 100000$. The methodology for sample partitioning is described in Section 2.3.2

Another takeaway from our method is that predicting the extent of the conflict trap both in the long run and in the short run should help target preventative policies. The model we estimate using a simple ML ensemble produces a highly imperfect forecast which could nonetheless help flag countries that would be affected by the conflict trap much more severely. If such a sample can be identified today, thinking about policies for escaping the conflict trap becomes a prerogative from a macroeconomic perspective.

Our model is modular in the sense that it allows using estimates derived from different methodologies regarding the economic costs of conflict. We use cross-country regressions with fixed effects to derive the cost of ongoing conflict and our results are consistent with comparable subnational studies (Abadie and Gardeazabal 2003; Mueller 2016). However, our framework is not restricted to using these results. If comparable evidence from, say, natural experiments on the cost of conflict becomes available, these can easily be used to simulate growth effects. Thus, the flexible structure of our model allows research to rely on the best available estimates.

Bibliography

Abadie, Alberto, and Javier Gardeazabal. 2003. "The economic costs of conflict: A case study of the Basque Country." *American Economic Review* 93 (1): 113–132.

Aguiar, Mark, and Gita Gopinath. 2007. "Emerging market business cycles: The cycle is the trend." *Journal of Political Economy* 115 (1): 69–102.

Akresh, Richard, Sonia Bhalotra, Marinella Leone, and Una Okonkwo Osili. 2012. "War and Stature: Growing Up during the Nigerian Civil War." *American Economic Review* 102 (3): 273–277.

Barro, Robert J., and Xavier Sala-i-Martin. 2004. *Economic Growth.* Second. Cambridge, Massachusetts and London, England: The MIT Press. ISBN: 978-0262025539.

Bazzi, Samuel, and Christopher Blattman. 2014. "Economic shocks and conflict: Evidence from commodity prices." *American Economic Journal: Macroeconomics* 6 (4): 1–38.

Berg, Andrew, Jonathan D. Ostry, and Jeromin Zettelmeyer. 2012. "What makes growth sustained?" *Journal of Development Economics* 98 (2): 149–166. ISSN: 0304-3878.

Berman, Nicolas, and Mathieu Couttenier. 2015. "External Shocks, Internal Shots: The Geography of Civil Conflicts." *The Review of Economics and Statistics* 97 (4): 758–776. Accessed June 10, 2024.

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

Besley, Timothy, Thiemo Fetzer, and Hannes Mueller. 2024. "How Big Is the Media Multiplier? Evidence from Dyadic News Data." *The Review of Economics and Statistics* (February): 1–45.

Besley, Timothy, and Hannes Mueller. 2012. "Estimating the peace dividend: The impact of violence on house prices in Northern Ireland." *American Economic Review* 102 (2): 810–833.

Besley, Timothy, and Torsten Persson. 2011a. *Pillars of Prosperity: The Political Economics of Development Clusters.* Princeton University Press.

———. 2011b. "The logic of political violence." *The Quarterly Journal of Economics* 126 (3): 1411–1445.

Blattman, Christopher, and Jeannie Annan. 2010. "The Consequences of Child Soldiering." *The Review of Economics and Statistics* 92 (4): 882–898.

Callaway, Brantly, and Pedro HC Sant'Anna. 2021. "Difference-in-differences with multiple time periods." *Journal of Econometrics* 225 (2): 200–230.

Cassar, Alessandra, Pauline A Grosjean, and Sam Whitt. 2011. "Civil war, social capital and market development: Experimental and survey evidence on the negative consequences of violence." *Social Capital and Market Development: Experimental and Survey Evidence on the Negative Consequences of Violence (November 10, 2011). UNSW Australian School of Business Research Paper,* no. 2011ECON14.

Cerra, Valerie, and Sweta Chaman Saxena. 2008. "Growth dynamics: the myth of economic recovery." *American Economic Review* 98 (1): 439–457.

Collier, Paul. 1999. "On the economic consequences of civil war." *Oxford Economic Papers* 51 (1): 168–183.

———. 2007. *The Bottom Billion: Why the Poorest Countries are Failing and What Can Be Done About It.* Oxford University Press.

Collier, Paul, V. L. Elliott, Håvard Hegre, Anke Hoeffler, Marta Reynal-Querol, and Nicholas Sambanis. 2003. *Breaking the conflict trap: Civil war and development policy.* Vol. 41181. 4. World Bank Publications.

Corral, Paul, Alexander Irwin, Nandini Krishnan, and Daniel Gerszon Mahler. 2020. *Fragility and conflict: On the front lines of the fight against poverty.* World Bank Publications.

Couttenier, Mathieu, Nathan Monnet, and Luca Piemontese. 2022. "The Economic Costs of Conflict: A Production Network Approach." CEPR Press, Paris & London, *CEPR Discussion Paper,* no. 16984.

Davis, Donald R., and David E. Weinstein. 2002. "Bones, Bombs, and Break Points: The Geography of Economic Activity." *American Economic Review* 92 (5): 1269–1289.

De Chaisemartin, Clément, and Xavier d'Haultfoeuille. 2020. "Two-way fixed effects estimators with heterogeneous treatment effects." *American Economic Review* 110 (9): 2964–2996.

———. 2023. "Two-way fixed effects and differences-in-differences with heterogeneous treatment effects: A survey." *The Econometrics Journal* 26 (3): C1–C30.

De Groot, Olaf J, Carlos Bozzoli, Anousheh Alamir, and Tilman Brück. 2022. "The global economic burden of violent conflict." *Journal of Peace Research* 59 (2): 259–276.

Dube, Oeindrila, and Juan F Vargas. 2013. "Commodity price shocks and civil conflict: Evidence from Colombia." *Review of Economic studies* 80 (4): 1384–1421.

Easterly, William, Michael Kremer, Lant Pritchett, and Lawrence H Summers. 1993. "Good policy or good luck?" *Journal of Monetary Economics* 32 (3): 459–483.

Esteban, Joan-Maria, and Debraj Ray. 1994. "On the measurement of polarization." *Econometrica: Journal of the Econometric Society,* 819–851.

Federle, Jonathan, André Meier, Gernot J. Müller, Willi Mutschler, and Moritz Schularick. 2024. "The price of war." (Kiel), no. 2262.

Garcia-Cicco, Javier, Roberto Pancrazi, and Martin Uribe. 2010. "Real business cycles in emerging countries?" *American Economic Review* 100 (5): 2510–2531.

Guiso, Luigi, Paola Sapienza, and Luigi Zingales. 2009. "Cultural biases in economic exchange?" *The Quarterly Journal of Economics* 124 (3): 1095–1131.

Hegre, Håvard, Joakim Karlsen, Håvard Mokleiv Nygård, Håvard Strand, and Henrik Urdal. 2013. "Predicting armed conflict, 2010–2050." *International Studies Quarterly* 57 (2): 250–270.

Hegre, Håvard, Håvard Mokleiv Nygård, and Ranveig Flaten Ræder. 2017. "Evaluating the scope and intensity of the conflict trap: A dynamic simulation approach." *Journal of Peace Research* 54 (2): 243–261.

Ichino, Andrea, and Rudolf Winter-Ebmer. 2004. "The Long-Run Educational Cost of World War II." *Journal of Labor Economics* 22 (1): 57–87.

Korovkin, Vasily, Alexey Makarin, and Yuhei Miyauchi. 2024. "Supply Chain Disruption and Reorganization: Theory and Evidence from Ukraine's War." *MIT Sloan Research Paper No. 7068-24.*

Leon, Gianmarco. 2012. "Civil Conflict and Human Capital Accumulation: The Long-Term Effects of Political Violence in Perú." *Journal of Human Resources* 47 (4): 991–1022.

Michalopoulos, Stelios, and Elias Papaioannou. 2016. "The long-run effects of the scramble for Africa." *American Economic Review* 106 (7): 1802–1848.

Miguel, Edward, and Gerard Roland. 2011. "The Long-Run Impact of Bombing Vietnam." *Journal of Development Economics* 96 (1): 1–15.

Montalvo, José G, and Marta Reynal-Querol. 2005. "Ethnic polarization, potential conflict, and civil wars." *American Economic Review* 95 (3): 796–816. Data retrieved from http://www.econ.upf.edu/~reynal/data_web.htm.

Mueller, Hannes. 2012. "Growth dynamics: The myth of economic recovery: Comment." *American Economic Review* 102 (7): 3774–3777.

———. 2016. "Growth and violence: argument for a per capita measure of civil war." *Economica* 83 (331): 473–497.

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

Mueller, Hannes, and Julia Tobias. 2016. "The cost of violence: Estimating the economic impact of conflict." *International Growth Centre.*

North, Douglass Cecil, John Joseph Wallis, and Barry R Weingast. 2009. *Violence and social orders: A conceptual framework for interpreting recorded human history.* Cambridge University Press.

Nunn, Nathan, and Diego Puga. 2012. "Ruggedness: The blessing of bad geography in Africa." *Review of Economics and Statistics* 94 (1): 20–36. Data retrieved from http://diegopuga.org/data/rugged/.

Organski, A.F.K., and Jacek Kugler. 1977. "The Costs of Major Wars: The Phoenix Factor." *American Political Science Review* 71 (4): 1347–1366.

Pinotti, Paolo. 2015. "The Economic Costs of Organised Crime: Evidence from Southern Italy." *The Economic Journal* 125 (586): F203–F232.

Ramey, Garey, and Valerie A. Ramey. 1995. "Cross-Country Evidence on the Link Between Volatility and Growth." *The American Economic Review* 85 (5): 1138–1151. Accessed June 27, 2024.

Robinson, James A, and Daron Acemoglu. 2012. *Why nations fail: The origins of power, prosperity and poverty.* Profile London.

Rodrik, Dani. 1999. "Where did all the growth go? External shocks, social conflict, and growth collapses." *Journal of Economic Growth* 4 (4): 385–412.

Rohner, Dominic. 2011. "Reputation, group structure and social tensions." *Journal of Development Economics* 96 (2): 188–199.

Rohner, Dominic, and Mathias Thoenig. 2021. "The elusive peace dividend of development policy: From war traps to macro complementarities." *Annual Review of Economics* 13:111–131.

Rohner, Dominic, Mathias Thoenig, and Fabrizio Zilibotti. 2013. "War signals: A theory of trade, trust, and conflict." *Review of Economic Studies* 80 (3): 1114–1147.

Steven Brakman, Harry Garretsen, and Marc Schramm. 2004. "The Strategic Bombing of German Cities during World War II and its Impact on City Growth." *Journal of Economic Geography* 4 (2): 201–218.

Sun, Liyang, and Sarah Abraham. 2021. "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects." *Journal of Econometrics* 225 (2): 175–199.

Tapsoba, Augustin. 2023. "The cost of fear: Impact of violence risk on child health during conflict." *Journal of Development Economics* 160:102975.

World Bank. 2020. *World Bank Group Strategy for Fragility, Conflict, and Violence 2020–2025.* World Bank, Washington, DC.

Chapter 3

Coordination across Tax Havens: A Global Games Approach

3.1 Introduction

Tax evasion has become a major concern in recent decades, as it undermines public revenues and exacerbates inequality (OECD 2024). A key channel, particularly among top-income earners, is using tax havens—jurisdictions that offer low tax rates and strong financial secrecy.[1] Estimates suggest that approximately 8% of global household financial wealth—equivalent to around 10% of world GDP—is held in tax havens (Zucman 2013, 2014; Alstadsæter et al. 2018).[2]

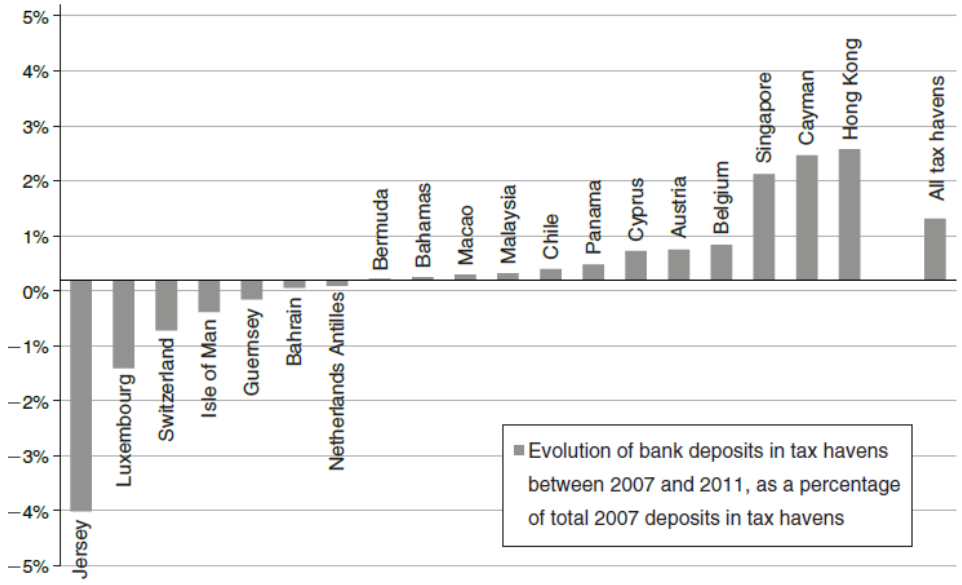

Figure 3.1: Evolution of foreign-owned deposits in tax havens.

Notes: The figure shows the evolution of foreign-owned deposits in each BIS-reporting tax haven. The comparison is based on the average deposits from the first semester of 2011 and 2007 (except for Cyprus, which began reporting in 2008:IV, and Malaysia, which started in 2007:IV). The difference is expressed as a fraction of the total deposits held in all tax havens in 2007 (2.6 trillion dollars).
Source: Bank for International Settlements (2002–2011), restricted bilateral locational banking statistics (Johannesen and Zucman 2014).

This has prompted supranational organizations such as the OECD and G20 to implement policies to curb these practices. Major initiatives were launched in 2008–2009, which pressured tax havens to sign a series of bilateral treaties of information exchange. However, overall evasion grew despite the international efforts (Johannesen and Zucman 2014). Figure 3.1 shows the evolution of foreign-owned before and after the implementation of the policies. While some jurisdictions experienced declines, others saw increases. On net, deposits held in tax havens continued to grow. The authors find that net deposits tend to decrease as more treaties are signed by each haven, suggesting that activity moved toward less compliant jurisdictions. This outcome reveals a tension between local improvements and global effectiveness, driven by a strategic relocation across jurisdictions.[3]

1. A tax haven is a jurisdiction that attracts foreign wealth through a combination of low or zero tax rates and limited cooperation in sharing information about foreign asset holders.

2. Zucman (2014) estimates $7.6 trillion in offshore wealth in 2013, corresponding to roughly $190 billion in annual tax revenue losses. Other estimates are higher: Boston Consulting Group (2014) reports $8.9 trillion for the same year, while Henry (2012) suggests the total could reach as much as $32 trillion.

3. The coexistence of multiple tax havens with different levels of compliance has long been recognized as a challenge

This paper develops a game-theoretical model to study tax evasion across multiple tax havens, capturing the strategic coordination effects that arise across them. The idea of tax evasion as a coordination game is well established in the literature (Bucovetsky 2014; Konrad and Stolper 2016). By lowering its tax rate, a small country may lose little revenue from its domestic base while gaining from attracting foreign taxable wealth. Therefore, the tax haven mechanism for evading is only available if a sufficient number of investors choose to evade.[4] The model builds on the concept of global games, a class of coordination games with incomplete information.[5] The global game structure emphasizes how agents' beliefs, shaped by expectations about others' actions, determine equilibrium outcomes. I then analyze how different policies that raise the cost of being a tax haven affect evader behavior and overall levels of evasion.

The model is a two-stage sequential game with homogeneous investors who can choose to pay taxes in a high-tax country or attempt to evade through one of two tax havens. The regime—whether a jurisdiction remains a tax haven or not—depends on the number of evaders who choose to operate through it. Agents face incomplete information about the number required to sustain each tax haven. In the first stage, investors decide whether to specialize in one tax haven or the other. Specialization can be interpreted as completing all necessary steps to transfer funds into the chosen jurisdiction. It involves no direct cost, but choosing one haven implies losing the opportunity to move to the other. This decision is based on a common public signal. In the second stage, investors receive an additional private signal and choose whether to evade or comply. If enough investors choose to evade, the tax haven remains viable, and evaders benefit relative to compliance. If not enough investors evade, the haven falls, and evaders are reported to their home authorities, facing a punishment greater than if they had complied.

In a partial equilibrium setting, I analyze how different types of policies affect overall evasion. Interventions are modeled as changes to the public signal, resembling economic sanctions that raise the cost of sustaining a tax haven, thereby requiring more evaders. I then conduct a comparative statics exercise to study the effects of targeting one or both havens.

The model captures the strategic interaction between tax havens, making policy outcomes dependent not only on the absolute safety of a haven but also on its relative attractiveness. When one haven is relatively more attractive, investors tend to concentrate there, increasing the likelihood of sustaining it. However, when havens are similar, investors split more evenly, reducing the likelihood of sustaining them. As a result, a policy targeting a single tax haven reduces evasion only if it affects the most attractive one; otherwise, it may even increase it. In contrast, equal punishment to all havens always reduces evasion. This mechanism helps explain the limited effectiveness of the OECD initiatives and provides a rationale for more homogeneous policies, such as the Global Minimum Tax.

This paper relates to at least two strands of literature. The first concerns tax evasion, specifically the literature on tax havens. Most existing models do not treat investors' decisions as a coordination problem with endogenous beliefs. Elsayyad and Konrad (2012) analyze a sequential game in which the OECD offers compensation or punishment to tax havens in exchange for ceasing their activities. They find that offers should be made simultaneously rather than sequentially, as remaining tax havens become

for policy effectiveness. Hines Jr (2005) already estimated that two-thirds of U.S. multinational investment was located in jurisdictions not classified as tax havens by the OECD's 1998 report.

4. For foundational models of tax competition, see the ZMW model (Zodrow and Mieszkowski 1986; Wilson 1986) and the KK model (Kanbur and Keen 1991), as summarized in Keen and Konrad (2013).

5. Global games were introduced by Carlsson and Van Damme (1993) which often leads to a unique, iterative dominant equilibrium. See Morris and Shin (2001) for a complete explanation.

stronger and more costly to deter. A similar result appears in Slemrod and Wilson (2009), which models tax havens as juridical entrepreneurs selling protection from national taxation. The paper most closely related to this one is Konrad and Stolper (2016), which presents a global game in which investors do not know the fixed cost of a tax haven. While they extend the model to multiple havens, each is treated independently, thereby ruling out coordination effects across tax havens.

The second is the global games literature, which has been used to model phenomena such as speculative attacks on currency pegs (Morris and Shin 1998), bank runs (Goldstein and Pauzner 2005), and revolutions against governments (Angeletos et al. 2007). This model belongs to the class of regime change global games, where payoffs change at a threshold rather than continuously. However, these models have focused on a single entity—one currency, one bank, or one government. In my setting, multiple global games happen simultaneously, and agents, at an ex-ante stage, choose which one they participate in. The first stage captures the broader coordination problem across entities, while the second stage captures the classic coordination problem within the entity.

The contribution of this paper lies in extending global games to a multi-entity setting, enabling the analysis of coordination effects across entities and introducing a layer of strategic interaction that has not been previously studied. This framework proves useful for understanding tax evasion through tax havens. The model generates equilibrium outcomes that would not arise—or would be reversed—in single-haven settings.

This paper is structured as follows. Section 3.2 presents the theoretical model. In Section 3.3, I perform a comparative statics exercise to analyze how different policies affect overall evasion. Section 3.4 extends the analysis to a setting with heterogeneous agents. In Section 3.5, I discuss the policy implications. The paper concludes with a summary of the results.

3.2 Model

The model is a two-stage sequential game. In the first stage, investors choose which of the two tax havens to specialize in. In the second stage, they decide whether to evade through the haven they selected or comply and pay taxes in their home country.

There is a continuum of homogeneous investors, indexed by $i \in I$, with total mass normalized to one. Each investor owns one unit of mobile capital. The home country taxes this capital at a rate $t \in [0, 1]$, while both tax havens apply a rate $p \in [0, 1]$, with $p < t$. The two tax havens are indexed by $j \in \{1, 2\}$ and referred to as TH_1 and TH_2.[6]

There is incomplete information about the number of investors required to sustain each tax haven, which is represented by θ_j. θ_j can be interpreted as the economic cost of being a tax haven due to, for example, the loss of domestic revenue and international sanctions. Agents do not know the true value of θ_j but have a common prior belief about it, which follows an independent normal distribution $\mathcal{N}(\mu_j, \sigma_\mu)$. Tax havens differ in their μ_j, representing the public signal.

In the first stage, each investor chooses which tax haven to specialize in, enabling evasion through it at a later point. This represents an "access stage", capturing jurisdiction-specific requirements for evasion, such as identifying legal loopholes, obtaining residence or citizenship, or setting up shell com-

6. The assumption of identical tax rates simplifies the analysis, allowing the model to focus on the difference of public signals. In practice, tax haven rates are all close to zero.

panies. Once an investor specializes in a given haven, they cannot use the other. I assume that the cost of specialization is negligible compared to the potential gains from evasion, so it is always profitable to specialize in one of the tax havens before deciding whether to evade or not. Denote the specialization decision of each agent as $s_i = \{1, 2\}$. The proportion of agents specialized in each tax haven is denoted by S_j, s.t. $S_1 + S_2 = 1$.

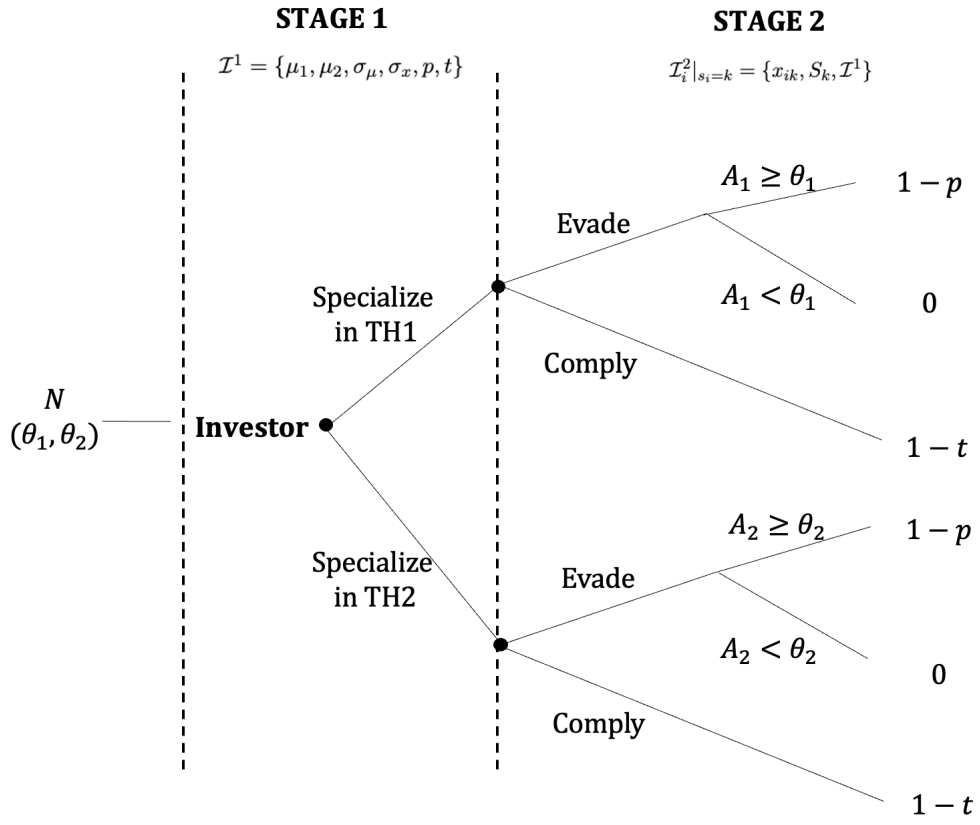

Figure 3.2: Structure of the game

Notes: The figure illustrates the sequential structure of the game. In the first stage, investors receive public information and choose which tax haven to specialize in. In the second stage, they receive additional private information about their selected haven and decide whether to evade taxes or comply. If the proportion of evaders A_j in a given haven is greater than or equal to θ_j, the haven remains active and taxes evaders at rate p. Otherwise, it collapses and reports evaders to their home country, where they are penalized. Investors who choose to comply are taxed at the standard rate t.
Source: Author's elaboration.

By specializing in a specific tax haven, investors also gain additional information about the corresponding θ_j. Each receives a private signal $x_{ij} = \theta_j + \varepsilon_i$, where $\varepsilon_i \sim \mathcal{N}(0, \sigma_x)$. They also observe the proportion of agents specialized in their chosen haven. Using this information, each agent decides whether to evade or comply and pay taxes in their home country. Let $a_i \in \{\text{Evade}, \text{Comply}\}$ denote the decision of agent i. The share of agents who choose to evade in each tax haven is denoted by A_j, with $A_j \leq S_j$ by construction.

If $A_j \geq \theta_j$, the tax haven country remains as a tax haven and taxes evaders according to p. However, if $A_j < \theta_j$, the tax haven's regime changes and reports evaders to their home country, which punishes evaders by expropriating all the capital. Thus, the payoff from compliance is $1 - t$, while the payoff from evasion is $1 - p$ if the haven survives, and 0 if it does not. The higher the value of θ_j, the more

agents are required to sustain the haven. As a result, agents' incentives to evade decrease in both signals and increase in the number of evaders. Their actions are therefore strategic complements.

I denote the information set of each agent by $\mathcal{I}_i^{\xi}$, where $\xi \in \{1, 2\}$ indicates the stage of the game. In the first stage, all agents have the same public information, so the information set is $\mathcal{I}^1 = \{\mu_1, \mu_2, \sigma_\mu, \sigma_x, p, t\}$. In the second stage, they accumulate information on the number of specialized agents and receive private signals, making the information set differ across agents: $\mathcal{I}_i^2|_{s_i=k} = \{x_{ik}, S_k, \mathcal{I}^1\}$. Figure 3.2 shows the structure of the game with the two stages and their different information sets.

The policies implemented by international organizations, such as economic sanctions, increase the cost of being a tax haven, raising the number of agents required to sustain it. This is captured by an increase in θ_j. Although agents do not observe the exact value of θ_j, the impact of these policies is reflected in the public signal μ_j. The following comparative statics exercise analyzes how changes in the public signals (μ_1, μ_2) affect the number of evaders in each haven, (A_1, A_2).

I solve the model by Backward induction.

3.2.1 Stage 2: evasion decision

The decision problem in the second stage takes the form of a global game. Consider the case of an agent who has chosen TH_k.

If the agent chooses to comply, her payoff is $1 - t$. If she chooses to evade, the payoff depends on the aggregate behavior of others: it is $1 - p$ if the selected tax haven survives ($A_k \geq \theta_k$), and 0 if it collapses ($A_k < \theta_k$). Table 3.1 presents the corresponding payoff structure.

	TH_k Remains ($A_k \geq \theta_k$)	TH_k Falls ($A_k < \theta_k$)
Evade	$1 - p$	0
Comply	$1 - t$	$1 - t$

Table 3.1: Stage 2 payoff table $\forall i : s_i = k$.

Notes: The table presents the payoffs of an agent who has chosen to specialize in TH_k during the second stage. θ_k is the minimum number of evaders required to sustain the haven's secrecy regime, while A_k is the actual number of evaders. The regime survives if $A_k \geq \theta_k$ and collapses otherwise. p is the tax rate applied by the haven when it survives, and t is the home-country tax rate applied in case of compliance.
Source: Author's elaboration.

An investor finds optimal to evade if

$$(1 - p) \Pr(A_k \geq \theta_k) \geq 1 - t. \tag{3.1}$$

Each agent forms a belief about the state of the world θ_k using the available signals. Since the risk of being reported increases with the signals, it is strictly dominant to evade when they are sufficiently low. According to this, suppose that they will adopt a *switching strategy* $a_k(x_{ik})$ based on a threshold value

of the private signal $\hat{x}_k$, s.t.

$$a_k(x_{ik}) = \begin{cases} \text{Evade} & \text{if } x_{ik} \leq \hat{x}_k \\ \text{Comply} & \text{if } x_{ik} > \hat{x}_k. \end{cases}$$

Given a threshold value $\hat{x}_k$, the probability of evading corresponds to the proportion of investors specialized in TH_k who decide to evade, that is,

$$\Pr(x_{ik} \leq \hat{x}_k|\theta) = \Phi\left(\frac{\hat{x}_k - \theta_k}{\sigma_x}\right) = \frac{A_k(\theta_k)}{S_k} \tag{3.2}$$

where Φ is the cumulative distribution function (CDF) of the standard normal. Since $A_k(\theta_k)$ decreases with θ_k, there exist a unique state of θ_k, say $\hat{\theta}_k$, that is equal to $A_k(\hat{\theta}_k)$. Using this fact, TH_k will survive if

$$\theta_k \leq \hat{\theta}_k = S_k \cdot \Phi\left(\frac{\hat{x}_k - \hat{\theta}_k}{\sigma_x}\right) \tag{3.3}$$

which characterizes a fixed point.

Given the value of the signals, investors can update their beliefs about θ_k. By Bayesian updating, the posterior belief about θ_k conditional on the signals is normal with mean $(\sigma_x^2\mu_k + \sigma_\mu^2 x_{ik})/(\sigma_x^2 + \sigma_\mu^2)$ and variance $(\sigma_x^2\sigma_\mu^2)/(\sigma_x^2 + \sigma_\mu^2)$. Then, considering Equation 3.3, the posterior probability of a tax haven surviving is

$$\Pr(\theta_k \leq \hat{\theta}_k|\mathcal{I}_i^2|_{s_i=k}) = \Phi\left(\frac{\hat{\theta}_k - \frac{\sigma_x^2\mu_k + \sigma_\mu^2 x_{ik}}{\sigma_x^2+\sigma_\mu^2}}{\sqrt{\frac{\sigma_x^2\sigma_\mu^2}{\sigma_x^2+\sigma_\mu^2}}}\right). \tag{3.4}$$

Remember that the incentives to evade decrease as the private signal increases. When the private signal is exactly at the threshold $\hat{x}_k$, an agent, in equilibrium, should be indifferent between evading and complying. Therefore, $\hat{x}_k$ can be pinned down as the value of the private signal that satisfies

$$(1-p)\Phi\left(\frac{\hat{\theta}_k - \frac{\sigma_x^2\mu_k + \sigma_\mu^2 \hat{x}_k}{\sigma_x^2+\sigma_\mu^2}}{\sqrt{\frac{\sigma_x^2\sigma_\mu^2}{\sigma_x^2+\sigma_\mu^2}}}\right) = 1 - t. \tag{3.5}$$

Using some algebra, the equilibrium threshold can be expressed as

$$\hat{x}_k^* = \alpha\hat{\theta}_k - \beta\Phi^{-1}\left(\frac{1-t}{1-p}\right) - (\alpha-1)\mu_k \tag{3.6}$$

where $\alpha = \frac{\sigma_x^2+\sigma_\mu^2}{\sigma_\mu^2}$ and $\beta = \frac{\sigma_x}{\sigma_\mu}\sqrt{\sigma_x^2+\sigma_\mu^2}$.

A monotone equilibrium $\hat{x}_k^*$ is thus identified by solving the system of equations formed by Equation 3.3 and Equation 3.6.

Proposition 3.1 *A Bayesian NE for each tax haven regime change global game $(\hat{x}_k^*)$ exists and is unique if and only if $\sigma_\mu^2 > \sigma_x/\sqrt{2\pi}$.*

The proof is in Appendix C.1.

Given that the equilibrium is defined by fixed points, the model needs to be solved computationally. However, we can state the relation between $\hat{x}_k^*$ and some parameters.

Corollary 3.1 *The equilibrium threshold $\hat{x}_k^*$ satisfies the following properties:*

- *$\hat{x}_k^*$ and $\hat{\theta}_k$ are complements.*
- *$\hat{x}_k^*$ is increasing in S_k.*
- *$\hat{x}_k^*$ is increasing in t but decreasing in p.*
- *$\hat{x}_k^*$ is decreasing in μ_k.*

The proof is in Appendix C.1. The equilibrium threshold $\hat{x}_k^*$ determines investors' incentives to evade; a higher threshold increases the likelihood of evasion, as it expands the set of private signals (x_{ik}) for which evasion is optimal. The first point reflects complementarity: a higher likelihood of evasion requires a higher state $\hat{\theta}_k$ for the tax haven to undergo a regime change, and conversely, a higher regime-change threshold encourages more evasion. The second point is intuitive: as more investors specialize in a tax haven, the mass of potential evaders increases, thus raising the likelihood of evasion. The third point captures comparative statics related to tax rates: a higher tax rate in the high-tax country (t) makes evasion more attractive, while higher tax haven rates (p) reduce incentives to evade. Finally, an increase in the public signal (μ_k) reduces the incentives to evade since it implies a higher expected value of θ_k.

3.2.2 Stage 1: specialization decision

In this stage, agents choose which tax haven to specialize in. The outcome is a specialization distribution (S_1^*, S_2^*) that is consistent with agents' optimal responses, based on their inference about the second stage. At the time of the decision, agents do not know the specialization proportions—which are determined endogenously—nor their private signals x_{ij}, which are revealed later. All agents share the same information set $\mathcal{I}^1 = \{\mu_1, \mu_2, \sigma_\mu, \sigma_x, p, t\}$. Then, the belief about θ_j conditional on the information set is normally distributed with mean μ_j and variance σ_μ^2, while for the signal x_{ij} is normally distributed with mean μ_j and variance $\sigma_\mu^2 + \sigma_x^2$.

Denote the expected payoff of specializing in TH_j as Π_j. An agent i will specialize in TH_1 if

$$\Pi_{i1}(\hat{x}_1^*, \hat{\theta}_1 | \mathcal{I}^1) \geq \Pi_{i2}(\hat{x}_2^*, \hat{\theta}_2 | \mathcal{I}^1). \tag{3.7}$$

with

$$\Pi_{ij}(\hat{x}_j^*, \hat{\theta}_j | \mathcal{I}^1) = (1-p)\Pr(x_{ij} \leq \hat{x}_j^* \ \cap \ \theta_j \leq \hat{\theta}_j | \mathcal{I}^1) + (1-t)\Pr(x_{ij} > \hat{x}_j^* | \mathcal{I}^1). \tag{3.8}$$

The specialization decision depends on the thresholds $(\hat{x}_j^*, \hat{\theta}_j)$, which determine the likelihood of successfully evading. These thresholds are functions of the tax haven parameters, as defined in Equations 3.3 and 3.6. The only factors differentiating expected payoffs are agents' beliefs about the specialization proportions (S_1, S_2) and the public signals (μ_1, μ_2).

Lemma 3.1 *The expected payoff of specializing in a TH_j (Π_{ij}) increases with the number of agents who specialize in it (S_j) and decreases in its public signal (μ_j) if*

$$(1-p)\frac{\sqrt{\sigma_x^2+\sigma_\mu^2}}{\sqrt{2\pi}\sigma_\mu\sigma_x} e^{\left(-\frac{1}{2(1-\rho^2)}\right)} \int_{-\infty}^{\hat{\theta}_j(S_j,\mu_j)} e^{\left(-\frac{1}{2(1-\rho^2)}\left(-\frac{2\rho(\hat{x}_j(S_j,\mu_j)-\mu_j)(\theta_j-\mu_j)}{\sigma_\mu\sqrt{\sigma_x^2+\sigma_\mu^2}}+\left(\frac{\theta_j-\mu_j}{\sigma_\mu}\right)^2\right)\right)} d\theta_j - (1-t) > 0 \tag{3.9}$$

The proof is in Appendix C.1. The inequality formalizes a sufficient condition under which the strategic complementarity of agents' choices holds. As S_j increases, the thresholds $\hat{x}_j$ and $\hat{\theta}_j$ also rise making evasion more likely since there are more potential evaders (Corollary 3.1). This raises the expected payoff by increasing the likelihood of successful evasion but simultaneously lowers it by reducing the likelihood of compliance. Since these probabilities are not exact complements—one is based on the joint distribution and the other on the marginals—the decrease in the first probability does not necessarily match the decrease in the second. The overall effect of a change in the thresholds on the expected payoff depends on the relative magnitude of the changes in the probabilities weighted by the tax rates. The condition ensures that the evasion payoff dominates. Conversely, an increase in μ_j lowers both thresholds, reducing the likelihood of evasion and thus the expected payoff.

With a continuum of agents, individual decisions have no impact on aggregate specialization proportions. Therefore, an agent's best response is only a function of the specialization distribution and the public signals. Let $br_i(S_1, S_2, \mu_1, \mu_2) = s_i$ denote the best response of agent i.

In equilibrium, the specialization proportions must be consistent with these best responses. Following the global games approach, assume that the strategy s_i is a switching rule that maps public signals into specialization choices, i.e., $s_i(\mu_1, \mu_2) \in \{1, 2\}$.

Definition 3.1 *A strategy profile $\{s_i^*, a_1^*, a_2^*\}_{i \in I}$ constitutes a **Perfect Bayesian Nash Equilibrium (PBNE)** of the two-stage game if:*

- ***(Stage 2 optimality)** For each agent i, given their information set $\mathcal{I}_i^2$, the evasion strategy $a_j(x_{ij})$ is a best response in the global game solved in Stage 2. That is, a_j^* solves the regime change problem with thresholds $(\hat{x}_j^*, \hat{\theta}_j)$.*

- ***(Stage 1 optimality)** For each agent i, given their information set $\mathcal{I}^1$ and the second-stage equilibrium thresholds $(\hat{x}_1^*, \hat{\theta}_1, \hat{x}_2^*, \hat{\theta}_2)$, the specialization strategy s_i is a best response, i.e.,:*

$$s_i^*(\mu_1, \mu_2) = br_i(S_1^*, S_2^*, \mu_1, \mu_2) = \arg \max_{j \in \{1,2\}} \Pi_{ij}(\hat{x}_j^*, \hat{\theta}_j | \mathcal{I}^1)$$

with the equilibrium proportions satisfying:

$$S_j^* = \int_I \mathbf{1}\{s_i^*(\mu_1, \mu_2) = j\}\, di \quad \text{for } j \in \{1, 2\}$$

Before introducing equilibrium strategies, it is useful to understand how public signals influence best responses. Given Lemma 3.1, if μ_1 is sufficiently low compared to μ_2, the expected payoff from specializing in TH_1 dominates for all values of (S_1, S_2). This arises because changes in $\mu_j \in (-\infty, \infty)$ can shift the threshold $\hat{x}_j^*$ across the entire real line, while $S_j \in [0, 1]$ has a much more limited influence on expected payoffs. Conversely, if μ_1 is sufficiently high, specialization in TH_2 dominates.

Accordingly, the set of equilibria is determined by two threshold values, $\underline{\mu_1}(\mu_2)$ and $\overline{\mu_1}(\mu_2)$, such that:

- If $\mu_1 < \underline{\mu_1}$, specializing in TH_1 strictly dominates for all (S_1, S_2), yielding a unique equilibrium where all agents choose TH_1 ($S_1^* = 1$, $S_2^* = 0$).

- If $\mu_1 > \overline{\mu_1}$, specializing in TH_2 strictly dominates for all (S_1, S_2), yielding a unique equilibrium where all agents choose TH_2 ($S_1^* = 0$, $S_2^* = 1$).
- If $\mu_1 \in [\underline{\mu_1}, \overline{\mu_1}]$, the equilibrium outcome depends on (S_1, S_2), generating multiple equilibria. In this case, there exists a unique threshold $\hat{S}_1(\mu_1, \mu_2) \in [0, 1]$ that makes agents indifferent between the two tax havens s.t.:
  - If $S_1 > \hat{S}_1$, then $\Pi_{i1} > \Pi_{i2}$, agents in TH_2 deviate, and the equilibrium converges to $(S_1^* = 1, S_2^* = 0)$.
  - If $S_1 < \hat{S}_1$, then $\Pi_{i1} < \Pi_{i2}$, agents in TH_1 deviate, and the equilibrium becomes $(S_1^* = 0, S_2^* = 1)$.
  - If $S_1 = \hat{S}_1$, then $\Pi_{i1} = \Pi_{i2}$, agents have no incentives to deviate, and the equilibrium is interior: $(S_1^* = \hat{S}_1, S_2^* = 1 - \hat{S}_1)$.

When the difference between the public signals is small, full coordination in the tax haven with the higher signal may still yield a higher expected payoff due to the strategic complementarity in specialization decisions. In other words, the mass of agents choosing the same haven can compensate for the unfavorable public signal. The threshold $\hat{S}_1(\mu_1, \mu_2)$ captures the precise share of agents required in TH_1 to make others indifferent between the two havens. This threshold is decreasing in $\mu_1 - \mu_2$: the more favorable the public signal of TH_1 relative to TH_2, the fewer agents are needed in TH_1 to sustain indifference.

Among all strategies that can support the equilibria described above, we focus on two symmetric switching strategies. The strategies are defined as follows:

- **Strategy 1 (Full coordination)** Agents specialize in the tax haven with the lower public signal. In the event of a tie, they coordinate on one tax haven. Without loss of generality, assume ties are broken in favor of TH_1:

$$s_i(\mu_1, \mu_2) = \begin{cases} \text{Specialize in } TH_1 & \text{if } \mu_1 \leq \mu_2 \\ \text{Specialize in } TH_2 & \text{if } \mu_1 > \mu_2. \end{cases} \tag{3.10}$$

- **Strategy 2 (Fair-Mixing)** Agents specialize in the tax haven with the lower public signal, and mix uniformly when the signals are equal:

$$s_i(\mu_1, \mu_2) = \begin{cases} \text{Specialize in } TH_1 & \text{if } \mu_1 < \mu_2 \\ \text{Specialize in } TH_2 & \text{if } \mu_1 > \mu_2 \\ \text{Mix (0.5, 0.5)} & \text{if } \mu_1 = \mu_2. \end{cases} \tag{3.11}$$

The full coordination strategy ensures coordination on the more favorable tax haven but introduces a discontinuous jump when public signals are equal. Finally, the hybrid strategy follows Strategy 1 when signals differ, but switches to Strategy 2 when they are equal, avoiding arbitrary tie-breaking.

Proposition 3.2 *If $\sigma_\mu^2 > \sigma_x/\sqrt{2\pi}$ and Lemma 3.1 holds, then Strategies 1 and 2 constitute Perfect Bayesian Nash Equilibria.*

The proof is in Appendix C.1. The first condition, established in Proposition 3.1, ensures that agents' second-stage strategies are optimal given the equilibrium thresholds and that the solution is unique. Lemma 1 ensures that expected payoffs are increasing in S_j and decreasing in μ_j, implying strategic complementarities and monotonicity in best responses. As a result, coordination emerges endogenously, and each strategy leads to equilibrium specialization consistent with the public signals.

Proposition 3.3 *Strategy 1 yields the equilibrium with the highest expected payoff for all agents; that is, it constitutes the payoff-dominant equilibrium.*

This result follows directly from Lemma 3.1 and the properties of Strategy 1. Since the expected payoff increases with the number of agents specializing in a given tax haven and decreases with the public signal, the highest payoff of the game is achieved when all agents coordinate on the tax haven with the lower public signal. In the case of a tie, full coordination on either haven still dominates any split. This is exactly what Strategy 1 prescribes. Hence, Strategy 1 attains the highest possible expected payoff, and, being an equilibrium strategy, it constitutes the payoff-dominant equilibrium.

3.3 Evasion analysis

For given switching strategies, the total number of agents who choose to evade in equilibrium is

$$A^* = A_1^* + A_2^* = S_1^* \Pr\left(x_{i1} \leq \hat{x}_1^* | \theta_1\right) + S_2^* \Pr\left(x_{i2} \leq \hat{x}_2^* | \theta_2\right). \tag{3.12}$$

Note that the realization of the private signals and the survival of the tax havens depend on the realizations of θ_1 and θ_2. As a result, A^* changes in each realization of the game.

We can construct an approximation of A^* by making a "first-stage" inference about the distribution of x_{ij}, i.e., using the public signals:

$$A^* \approx \tilde{A} = S_1^* \Phi\left(\frac{\hat{x}_1^* - \mu_1}{\sqrt{\sigma_x^2 + \sigma_\mu^2}}\right) + S_2^* \Phi\left(\frac{\hat{x}_2^* - \mu_2}{\sqrt{\sigma_x^2 + \sigma_\mu^2}}\right) \tag{3.13}$$

The higher σ_x and the lower σ_μ, the more accurate the approximation becomes, as the public signals provide a better estimate of θ_j.

Comparative statics on $\tilde{A}$ are conducted with respect to the public signals. An increase in only one of the signals, say μ_1, affects evasion through multiple channels: (i) it shifts the mean of the distribution of x_{i1} to the right, making it less likely that agents fall below the evasion threshold—thereby reducing evasion; (ii) it raises the evasion threshold $\hat{x}_1^*$ (as shown in Lemma 3.1), further decreasing the likelihood of evasion in TH_1; and (iii) depending on the specialization switching strategy, the increase in μ_1 may reduce S_1^* and increase S_2^*, shifting more agents toward TH_2. These changes feed back into the thresholds: $\hat{x}_1^*$ decreases due to the lower S_1^*, while $\hat{x}_2^*$ increases due to the higher S_2^*. The net effect of this third channel is ambiguous. While evasion in TH_1 becomes less attractive, the reallocation of agents toward TH_2 makes evasion there more appealing, potentially offsetting the initial reduction and even increasing total evasion.

The aggregate effect can be seen formally in this

$$\frac{\partial \tilde{A}}{\partial \mu_1} = \underbrace{S_1^* \frac{\phi_1(\cdot)}{\sqrt{\sigma_x^2+\sigma_\mu^2}} \left(\frac{\partial \hat{x}_1^*}{\partial S_1^*} \frac{\partial S_1^*}{\partial \mu_1} - 1 \right)}_{\leq 0} + \underbrace{S_2^* \frac{\phi_2(\cdot)}{\sqrt{\sigma_x^2+\sigma_\mu^2}} \left(\frac{\partial \hat{x}_2^*}{\partial S_2^*} \frac{\partial S_2^*}{\partial S_1^*} \frac{\partial S_1^*}{\partial \mu_1} \right)}_{\geq 0} + \underbrace{\frac{\partial S_1^*}{\partial \mu_1} \left[\Phi_1(\cdot) - \Phi_2(\cdot) \right]}_{\substack{<0 \text{ if } \mu_1<\mu_2 \\ >0 \text{ if } \mu_1>\mu_2 \\ =0 \text{ if } \mu_1=\mu_2}} \tag{3.14}$$

$\Phi_j(\cdot)$ and $\phi_j(\cdot)$ denote the CDF and PDF corresponding to TH_j. The first term captures the direct effect in TH_1: a higher μ_1 shifts the signal distribution (i) and raises the evasion threshold (ii), both reducing evasion. The second term reflects the reallocation effect (iii)—as agents switch from TH_1 to TH_2, S_2^* increases, potentially lowering $\hat{x}_2^*$ and raising evasion. The third term captures a composition effect: reallocating agents alters total evasion depending on which haven has the higher evasion rate.

When the change in μ_1 does not affect the specialization decision—that is, $\frac{\partial S_1^*}{\partial \mu_1} = 0$—the expression simplifies to:

$$\frac{\partial \tilde{A}}{\partial \mu_1} = -S_1^* \frac{\phi_1(\cdot)}{\sqrt{\sigma_x^2+\sigma_\mu^2}} \leq 0 \tag{3.15}$$

This confirms that, holding specialization fixed, increasing μ_1 always reduces evasion through TH_1.

Proposition 3.4 *Given Strategy 1 and 2, when the public signal of one tax haven increases, the effect on evasion depends on the relative signal levels:*

(i) If the increase occurs in the more favorable haven (i.e., the one with the lower public signal), evasion decreases under both strategies.

(ii) If the increase brings the public signals to equality ($\mu_1 = \mu_2$), Strategy 2—which prescribes mixing—results in lower evasion than full coordination (Strategy 1).

(iii) If the signal increases beyond the point of equality, evasion rises under Strategy 2 due to a discontinuous shift in specialization, while it remains unchanged under Strategy 1.

(iv) If the increase occurs in the less favorable haven, it has no effect on evasion.

The proof is in Appendix C.1. When an increase in a public signal does not induce a shift in the specialization distribution, it can only reduce evasion. The interesting dynamics arise under Strategy 2 when the signals become equal. Relative to full coordination (Strategy 1), splitting the population across tax havens reduces evasion. However, due to the discontinuous nature of Strategy 2, a marginal increase beyond the equality point causes all agents to switch to the previously less favorable haven. This concentration effect reverses the earlier decline and leads to a sharp rise in evasion relative to the mixed case. By contrast, if the increase occurs in the haven with the higher public signal—one that no agents were choosing—then both strategies still prescribe full specialization in the other haven, and the increase has no effect on evasion.

Proposition 3.5 *An equal increase in both public signals leads to a monotonic decrease in evasion under both strategies.*

This result follows from the fact that when both public signals increase simultaneously while maintaining a constant difference (e.g., $\mu_2 = \mu_1 + c$), the specialization choices of agents remain unchanged. Consequently, the third channel—feedback from switching in specialization—is neutralized. Evasion decreases purely due to the direct effects: the shift in the distribution of private signals and the increase in the evasion thresholds, both of which reduce the probability of evading taxes.

Figure 3.3 illustrates how changes in public signals affect equilibrium evasion outcomes under the two specialization strategies, showcasing the results from Propositions 3.4 and 3.5. Panel (a) considers the case where μ_2 is fixed and μ_1 increases. Panel (b) shows the case where both μ_1 and μ_2 increase equally.

(a) Increasing μ_1 (b) Increasing μ_1 and μ_2

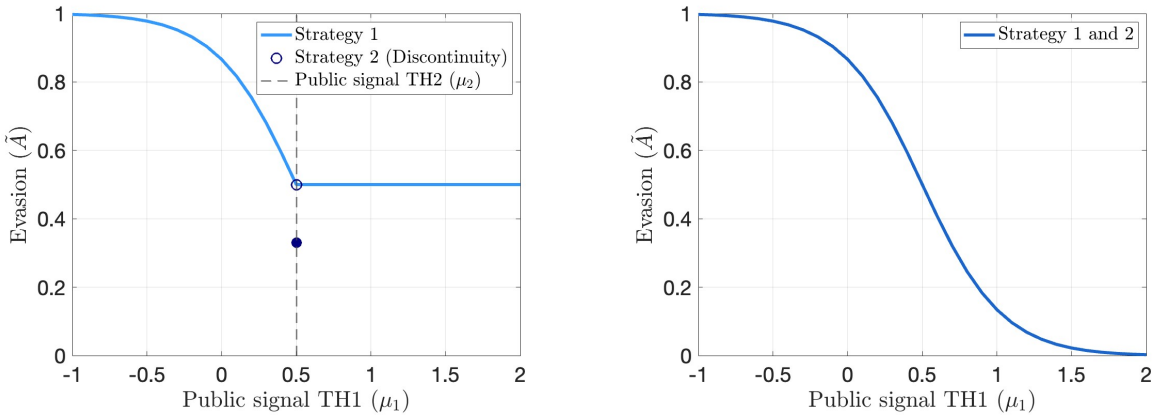

Figure 3.3: Comparative statics on evasion

Notes: The figure shows how equilibrium evasion levels vary under Strategies 1 and 2 for different values of the public signals. Panel (a) analyzes the effect of increasing μ_1 while holding μ_2 fixed at 0.5. Panel (b) considers the case where both public signals increase jointly such that $\mu_2 = \mu_1 + 1.5$. The results are computed under the following parameter values: $(p, t, \sigma_\mu, \sigma_x) = (0, 0.5, 1, 1)$.
Source: Author's elaboration.

3.4 Extension: investors with specialization bias

The main results rely on equilibrium switching strategies that generate sharp transitions in specialization, with discontinuous shifts around signal thresholds. In this section, I explore a more flexible framework that allows for smooth switching behavior by introducing individual biases toward one tax haven. While this extension does not yield a full equilibrium characterization, it provides useful insights into how gradual transitions affect evasion outcomes.

A simple example of a strategy that allows for smooth transitions is given by

$$s_i(\mu_1, \mu_2, \delta_i) = \begin{cases} \text{Specialize in } TH_1 & \text{if } \mu_1 - \delta_i < \mu_2, \\ \text{Mix (0.5, 0.5)} & \text{if } \mu_1 - \delta_i = \mu_2, \\ \text{Specialize in } TH_2 & \text{if } \mu_1 - \delta_i > \mu_2, \end{cases} \tag{3.16}$$

where $\delta_i \sim \mathcal{N}(0, \sigma_\delta^2)$ represents an idiosyncratic bias toward TH_1. By centering the distribution at zero, agents are on average indifferent between the two tax havens. The symmetry of the normal distribution ensures that half the population favors TH_1, while the other half prefers TH_2. The parameter σ_δ captures the extent of heterogeneity: higher values lead to greater dispersion in preferences.

Figure 3.4 presents the comparative statics of the base model and this extension, using alternative distributions of δ ($\sigma_\delta = 0.25$ and $\sigma_\delta = 0.5$). A key insight is that lower volatility in δ (i.e., smaller σ_δ) leads to behavior that more closely mirrors the base model, as specialization shifts more abruptly between tax havens.

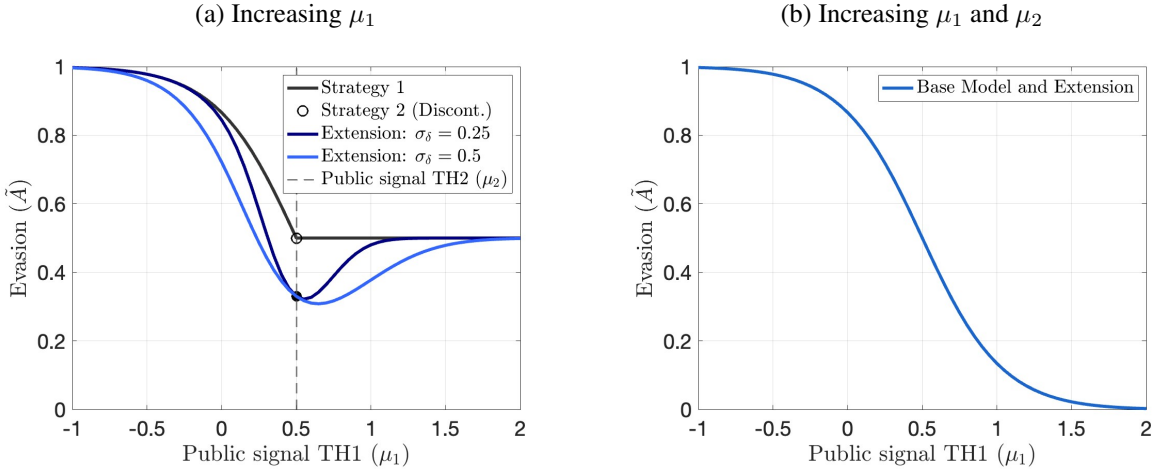

Figure 3.4: Comparative statics on evasion in extension

Notes: The figure illustrates how equilibrium evasion levels respond to changes in the public signals, comparing the base model from the previous section with an extension under alternative distributions ($\sigma_\delta = 0.25$ and $\sigma_\delta = 0.5$). Panel (a) examines the effect of increasing μ_1 while holding μ_2 fixed at 0.5. Panel (b) considers a joint increase in both signals such that $\mu_2 = \mu_1 + 1.5$. Results are computed using the following parameters: $(p, t, \sigma_\mu, \sigma_x) = (0, 0.5, 1, 1)$.
Source: Author's elaboration.

In the extended model, agents begin switching before the signals reach equality, which introduces a smoothing effect that reduces evasion relative to the base model. This reduction occurs gradually, reaching its lowest point precisely at equality—where, due to the symmetry of the distribution, half of the agents specialize in each tax haven, replicating Strategy 2 in the base model.

Interestingly, as the difference between μ_1 and μ_2 increases further, evasion continues to fall, reaching a minimum, and then begins to rise again, eventually converging to the base model values. This pattern suggests that, initially, the policy effect dominates—evasion lost in the taxed haven is not fully compensated by gains in the alternative haven. However, once the gap in public signals becomes large enough, the crowding-in effect takes over, and evasion begins to rise as more agents coordinate on the better tax haven.

3.5 Policy implications

The model reveals a central insight: when coordination problems arise across multiple tax havens, agents respond not only to the absolute conditions of each jurisdiction but also to their relative attractiveness.

If one jurisdiction stands out as more secure, agents concentrate there, making coordination easier and evasion more likely. When havens appear similar, agents are more uncertain about others' choices, coordination weakens, and aggregate evasion falls.

This has direct implications for policy design. A policy targeting a single tax haven reduces evasion only if it undermines the most attractive one; otherwise, it may simply redirect evaders to other jurisdictions, leaving overall evasion unchanged or even increased. By contrast, symmetric enforcement weakens coordination in all of them and reduces evasion.

This mechanism may help explain the limited effectiveness of past initiatives, such as those led by the OECD. Uneven enforcement created relocation rather than deterrence. In this light, the Global Minimum Tax appears more promising, as long as it is implemented uniformly and with broad international participation.

3.6 Conclusions

This paper studies tax evasion through the lens of a global game with multiple tax havens. By modeling evasion as a regime change problem occurring across several tax havens, it captures how investors coordinate both within and across jurisdictions under incomplete information. The key innovation is to allow for cross-haven strategic interactions, which are absent in single-entity models. I then analyze how different types of policies that raise the cost of being a tax haven affect evasion outcomes.

Coordination is easier when tax havens differ in perceived safety, as investors tend to concentrate on the most attractive one. This makes uneven enforcement potentially counterproductive, as it may increase overall evasion by shifting it across jurisdictions. In contrast, applying pressure evenly reduces coordination incentives in all havens and is more effective. This helps explain the limited success of the OECD's 2008–2009 initiatives and highlights the potential of the Global Minimum Tax, provided it is implemented uniformly and with broad coverage.

As a direction for future research, the model could be extended to include more than two tax havens with differing tax rates, or to a general equilibrium setting where both the OECD and tax havens are strategic players. Additionally, since individuals are assumed to be risk-neutral, the crowding-in effect arises from differences in expected value rather than risk reduction. Extending the model to risk-averse investors would offer new insights into how coordination effects operate under different preferences. Finally, the framework could be applied to other coordination settings where individuals choose among multiple entities—such as banks or currencies, as studied in other global game contexts. Ignoring cross-entity interactions in these environments may miss important strategic interactions.

Bibliography

Alstadsæter, Annette, Niels Johannesen, and Gabriel Zucman. 2018. "Who owns the wealth in tax havens? Macro evidence and implications for global inequality." *Journal of Public Economics* 162:89–100.

Angeletos, George-Marios, Christian Hellwig, and Alessandro Pavan. 2007. "Dynamic global games of regime change: Learning, multiplicity, and the timing of attacks." *Econometrica* 75 (3): 711–756.

Boston Consulting Group. 2014. *Global Wealth Report 2014: Riding a Wave of Growth.* Boston, MA: Boston Consulting Group.

Bucovetsky, Sam. 2014. "Honor among tax havens." *Journal of Public Economics* 110:74–81.

Carlsson, Hans, and Eric Van Damme. 1993. "Global games and equilibrium selection." *Econometrica: Journal of the Econometric Society,* 989–1018.

Elsayyad, May, and Kai A Konrad. 2012. "Fighting multiple tax havens." *Journal of International Economics* 86 (2): 295–305.

Goldstein, Itay, and Ady Pauzner. 2005. "Demand–deposit contracts and the probability of bank runs." *The Journal of Finance* 60 (3): 1293–1327.

Henry, James S. 2012. "The price of offshore revisited." *Tax Justice Network* 22:57–168.

Hines Jr, James R. 2005. "Do tax havens flourish?" *Tax Policy and the Economy* 19:65–99.

Johannesen, Niels, and Gabriel Zucman. 2014. "The end of bank secrecy? An evaluation of the G20 tax haven crackdown." *American Economic Journal: Economic Policy* 6 (1): 65–91.

Kanbur, SM Ravi, and Michael Keen. 1991. *Tax competition and tax coordination when countries differ in size.* Vol. 738. World Bank Publications.

Keen, Michael, and Kai A Konrad. 2013. "The theory of international tax competition and coordination." In *Handbook of Public Economics,* 5:257–328. Elsevier.

Konrad, Kai A, and Tim BM Stolper. 2016. "Coordination and the fight against tax havens." *Journal of International Economics* 103:96–107.

Morris, Stephen, and Hyun Song Shin. 1998. "Unique equilibrium in a model of self-fulfilling currency attacks." *American Economic Review,* 587–597.

———. 2001. "Global Games: Theory and Applications." Cowles Foundation Discussion Paper No. 1275R, no. 1275R (August).

OECD. 2024. *Taxation and Inequality: OECD Report to G20 Finance Ministers and Central Bank Governors.* Paris.

Slemrod, Joel, and John D Wilson. 2009. "Tax competition with parasitic tax havens." *Journal of Public Economics* 93 (11-12): 1261–1270.

Wilson, John D. 1986. "A theory of interregional tax competition." *Journal of Urban Economics* 19 (3): 296–315.

Zodrow, George R, and Peter Mieszkowski. 1986. "Pigou, Tiebout, property taxation, and the underprovision of local public goods." *Journal of Urban Economics* 19 (3): 356–370.

Zucman, Gabriel. 2013. "The missing wealth of nations: Are Europe and the US net debtors or net creditors?" *The Quarterly Journal of Economics* 128 (3): 1321–1364.

———. 2014. "Taxing across borders: Tracking personal wealth and corporate profits." *Journal of Economic Perspectives* 28 (4): 121–48.

Appendix A

Appendix for Chapter 1

A.1 Derivations and proofs

A.1.1 Derivation of the price equation

For the 1-periods bond, the price is given by

$$\begin{aligned} Q_{1t} =&\beta\mathbb{E}_t\left[\frac{F_{1t}}{G_{t+1}^{\theta}\Pi_{t+1}}\right]=\beta\mathbb{E}_t\left[\frac{Z_{t+1}}{(\alpha_G G_t^{\phi_G}\varepsilon_{t+1}V_{t+1})^{\theta}\alpha_{\Pi}\Pi_t^{\phi_{\Pi}}\eta_{t+1}W_{t+1}}\right] \\ &=\beta\frac{1}{(\alpha_G G_t^{\phi_G})^{\theta}\alpha_{\Pi}\Pi_t^{\phi_{\Pi}}}\mathbb{E}_t\left[\frac{1}{\varepsilon_{t+1}^{\theta}}\right]\mathbb{E}_t\left[\frac{1}{\eta_{t+1}}\right]\mathbb{E}_t\left[\frac{1}{V_{t+1}^{\theta}W_{t+1}}\right] \\ &=\beta\frac{1}{(\alpha_G G_t^{\phi_G})^{\theta}\alpha_{\Pi}\Pi_t^{\phi_{\Pi}}}\mathbb{E}_t\left[e^{-\log(\theta\varepsilon_{t+1})}\right]\mathbb{E}_t\left[e^{-\log(\eta_{t+1})}\right]\left(1-\delta_{1,t}+\delta_{1,t}\frac{1-\gamma J_F}{J_G^{\theta}J_{\Pi}}\right) \\ &=\beta\frac{e^{\frac{1}{2}\left((\theta\sigma_{\varepsilon})^2+\sigma_{\eta}^2\right)}}{(\alpha_G G_t^{\phi_G})^{\theta}\alpha_{\Pi}\Pi_t^{\phi_{\Pi}}}\left(1+\delta_{1,t}\left(\frac{1-\gamma J_F}{J_G^{\theta}J_{\Pi}}-1\right)\right) \end{aligned} \tag{A.1}$$

For the 2-periods bond, the price is given by

$$\begin{aligned} Q_{2t} =&\beta^2\mathbb{E}_t\left[\frac{1}{G_{t+1}^{\theta}G_{t+2}^{\theta}\Pi_{t+1}\Pi_{t+2}}\right] \\ &=\beta^2\mathbb{E}_t\left[\frac{1}{G_{t+1}^{\theta}(\alpha_G G_{t+1}^{\phi_G}\varepsilon_{t+2}V_{t+2})^{\theta}\Pi_{t+1}(\alpha_{\Pi}\Pi_{t+1}^{\phi_{\Pi}}\eta_{t+2}W_{t+2})}\right] \\ &=\beta^2\frac{1}{(\alpha_G^{2+\phi_G}G_t^{\phi_G+\phi_G^2})^{\theta}\alpha_{\Pi}^{2+\phi_{\Pi}}\Pi_t^{\phi_{\Pi}+\phi_{\Pi}^2}}\mathbb{E}_t\left[\frac{1}{\varepsilon_{t+1}^{(1+\phi_G)\theta}}\right]\mathbb{E}_t\left[\frac{1}{\varepsilon_{t+2}^{\theta}}\right]\mathbb{E}_t\left[\frac{1}{\eta_{t+1}^{1+\phi_{\Pi}}}\right]\mathbb{E}_t\left[\frac{1}{\eta_{t+2}}\right] \\ &\mathbb{E}_t\left[\frac{Z_{t+1}}{V_{t+1}^{\theta}W_{t+1}}\right]\mathbb{E}_t\left[\frac{Z_{t+2}}{V_{t+2}^{\theta}W_{t+2}}\right] \\ &=\beta^2\frac{e^{\frac{1}{2}\left((1+(1+\phi_G)^2)\theta^2\sigma_{\varepsilon}^2+(1+(1+\phi_{\Pi})^2)\sigma_{\eta}^2\right)}}{(\alpha_G^{2+\phi_G}G_t^{\phi_G+\phi_G^2})^{\theta}\alpha_{\Pi}^{2+\phi_{\Pi}}\Pi_t^{\phi_{\Pi}+\phi_{\Pi}^2}}\left(1+\delta_{1,t}\left(\frac{1-\gamma J_F}{J_G^{(1+\phi_G)\theta}J_{\Pi}^{1+\phi_{\Pi}}}-1\right)\right)\left(1+\delta_{2,t}\left(\frac{1-\gamma J_F}{J_G^{\theta}J_{\Pi}}-1\right)\right) \end{aligned} \tag{A.2}$$

For the 3-periods bond, the price is given by

$$\begin{aligned} Q_{3t} =&\beta^3\frac{e^{\frac{1}{2}\left((1+(1+\phi_G)^2+(1+\phi_G+\phi_G^2)^2)\theta^2\sigma_{\varepsilon}^2+(1+(1+\phi_{\Pi})^2+(1+\phi_{\Pi}+\phi_{\Pi}^2)^2)\sigma_{\eta}^2\right)}}{(\alpha_G^{3+2\phi_G+\phi_G^2}G_t^{\phi_G+\phi_G^2+\phi_G^3})^{\theta}\alpha_{\Pi}^{3+2\phi_{\Pi}+\phi_{\Pi}^2}\Pi_t^{\phi_{\Pi}+\phi_{\Pi}^2+\phi_{\Pi}^3}}\left(1+\delta_{1,t}\left(\frac{1-\gamma J_F}{J_G^{(1+\phi_G+\phi_G^2)\theta}J_{\Pi}^{1+\phi_{\Pi}+\phi_{\Pi}^2}}-1\right)\right) \\ &\left(1+\delta_{2,t}\left(\frac{1-\gamma J_F}{J_G^{(1+\phi_G)\theta}J_{\Pi}^{1+\phi_{\Pi}}}-1\right)\right)\left(1+\delta_{3,t}\left(\frac{1-\gamma J_F}{J_G^{\theta}J_{\Pi}}-1\right)\right) \end{aligned} \tag{A.3}$$

Then, for the N-period bond,

$$\begin{aligned} Q_{Nt} =& \beta^N \frac{e^{\frac{1}{2}\left(\sum_{i=1}^{N}(\sum_{j=0}^{i-1}\phi_G^j)^2\theta^2\sigma_\varepsilon^2+\sum_{i=1}^{N}(\sum_{j=0}^{i-1}\phi_\Pi^j)^2\sigma_\eta^2\right)}}{\left(\alpha_G^{\sum_{i=1}^{N} i\phi_G^{N-i}} G_t^{\sum_{i=1}^{N}\phi_G^i}\right)^\theta \alpha_\Pi^{\sum_{i=1}^{N} i\phi_\Pi^{N-i}} \Pi_t^{\sum_{i=1}^{N}\phi_\Pi^i}} \\ & \prod_{\tau=1}^{N}\left(1+\delta_{\tau,t}\left(\frac{1}{J_G^{\sum_{j=1}^{N+1-i}\theta\phi_G^{j-1}} J_\Pi^{\sum_{j=1}^{N+1-i}\phi_\Pi^{j-1}}}-1\right)\right) \end{aligned} \tag{A.4}$$

A.1.2 Proof of Proposition 1.1

I use the logarithmic differentiation trick. Taking the logarithm of Equation 1.11:

$$\log(Q_{Nt}) = \log(Q^{ND}) + \sum_{\tau=1}^{N}\log\left(1+\phi_\delta^{\tau-1}\delta_{1,t}\left(J_{\tau,N}-1\right)\right) \tag{A.5}$$

Then, differentiating with respect to $\delta_{1,t}$,

$$\begin{aligned} \frac{\partial\log(Q_{Nt})}{\partial\delta_{1,t}} =& \sum_{\tau=1}^{N}\frac{\phi_\delta^{\tau-1}\left(J_{\tau,N}-1\right)}{1+\phi_\delta^{\tau-1}\delta_{1,t}\left(J_{\tau,N}-1\right)} \\ \frac{1}{Q_{Nt}}\frac{\partial Q_{Nt}}{\partial\delta_{1,t}} =& \sum_{\tau=1}^{N}\frac{\phi_\delta^{\tau-1}\left(J_{\tau,N}-1\right)}{1+\phi_\delta^{\tau-1}\delta_{1,t}\left(J_{\tau,N}-1\right)} \\ \frac{\partial Q_{Nt}}{\partial\delta_{1,t}} =& Q_{Nt}\sum_{\tau=1}^{N}\frac{\phi_\delta^{\tau-1}\left(J_{\tau,N}-1\right)}{1+\phi_\delta^{\tau-1}\delta_{1,t}\left(J_{\tau,N}-1\right)} \end{aligned} \tag{A.6}$$

Since $Q_{Nt} > 0$, the sign of $\frac{\partial Q_{Nt}}{\partial\delta_{1,t}}$ is determined by the sign of the second element, which proves the first part of the proposition.

Because the denominator in each term is always positive ($1+\phi_\delta^{\tau-1}\delta_{1,t}\left(J_{\tau,N}-1\right) > 0$ and $\phi_\delta > 0$), the sign of each element in the sum depends on the sign of $J_{\tau,N}-1$. The sum includes all periods until maturity, with each term weighted by a positive denominator. Therefore, if all $J_{\tau,N} < 1$, then each term in the sum is negative, and thus the price decreases with an increase in $\delta_{1,t}$.

A.2 Appendix Figures

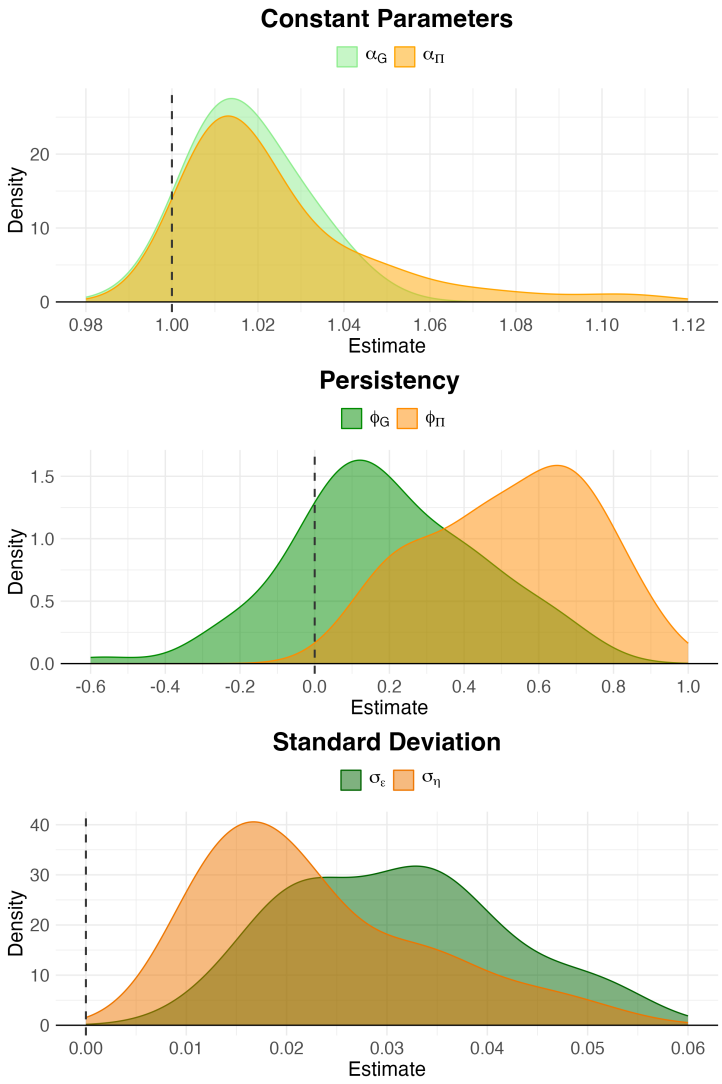

Figure A.1: Distribution of estimates from laws of motion

Notes: Notes: The figure shows the kernel density of the estimates for α_G, α_Π, ϕ_G, ϕ_Π, σ_ε, and σ_η. The density plots for the constant parameters use a bandwidth of 0.01, for the persistency parameters a bandwidth of 0.1, and for the residual standard deviations a bandwidth of 0.005.

Source: Author's calculations.

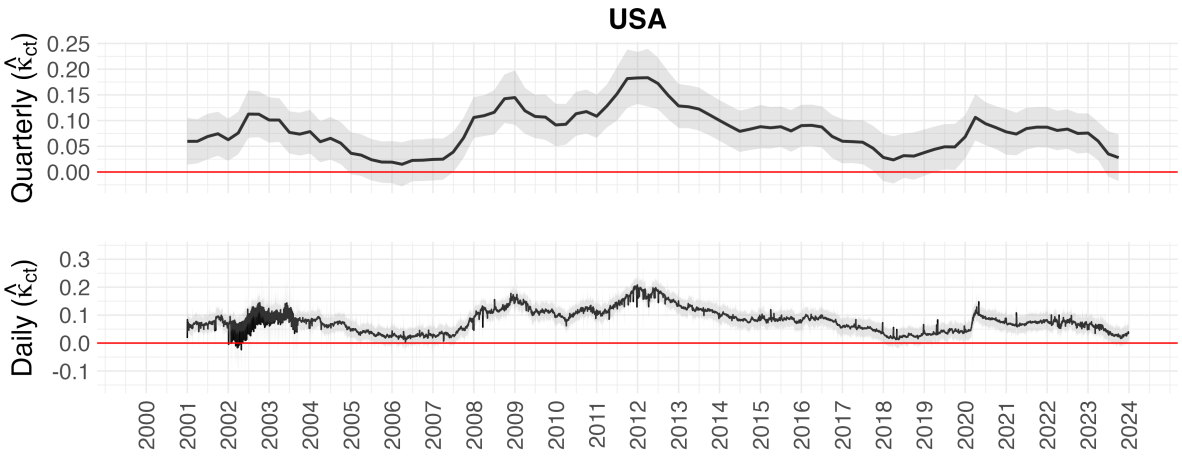

Figure A.2: Comparison between daily and quarterly model

Notes: The figure shows the evolution of $\hat{\kappa}_{ct}$ with 95% confidence intervals for the US using daily and quarterly data. The line at 0 represents a reference threshold.

Source: Author's calculations.

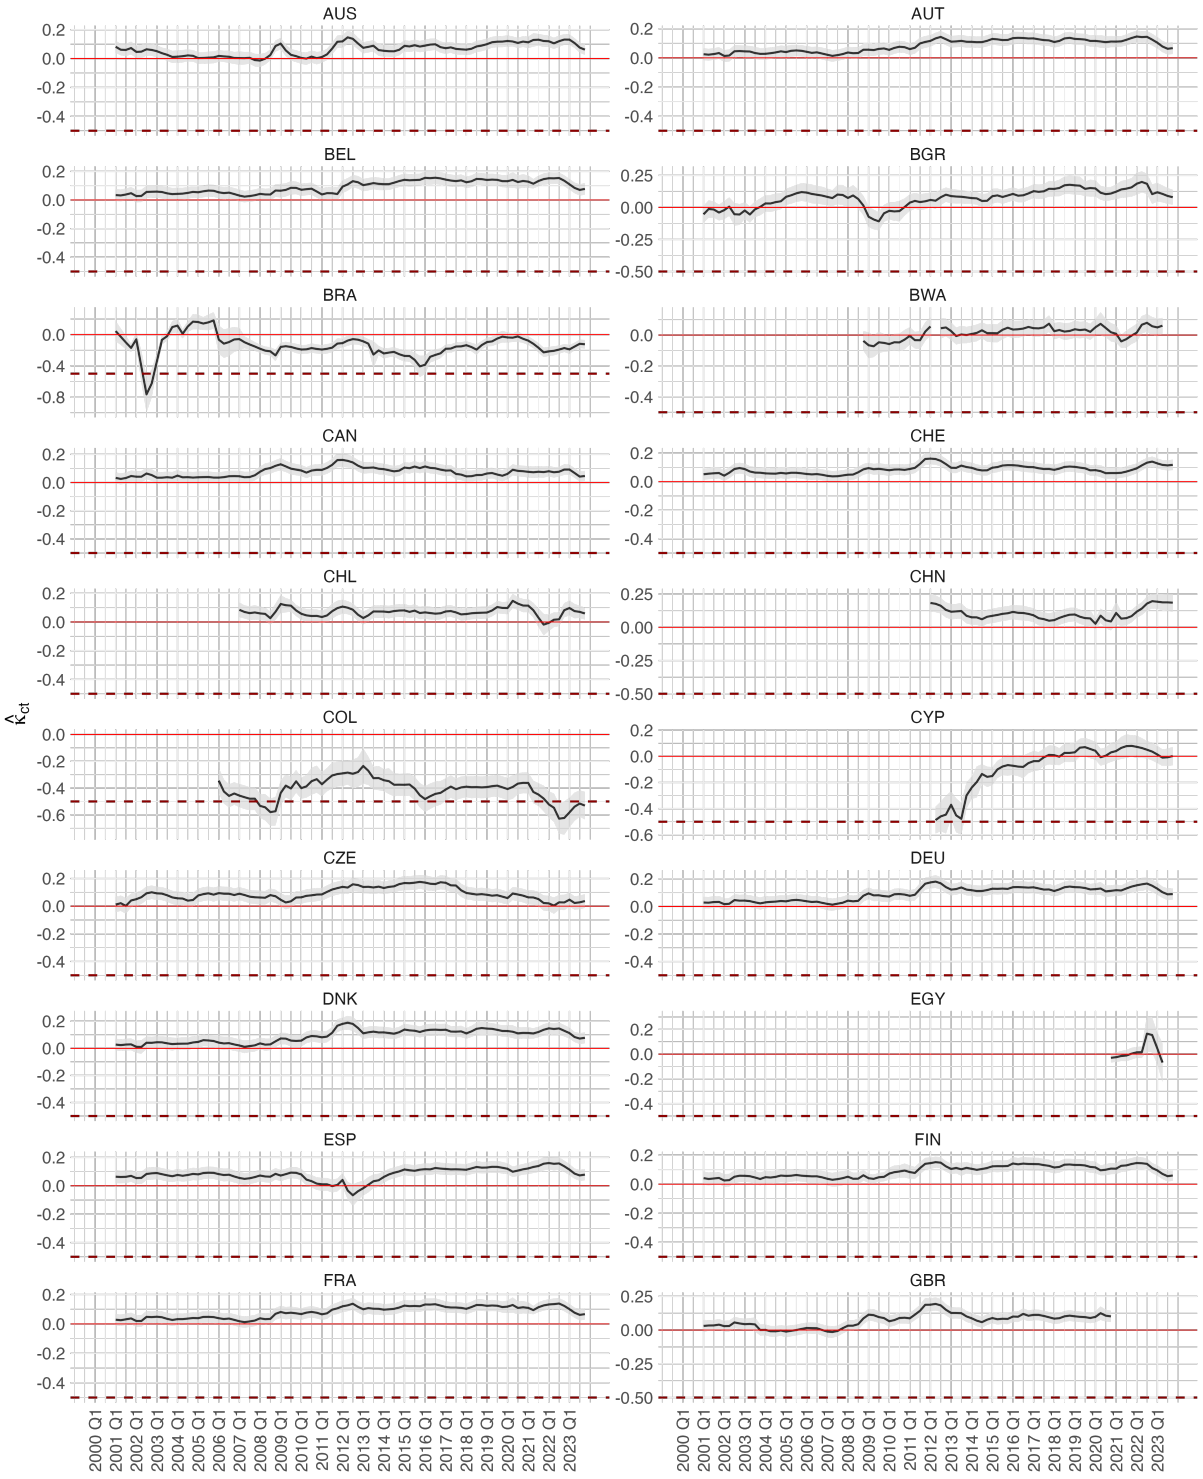

Figure A.3: $\hat{\kappa}_{ct}$ for all countries - part 1

Notes: The figure shows the evolution of $\hat{\kappa}_{ct}$ with 95% confidence intervals. The lines at 0 and -0.5 represent reference thresholds. The values are estimated based on the quarterly model.

Source: Author's calculations.

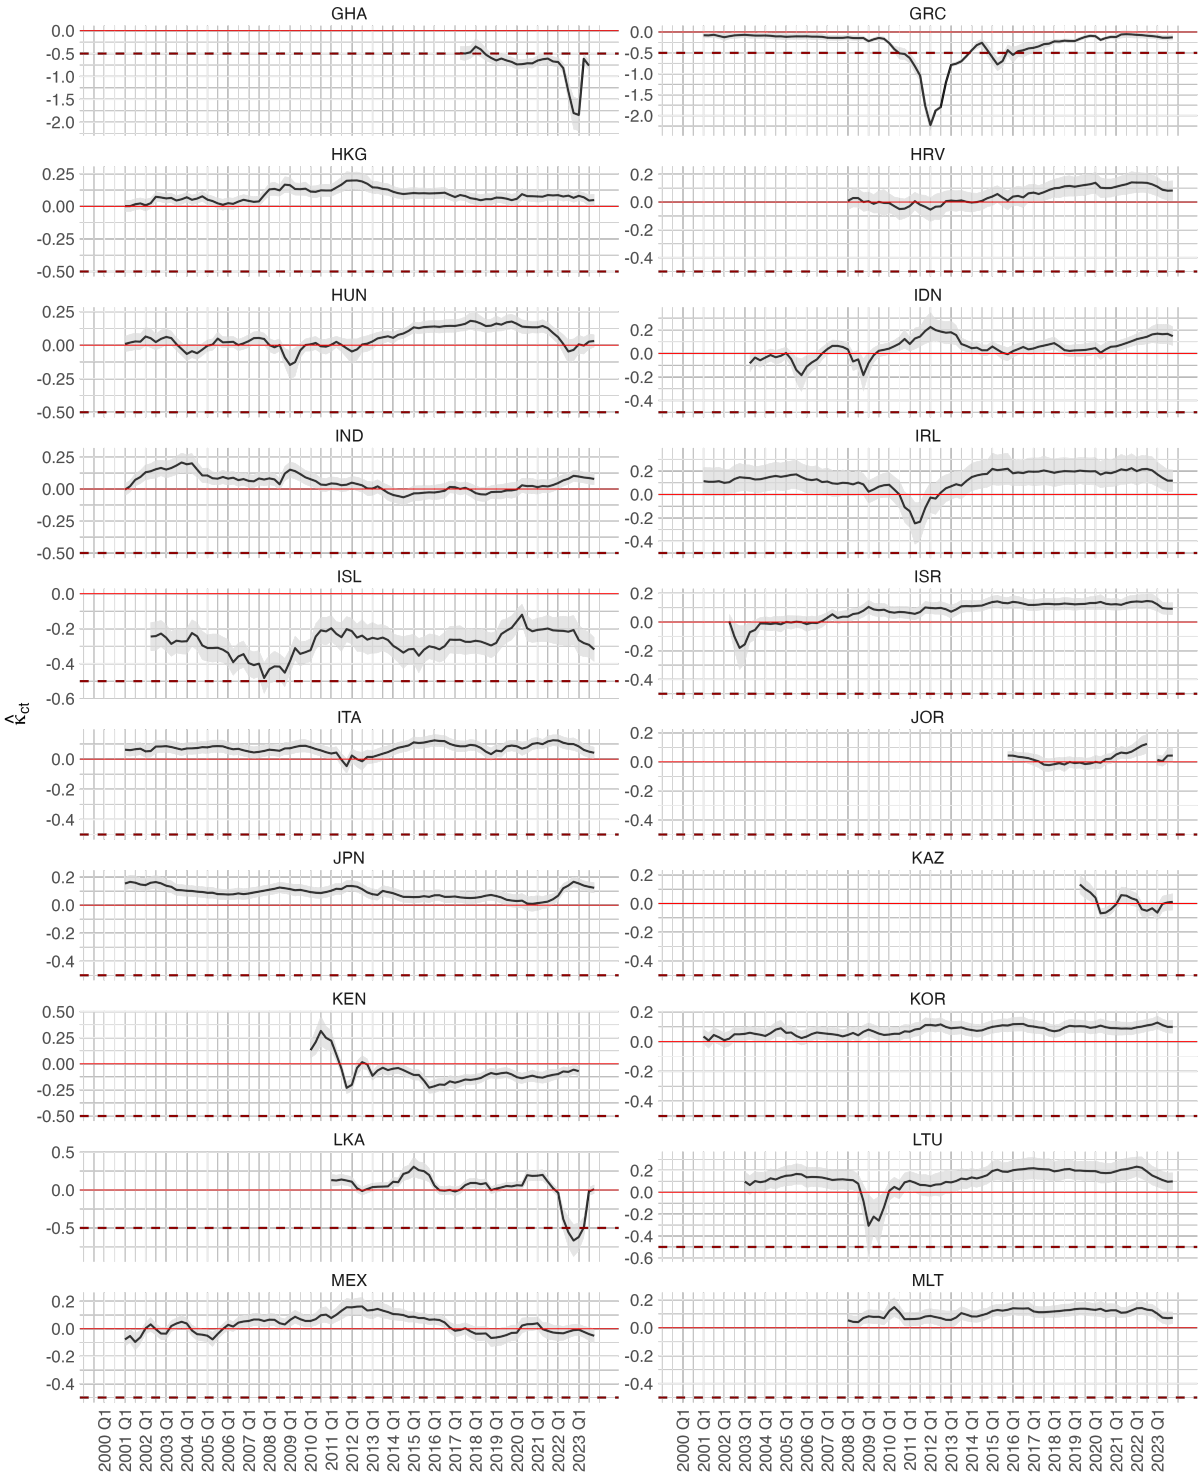

Figure A.4: $\hat{\kappa}_{ct}$ for all countries - part 2

Notes: The figure shows the evolution of $\hat{\kappa}_{ct}$ with 95% confidence intervals. The lines at 0 and -0.5 represent reference thresholds. The values are estimated based on the quarterly model.

Source: Author's calculations.

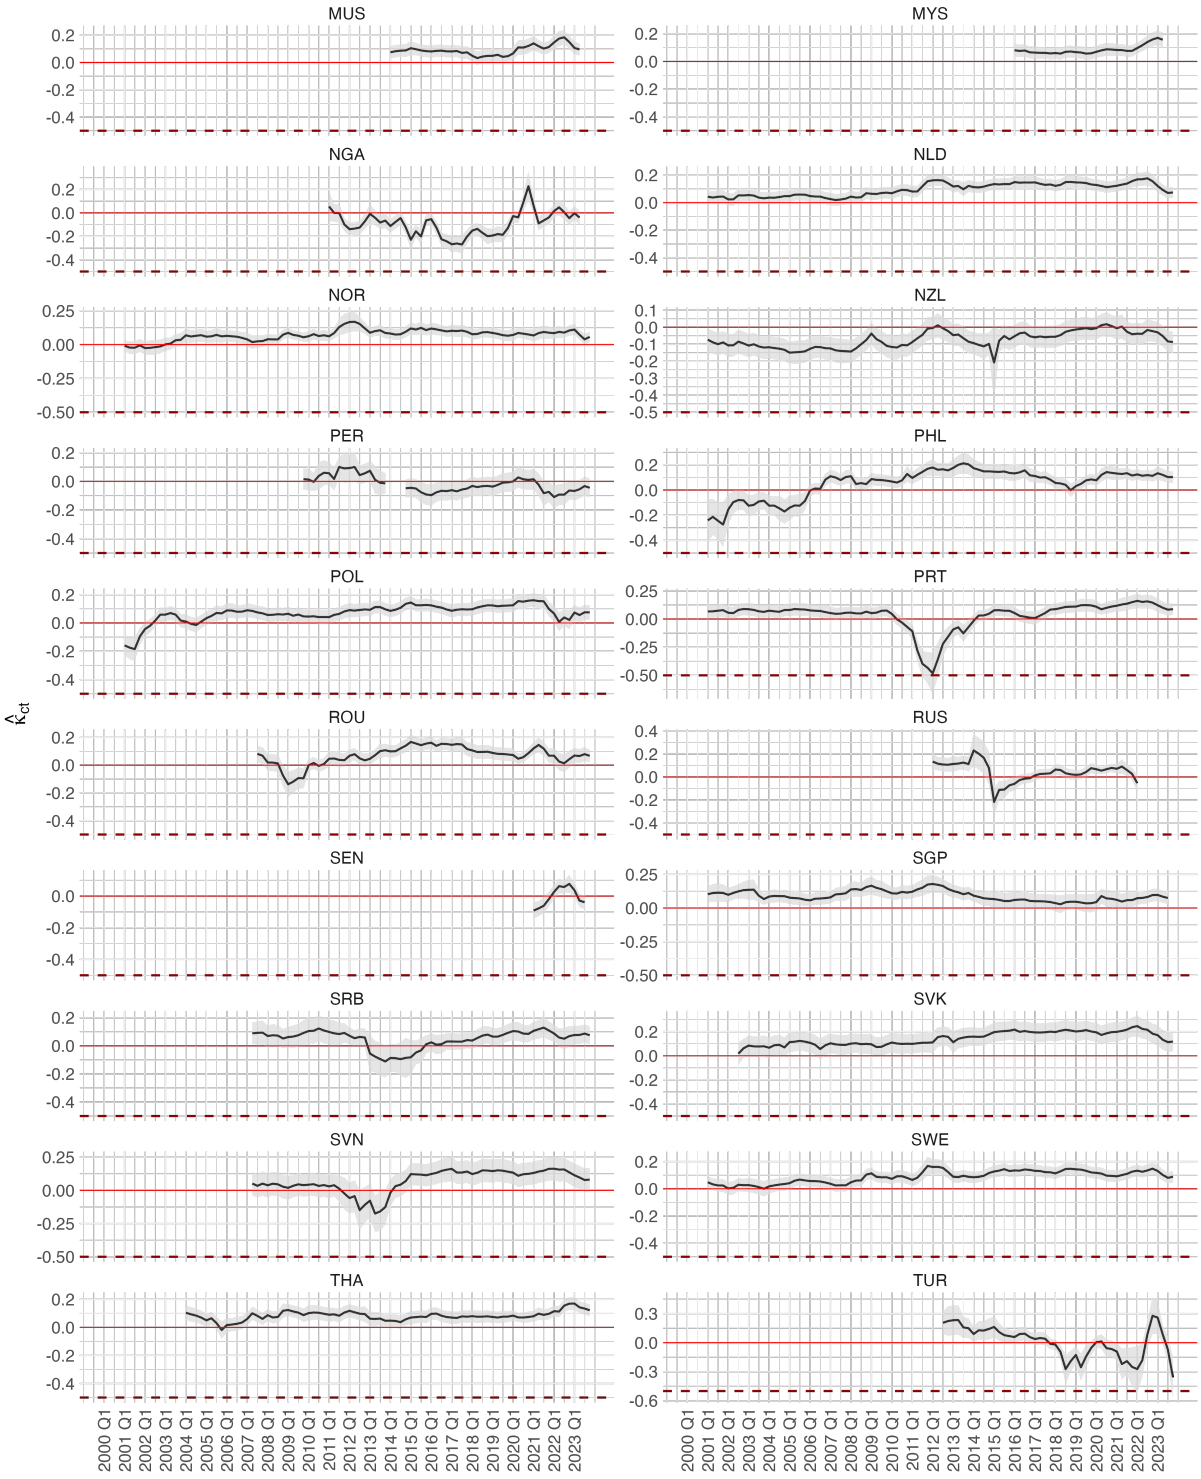

Figure A.5: $\hat{\kappa}_{ct}$ for all countries - part 3

Notes: The figure shows the evolution of $\hat{\kappa}_{ct}$ with 95% confidence intervals. The red line at 0 represents reference thresholds. The values are estimated based on the quarterly model.

Source: Author's calculations.

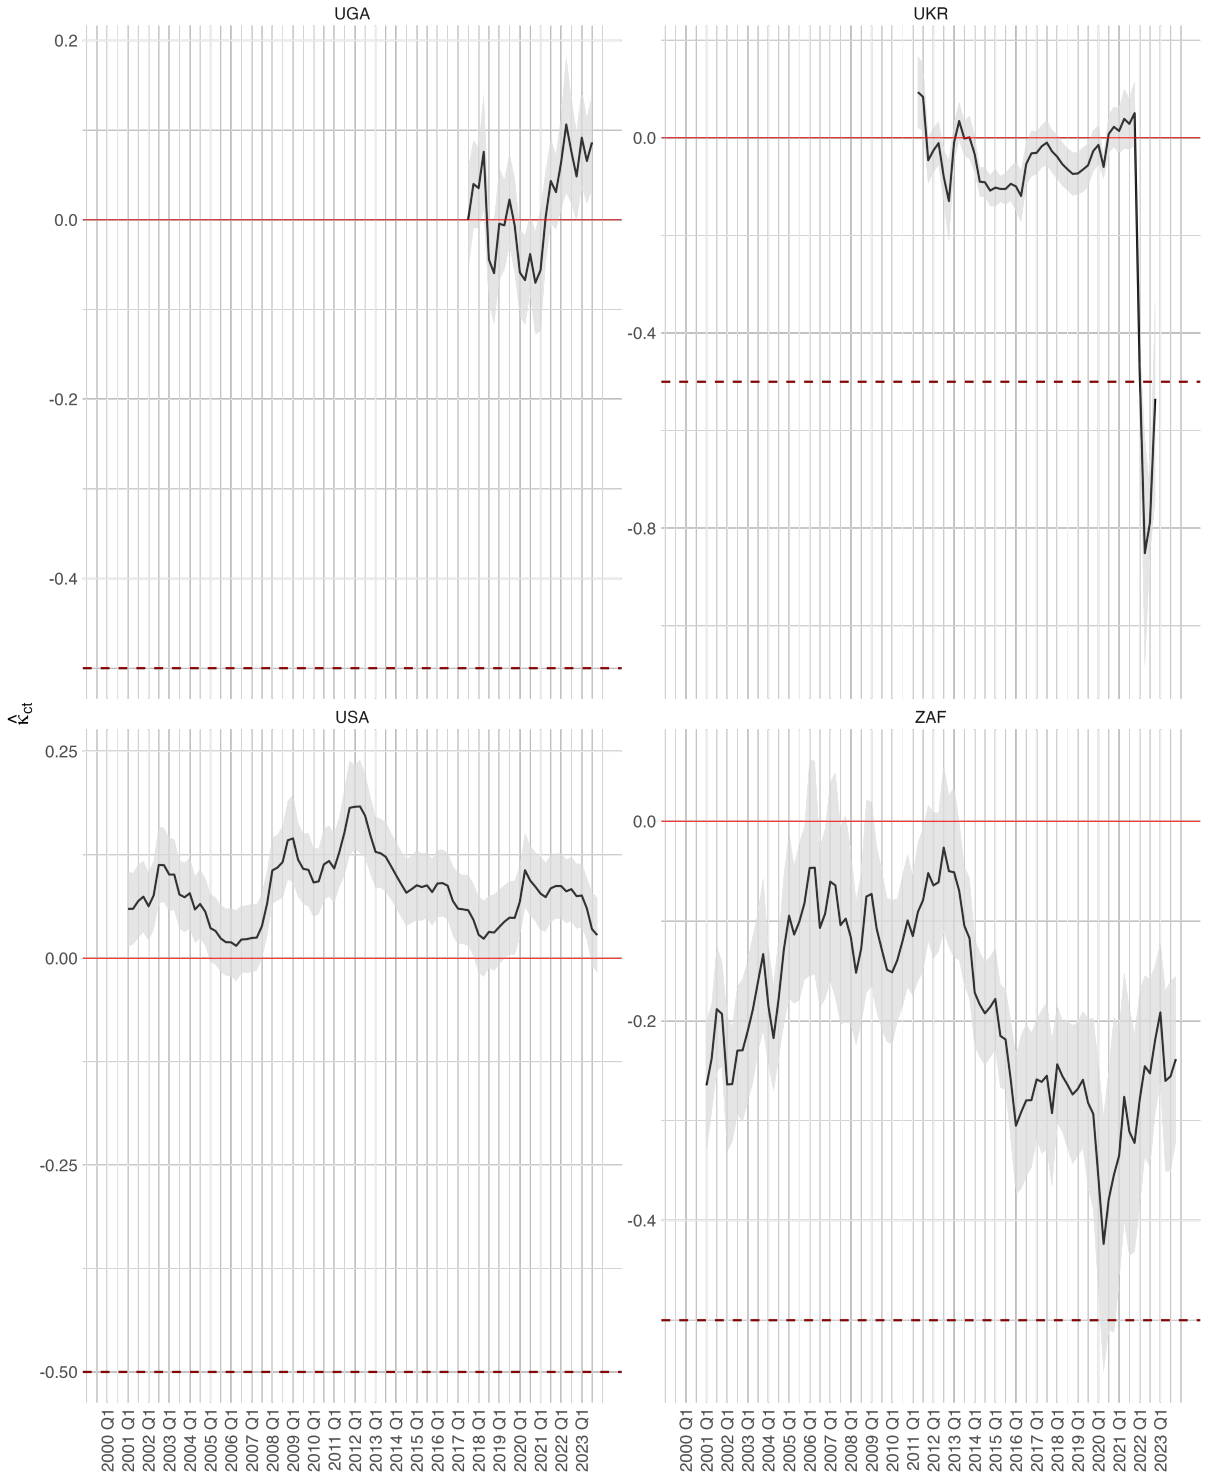

Figure A.6: $\hat{\kappa}_{ct}$ for all countries - part 4

Notes: The figure shows the evolution of $\hat{\kappa}_{ct}$ with 95% confidence intervals. The lines at 0 and -0.5 represent reference thresholds. The values are estimated based on the quarterly model.

Source: Author's calculations.

A.3 Appendix Tables

Table A.1: Yield curve data overview: country coverage, time span, and maturities

Isocode	Start date	End date	Maturities	Isocode	Start date	End date	Maturities
AUS	2001-01-01	2023-12-29	1-10	AUT	2001-01-01	2023-12-29	1-10
BEL	2001-01-02	2023-12-29	2-10	BGR	2001-02-19	2023-12-29	1-5, 7, 10
BRA	2001-01-05	2023-12-28	1-3, 5, 8, 10	BWA	2008-10-27	2023-06-02	3
CAN	2001-01-02	2023-12-29	1-5, 7, 10	CHE	2001-01-02	2023-12-29	1-10
CHL	2007-03-26	2023-12-29	1, 2, 4, 5, 8, 10	CHN	2012-01-04	2023-12-29	1-3, 5, 7, 10
COL	2006-01-02	2023-12-29	2, 4, 5, 10	CYP	2012-04-27	2023-12-29	2-4, 7, 10
CZE	2001-01-02	2023-12-29	1-10	DEU	2001-01-01	2023-12-29	1-10
DNK	2001-01-02	2023-12-29	2, 3, 5, 8, 10	EGY	2020-10-01	2023-06-30	1-3, 5, 7, 10
ESP	2001-01-02	2023-12-29	1-10	FIN	2001-01-02	2023-12-29	2-6, 8, 10
FRA	2001-01-01	2023-12-29	1-10	GBR	2001-01-01	2020-12-31	1-10
GHA	2017-04-20	2023-09-29	4, 5, 7, 8, 10	GRC	2001-01-03	2023-12-29	1-3, 5, 7, 10
HKG	2001-01-02	2023-12-29	1-3, 5, 7, 10	HRV	2008-01-30	2023-12-29	1-5, 10
HUN	2001-01-02	2023-12-29	1, 3, 5, 10	IDN	2003-05-14	2023-12-29	1, 3, 5, 10
IND	2001-01-01	2023-12-29	1-10	IRL	2001-01-02	2023-12-29	1-10
ISL	2002-04-15	2023-12-29	2, 5, 10	ISR	2002-04-09	2023-12-28	1-3, 5, 10
ITA	2001-01-02	2023-12-29	1-10	JOR	2015-10-06	2023-12-28	1-3, 5, 7, 10
JPN	2001-01-04	2023-12-29	1-10	KAZ	2019-04-19	2023-12-29	1-10
KEN	2010-01-04	2023-03-31	1-10	KOR	2001-01-02	2023-12-29	1-5, 10
LKA	2011-01-03	2023-12-29	1-10	LTU	2003-01-20	2023-12-29	3, 5, 10
MEX	2001-01-02	2023-12-29	1, 3, 5, 7, 10	MLT	2008-02-29	2023-12-29	1, 3, 5, 10
MUS	2014-01-22	2023-06-30	1-5, 10	MYS	2016-01-04	2023-06-30	1, 3, 5, 7, 10
NGA	2011-01-04	2023-06-30	1-5, 7, 10	NLD	2001-01-01	2023-12-29	2-10
NOR	2001-01-01	2023-12-29	1, 3, 5, 10	NZL	2001-01-03	2023-12-29	1, 2, 5, 7, 10
PER	2009-10-28	2023-12-29	2, 5, 10	PHL	2001-01-02	2023-12-29	1-5, 7, 10
POL	2001-01-02	2023-12-29	1-8, 10	PRT	2001-01-01	2023-12-29	1-10
ROU	2007-08-16	2023-12-29	1-5, 7, 10	RUS	2012-01-04	2022-03-31	1-3, 5, 7, 10
SEN	2021-03-24	2023-09-29	3, 5, 7	SGP	2001-01-02	2023-09-29	1, 2, 5, 10
SRB	2007-05-04	2023-12-28	1-3, 5, 10	SVK	2002-07-26	2023-12-29	2, 5, 6, 8-10
SVN	2007-04-03	2023-12-29	1, 3-5, 7, 8, 10	SWE	2001-01-02	2023-12-29	2, 5, 7, 10
THA	2004-01-05	2023-12-28	1-5, 7, 10	TUR	2012-07-02	2023-12-29	1-3, 5, 10
UGA	2017-07-03	2023-09-29	1-3, 5, 10	UKR	2011-04-01	2022-12-30	1-3, 6
USA	2001-01-01	2023-12-29	1-3, 5, 7, 10	ZAF	2001-01-02	2023-12-29	5, 10

Notes: The table provides an overview of yield curve data for 64 countries, showing the respective time spans and bond maturities between 1 and 10 years used in the analysis. The data is based on benchmark bonds.
Source: Datastream.

Table A.2: Fixed effects regression: impact of war on consumption growth and inflation

	Growth_t (1)	Inflation_t (2)
War_t	-0.023^{***}	0.018^{***}
	(0.004)	(0.004)
Growth_{t-1}	0.164^{***}	
	(0.014)	
Inflation_{t-1}		0.526^{***}
Country FE	✓	✓
Time FE	✓	✓
Observations	5,730	5,242
Adjusted R^2	0.003	0.290

Note: The table presents fixed effects regressions examining the impact of war on consumption growth (Growth_t) and inflation (Inflation_t). The regression equations are: $\log(\text{Growth}_t) = \beta_1^G \text{War}_t + \beta_2^G \log(\text{Growth}_{t-1}) + \kappa_c + \kappa_t + \epsilon_{it}$, $\log \text{Inflation}_t) = \beta_1^\Pi \text{War}_t + \beta_2^\Pi \log(\text{Inflation}_{t-1}) + \kappa_c + \kappa_t + \epsilon_{it}$. Robust standard errors are reported in parentheses, with $^{*}p < 0.1$, $^{**}p < 0.05$, and $^{***}p < 0.01$ indicating significance levels.
Source: WB/WDI and UCDP/GED data.

Table A.3: Fixed effect regression: quarterly model

	Observed price (q_{Nct})				
	(1)	(2)	(3)	(4)	(5)
Non-disaster price ($\hat{q}_{Nct}^{ND}$)	0.108^{***}	0.206^{***}	0.213^{***}	0.316^{***}	0.294^{***}
	(0.008)	(0.010)	(0.003)	(0.001)	(0.001)
Country-time FE	✓	✓	✓	✓	
Maturity-country FE	✓	✓			
Maturity-time FE	✓		✓		
Observations	28,725	28,725	28,725	28,725	28,725
Adjusted R^2	0.973	0.947	0.846	0.828	0.351

Note: The table presents a fixed effects regression of the observed log bond price ($q_{Nct} = \log(Q_{Nct})$) on the log of the theoretical non-disaster price ($q_{Nct}^{ND} = \log(Q_{Nct}^{ND})$). Fixed effects are denoted as κ_{Nc} (Maturity-Country), κ_{Nt} (Maturity-Time), and κ_{ct} (Country-Time). Models differ by their inclusion of these fixed effects. Robust standard errors are reported in parentheses, with $^{*}p < 0.1$, $^{**}p < 0.05$, and $^{***}p < 0.01$ indicating significance levels.
Source: Author's calculations using Datastream data for observed prices.

Appendix B

Appendix for Chapter 2

B.1 Filling missing data in the cross-country sample

Several variables are not available for the entire sample like the share of GDP produced in the natural resource sector in the previous years, the earlier average over the polity2 score, executive constraints, executive openness and executive competitiveness from the Polity5 dataset, and ethnic and religious fractionalization and polarization scores from Montalvo and Reynal-Querol (2005).

We impute the average values for all variables with missing values. We impute missing values through a clustering method that uses continent, longitude, and latitude to match countries. Note, this imputation imposes some measurement error but experiments with the smaller sample suggest performance is not affected substantially.

B.2 Appendix Figures

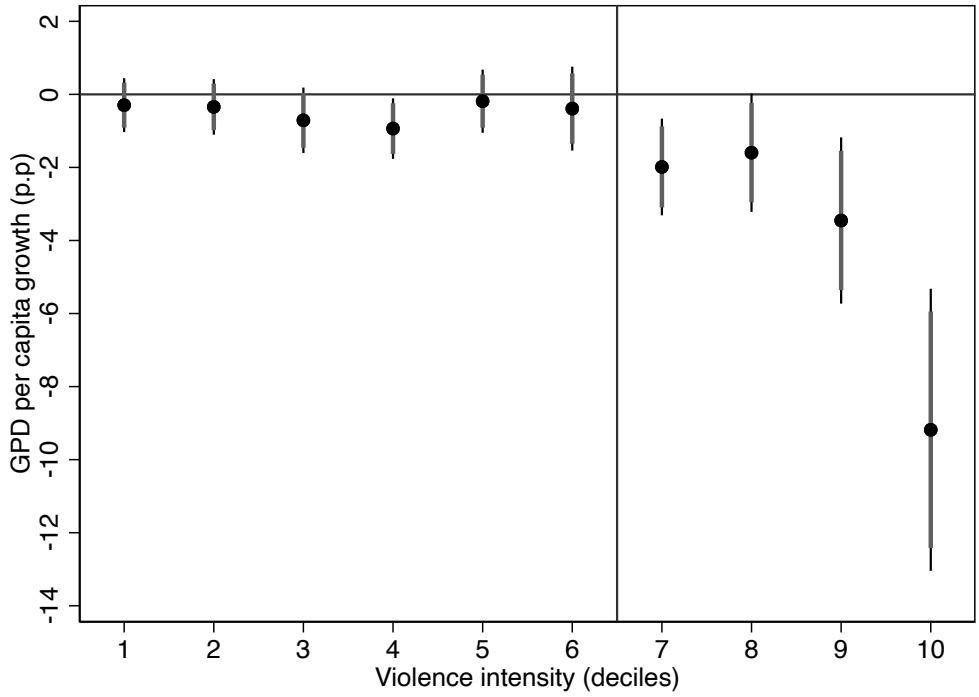

Figure B.1: Violence intensity and economic growth

Notes: The figure depicts the impact on GDP per capita growth for different levels of violence intensity, categorized by deciles. A country fixed effects model is used, where GDP per capita growth is regressed on a set of dummies representing each decile while controlling for year fixed effects. Violence intensity is quantified by the number of battle-related deaths per capita. Point estimates are displayed as circles where grey bars and black bars indicate significance at 5% and 10% respectively using robust standard errors.
Source: Authors' calculations based on GDP (constant 2015 US$) and population data from WB/WDI, and number of battle-related fatalities from UCDP/GED.

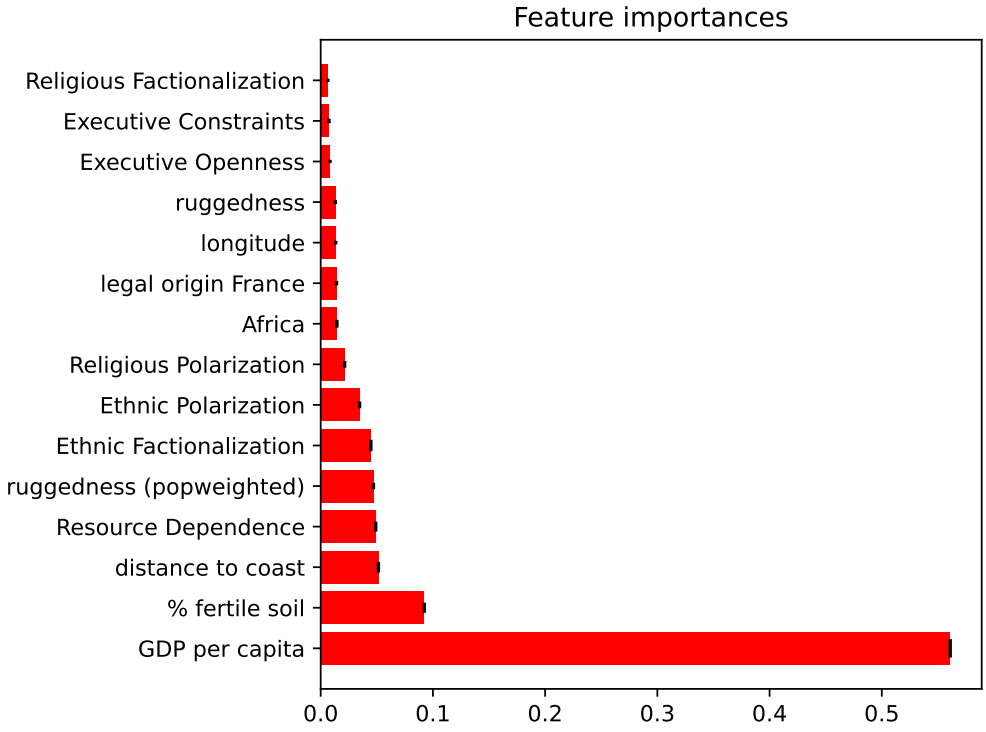

Figure B.2: Importances of random forest when predicting conflict trap

Notes: The figure shows the variables' importances in the random forest model for the 15 most important variables according to the random forest. A variable is deemed important if it appears often and tends to be chosen towards the top of the decision trees in the forest. In sklearn, feature importances are provided by the fitted attribute feature importances.

Source: Authors' calculations based on replication dataset from Nunn and Puga (2012) (legal origin France, Africa dummy, ruggedness, % fertile soil, % desert, distance to coast, longitude and latitude), Montalvo and Reynal-Querol (2005) (ethnic polarization/fractionalization and religious polarization/fractionalization), Polity5 dataset (polity2 score, executive constraints, executive openness, and executive competitiveness) and the WB/WDI (natural resource dependence of GDP, GDP, and population).

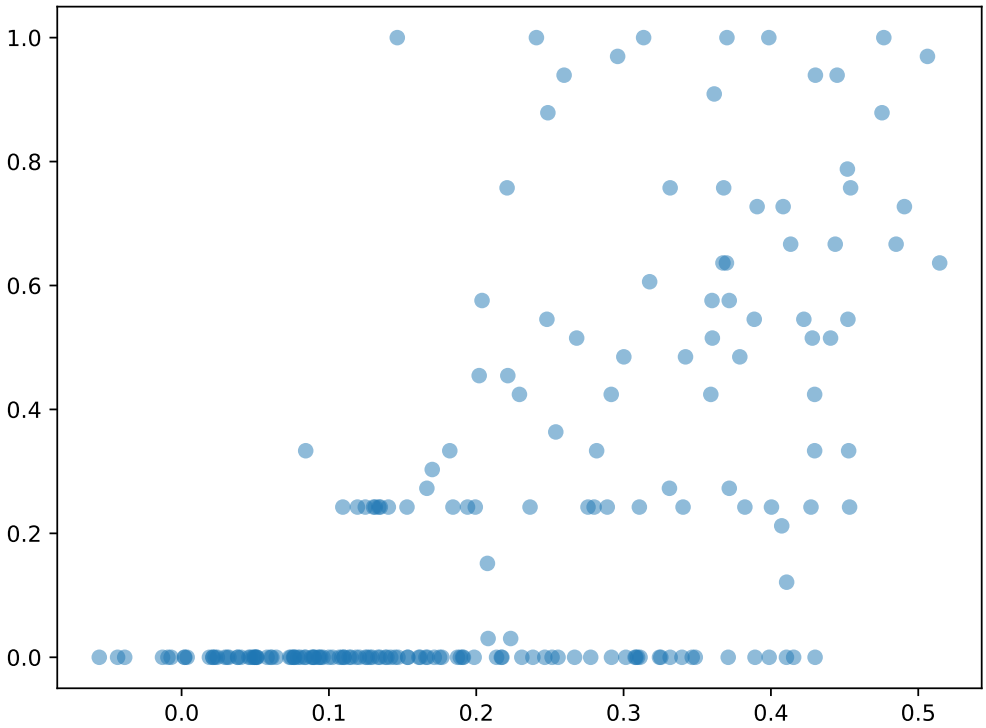

Figure B.3: Fitted values of ensemble model and actual values

Notes: The figure shows the within-fitted values (x-axis) when compared to the actual variation (y-axis). Note that the dependent variable here is the share spent inside the conflict trap and is therefore a value between 0 and 1. The ensemble is the mean of two prediction scores one of which is a linear regression model and is therefore not guaranteed to lie between 0 and 1. Note, this does not affect our classification into two classes.

Source: Authors' calculations based on replication dataset from Nunn and Puga (2012) (legal origin France, Africa dummy, ruggedness, % fertile soil, % desert, distance to coast, longitude and latitude), Montalvo and Reynal-Querol (2005) (ethnic polarization/fractionalization and religious polarization/fractionalization), Polity5 dataset (polity2 score, executive constraints, executive openness, and executive competitiveness) and the WB/WDI (natural resource dependence of GDP, GDP, and population).

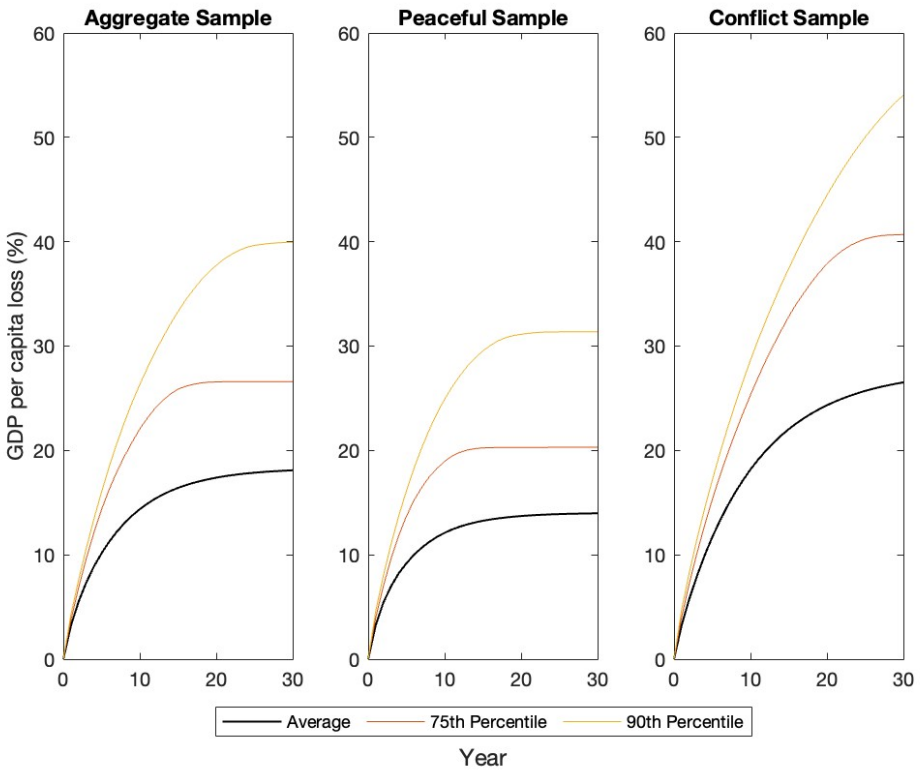

Figure B.4: Evolution of GDP per capita loss due with $\tau = 4$

Notes: The figure describes the evolution of GDP per capita loss due to entering into conflict when $\tau = 4$ for each specification: the aggregate sample, the peaceful sample, and the conflict sample. Growth paths are simulated for countries as they transition through the state space, utilizing a re-estimated transition matrix (see Table B.8) and a re-estimated growth vector (see Table B.9). Countries start in conflict and stable peace is absorbing. Conflict is defined as surpassing the threshold of violence intensity that significantly harms the economy, which corresponds to having more than 9.35 battle-related deaths per million inhabitants during a year. The distribution at each period is described by the average, 75th percentile, and 90th percentile. The horizon for simulation is $T = 30$, and the number of simulations conducted is $N = 100000$. The methodology for sample partitioning is described in Section 2.3.2

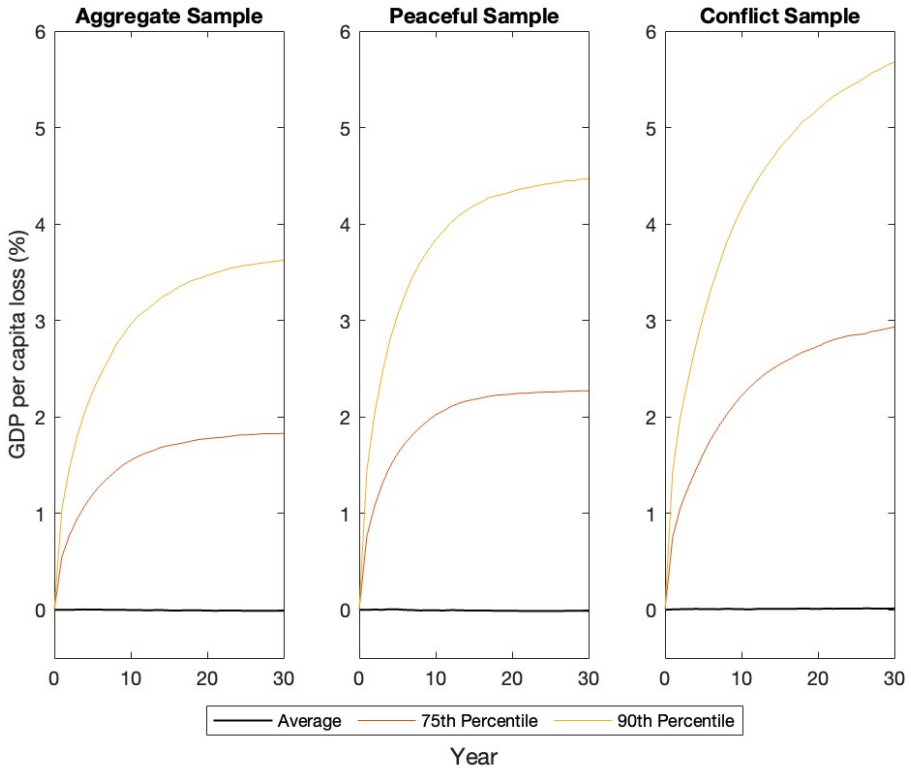

Figure B.5: Evolution of GDP per capita loss due to the variation component of the estimation

Notes: The figure describes the evolution of GDP per capita loss due to entering into conflict caused solely by the variation component of the estimation for each specification: the aggregate sample, the peaceful sample, and the conflict sample. To achieve this, the effect coming from the coefficients is eliminated, which represents the means in the growth vector. Formally, for $h < \tau + 1$, $\hat{f}_h = \mathcal{N}(0, \hat{\sigma}_h^{SE})$. The transition matrix for each specification is the same as in the main setting (Table 2.1) and the standard errors come from Table 2.2. Countries start in conflict and stable peace is absorbing. Conflict is defined as surpassing the threshold of violence intensity that significantly harms the economy, which corresponds to having more than 9.35 battle-related deaths per million inhabitants during a year. The distribution at each period is described by the average, 75th percentile, and 90th percentile. The horizon for simulation is $T = 30$, and the number of simulations conducted is $N = 100000$. The methodology for sample partitioning is described in Section 2.3.2

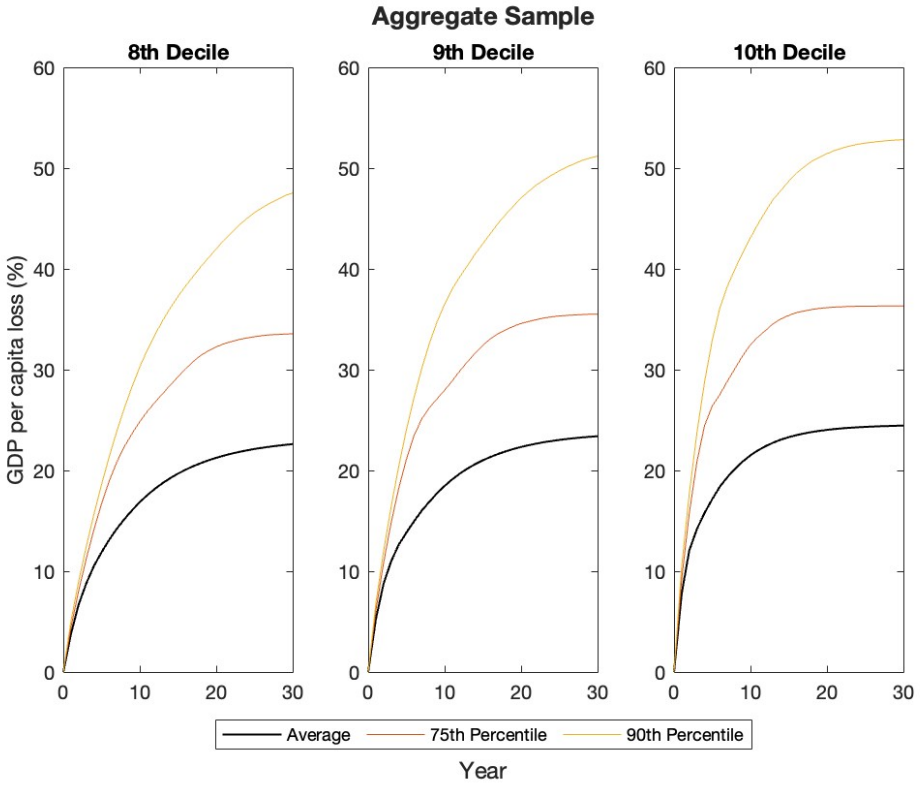

Figure B.6: Evolution of GDP per capita loss for more restrictive definitions of conflict

Notes: The figure describes the evolution of GDP per capita loss due to entering into conflict for different conflict definitions, each associated with surpassing different deciles of violence. Growth paths are simulated for countries as they transition through the state space, utilizing a re-estimated transition matrix (see Table B.6) and a re-estimated growth vector (see Table B.7) based on the aggregate sample. Countries start in conflict and stable peace is absorbing. Conflict is defined as having more than 21.71 (8th decile), 51.77 (9th decile), and 155.56 (10th decile) battle-related deaths per million inhabitants during a year. The distribution at each period is described by the average, 75th percentile, and 90th percentile. The horizon for simulation is $T = 30$, and the number of simulations conducted is $N = 100000$.

B.3 Appendix Tables

Table B.1: Country data overview

ISO	Years	CS	ISO	Years	CS	ISO	Years	CS	ISO	Years	CS
AFG	2003 - 2020	0	AGO	1990 - 2020	0	ALB	1990 - 2020	0	AND	1990 - 2020	0
ARE	1990 - 2020	0	ARG	1990 - 2020	0	ARM	1991 - 2020	0	ATG	1990 - 2020	0
AUS	1990 - 2020	0	AUT	1990 - 2020	0	AZE	1991 - 2020	0	BDI	1990 - 2020	1
BEL	1990 - 2020	0	BEN	1990 - 2020	0	BFA	1990 - 2020	1	BGD	1990 - 2020	0
BGR	1990 - 2020	0	BHR	1990 - 2020	0	BHS	1990 - 2020	0	BIH	1995 - 2020	0
BLR	1991 - 2020	0	BLZ	1990 - 2020	0	BMU	1990 - 2020	0	BOL	1990 - 2020	0
BRA	1990 - 2020	0	BRB	1990 - 2020	0	BRN	1990 - 2020	0	BTN	1990 - 2020	0
BWA	1990 - 2020	0	CAF	1990 - 2020	1	CAN	1998 - 2020	0	CHE	1990 - 2020	0
CHL	1990 - 2020	0	CHN	1990 - 2020	1	CIV	1990 - 2020	1	CMR	1990 - 2020	1
COD	1990 - 2020	1	COG	1990 - 2020	1	COL	1990 - 2020	0	COM	1990 - 2020	0
CPV	1990 - 2020	1	CRI	1990 - 2020	0	CUB	1990 - 2020	0	CYP	1990 - 2020	0
CZE	1991 - 2020	0	DEU	1990 - 2020	0	DJI	2014 - 2020	1	DMA	1990 - 2020	0
DNK	1990 - 2020	0	DOM	1990 - 2020	0	DZA	1990 - 2020	0	ECU	1990 - 2020	0
EGY	1990 - 2020	1	ERI	1993 - 2011	1	ESP	1990 - 2020	0	EST	1996 - 2020	0
ETH	1990 - 2020	1	FIN	1990 - 2020	0	FJI	1990 - 2020	0	FRA	1990 - 2020	0
FSM	1990 - 2020	0	GAB	1990 - 2020	0	GBR	1990 - 2020	0	GEO	1990 - 2020	0
GHA	1990 - 2020	0	GIN	1990 - 2020	1	GMB	1990 - 2020	0	GNB	1990 - 2020	1
GNQ	1990 - 2020	1	GRC	1990 - 2020	0	GRD	1990 - 2020	0	GTM	1990 - 2020	0
GUY	1990 - 2020	0	HKG	1990 - 2020	0	HND	1990 - 2020	0	HRV	1996 - 2020	0
HTI	1990 - 2020	0	HUN	1992 - 2020	0	IDN	1990 - 2020	0	IND	1990 - 2020	0
IRL	1990 - 2020	0	IRN	1990 - 2020	0	IRQ	1990 - 2020	1	ISL	1996 - 2020	0
ISR	1996 - 2020	0	ITA	1990 - 2020	0	JAM	1990 - 2020	0	JOR	1990 - 2020	0
JPN	1990 - 2020	0	KAZ	1991 - 2020	0	KEN	1990 - 2020	0	KGZ	1990 - 2020	1
KHM	1994 - 2020	0	KIR	1990 - 2020	0	KNA	1990 - 2020	0	KOR	1990 - 2020	0
KWT	1993 - 2020	0	LAO	1990 - 2020	0	LBN	1990 - 2020	0	LBR	2001 - 2020	1
LBY	2000 - 2020	1	LCA	1990 - 2020	0	LKA	1990 - 2020	0	LSO	1990 - 2020	0
LTU	1996 - 2020	0	LUX	1990 - 2020	0	LVA	1996 - 2020	0	MAC	1990 - 2020	0
MAR	1990 - 2020	1	MCO	1990 - 2019	0	MDA	1996 - 2020	0	MDG	1990 - 2020	0
MDV	1996 - 2020	0	MEX	1990 - 2020	0	MHL	1990 - 2020	0	MKD	1991 - 2020	0
MLI	1990 - 2020	1	MLT	1990 - 2020	0	MMR	1990 - 2020	0	MNE	1998 - 2020	0
MNG	1990 - 2020	0	MOZ	1990 - 2020	0	MRT	1990 - 2020	1	MUS	1990 - 2020	0
MWI	1990 - 2020	0	MYS	1990 - 2020	0	NAM	1990 - 2020	0	NER	1990 - 2020	1
NGA	1990 - 2020	0	NIC	1990 - 2020	0	NLD	1990 - 2020	0	NOR	1990 - 2020	0
NPL	1990 - 2020	1	NRU	2005 - 2020	0	NZL	1990 - 2020	0	OMN	1990 - 2019	0
PAK	1990 - 2020	1	PAN	1990 - 2020	0	PER	1990 - 2020	0	PHL	1990 - 2020	0
PLW	2001 - 2020	0	PNG	1990 - 2020	0	POL	1991 - 2020	0	PRI	1990 - 2020	0
PRT	1990 - 2020	0	PRY	1990 - 2020	0	QAT	2001 - 2020	0	ROU	1991 - 2020	0
RUS	1990 - 2020	0	RWA	1990 - 2020	0	SAU	1990 - 2020	0	SDN	1990 - 2020	1
SEN	1990 - 2020	0	SGP	1990 - 2020	0	SLB	1990 - 2020	0	SLE	1990 - 2020	1
SLV	1990 - 2020	0	SMR	1998 - 2019	0	SOM	2014 - 2020	1	SRB	1996 - 2020	0
SSD	2009 - 2015	1	STP	2002 - 2020	0	SUR	1990 - 2020	0	SVK	1993 - 2020	0
SVN	1996 - 2020	0	SWE	1990 - 2020	0	SWZ	1990 - 2020	0	SYC	1990 - 2020	0
SYR	1990 - 2019	0	TCD	1990 - 2020	1	TGO	1990 - 2020	1	THA	1990 - 2020	0
TJK	1990 - 2020	1	TKM	1990 - 2019	0	TLS	2001 - 2020	0	TON	1990 - 2020	0
TTO	1990 - 2020	0	TUN	1990 - 2020	0	TUR	1990 - 2020	0	TUV	1991 - 2020	0
TZA	1990 - 2020	0	UGA	1990 - 2020	1	UKR	1990 - 2020	0	URY	1990 - 2020	0
USA	1990 - 2020	0	UZB	1990 - 2020	1	VCT	1990 - 2020	0	VNM	1990 - 2020	0
VUT	1990 - 2020	0	WSM	1990 - 2020	0	XKX	2009 - 2020	1	YEM	1991 - 2018	0
ZAF	1990 - 2020	0	ZMB	1990 - 2020	1	ZWE	1990 - 2020	0			

Notes: The table shows the ISO codes of the countries included in the dataset, the years covered, and if it belongs to the Peaceful Sample (CS=0) or the Conflict Sample (CS=1).
Source: GDP (constant 2015 US$) and population data from WB/WDI, and number of battle-related fatalities from UCDP/GED. The methodology for sample partitioning is described in Section 2.3.2

Table B.2: Summary statistics for aggregate sample (N=5718)

	Mean	Std. Dev.	Min	Max
Population (Mill.)	34.65	131.05	0.01	1402.11
GDP (Cons. 2015 US$, Bill.)	306.07	1326.71	0.02	19974.54
Growth Rate	0.02	0.06	-0.65	1.40
Num. Battle-Related Fatalities	426.44	7506.77	0.00	533436.00
Num. Battle-Related Fatalities per Pop	0.03	1.20	0.00	89.86
Conflict Status Dummy	0.10	0.30	0.00	1.00
ML Predicted Score	0.22	0.14	-0.06	0.51

Notes: The table shows summary statistics for the Aggregate Sample. Conflict is defined as surpassing the threshold of violence intensity that significantly harms the economy, which corresponds to having more than 9.35 battle-related deaths per million inhabitants during a year.
Source: Authors' calculations based on GDP (constant 2015 US$) and population data from WB/WDI, and number of battle-related fatalities from UCDP/GED.

Table B.3: Summary statistics for peaceful sample (N=4707)

	Mean	Std. Dev.	Min	Max
Population (Mill.)	28.88	100.63	0.01	1380.00
GDP (Cons. 2015 US$, Bill.)	326.02	1337.23	0.02	19974.54
Growth Rate	0.02	0.05	-0.55	0.92
Num. Battle-Related Fatalities	361.16	8144.52	0.00	533436.00
Num. Battle-Related Fatalities per Pop	0.03	1.32	0.00	89.86
Conflict Status Dummy	0.06	0.24	0.00	1.00
ML Predicted Score	0.17	0.11	-0.06	0.37

Notes: The table shows summary statistics for the A Peaceful sample. Conflict is defined as surpassing the threshold of violence intensity that significantly harms the economy, which corresponds to having more than 9.35 battle-related deaths per million inhabitants during a year.
Source: Authors' calculations based on GDP (constant 2015 US$) and population data from WB/WDI, and number of battle-related fatalities from UCDP/GED.

Table B.4: Summary statistics for conflict sample (N=1011)

	Mean	Std. Dev.	Min	Max
Population (Mill.)	61.52	221.70	0.34	1402.11
GDP (Cons. 2015 US$, Bill.)	213.17	1273.09	0.17	14631.84
Growth Rate	0.02	0.10	-0.65	1.40
Num. Battle-Related Fatalities	730.35	3128.39	0.00	49856.00
Num. Battle-Related Fatalities per Pop	0.05	0.29	0.00	7.69
Conflict Status Dummy	0.27	0.45	0.00	1.00
ML Predicted Score	0.43	0.04	0.37	0.51

Notes: The table shows summary statistics for the Conflict Sample. Conflict is defined as surpassing the threshold of violence intensity that significantly harms the economy, which corresponds to having more than 9.35 battle-related deaths per million inhabitants during a year.
Source: Authors' calculations based on GDP (constant 2015 US$) and population data from WB/WDI, and number of battle-related fatalities from UCDP/GED.

Table B.5: Violence intensity and economic growth

	GDP per capita Growth
1st Decile	-0.295
	(0.374)
2nd Decile	-0.342
	(0.385)
3rd Decile	-0.710
	(0.454)
4th Decile	-0.936*
	(0.420)
5th Decile	-0.187
	(0.438)
6th Decile	-0.389
	(0.582)
7th Decile	-1.987**
	(0.671)
8th Decile	-1.597
	(0.821)
9th Decile	-3.453**
	(1.154)
10th Decile	-9.183***
	(1.958)
Observations	5730
Country FE	Yes
Time FE	Yes
(Within country) R^2	0.142

Notes: The table shows the impact on GDP per capita growth for different levels of violence intensity, categorized by deciles. A country fixed effects model is used, where GDP per capita growth is regressed on a set of dummies representing each decile while controlling for year fixed effects. Year fixed effects were also included as control variables. Violence intensity is quantified by the number of battle-related deaths per capita. Robust standard errors are in parentheses. $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.
Source: Authors' calculations based on GDP (constant 2015 US$) and population data from WB/WDI, and number of battle-related fatalities from UCDP/GED.

Table B.6: Estimated transition probabilities for more restrictive definitions of conflict

		Stay in Conflict	Unstable Peace to Conflict						Stable Peace to Conflict
		$\hat{\pi}_0$	$\hat{\pi}_1$	$\hat{\pi}_2$	$\hat{\pi}_3$	$\hat{\pi}_4$	$\hat{\pi}_5$	$\hat{\pi}_6$	$\hat{\pi}_7$	$\hat{\pi}_8$
AS	D8	0.76	0.16	0.13	0.09	0.06	0.04	0.07	0.06	0.01
	D9	0.69	0.24	0.08	0.08	0.1	0.09	0	0.09	0.01
	D10	0.64	0.1	0.15	0.07	0.1	0.03	0.03	0.03	0.01

Notes: The table shows the estimated transition probabilities for different conflict definitions, each associated with surpassing different deciles of violence. The analysis is performed using the 8th decile (D8), 9th decile (D9), and 10th decile (D10) on the aggregate sample (AS). They are estimated using the proportion of transitions observed in the data. $\hat{\pi}_i$ denotes the probability of transitioning to conflict from state i, where $i = 0$ is conflict, $i = k$ s.t. $k \in [1, 7]$ is the kth number of consecutive years in post-conflict peace, and $i = 8$ is stable peace. Conflict is defined as having more than 21.71 (8th decile), 51.77 (9th decile), and 155.56 (10th decile) battle-related deaths per million inhabitants during a year.
Source: Authors' calculations based on population data from WB/WDI, and number of battle-related fatalities from UCDP/GED.

Table B.7: Estimation of the growth vector for more restrictive definitions of conflict

	Aggregate Sample		
	8th Decile	9th Decile	10th Decile
	GDP per capita Growth		
Conflict	-0.0378***	-0.0522***	-0.0779***
	(0.00996)	(0.0123)	(0.0210)
1st Year Post-conflict Peace	-0.00579	-0.00211	0.00887
	(0.00929)	(0.0133)	(0.0248)
2nd Year Post-conflict Peace	0.00755	0.0182	0.0219
	(0.00770)	(0.0101)	(0.0159)
3rd Year Post-conflict Peace	-0.000627	0.00677	0.00569
	(0.00671)	(0.00874)	(0.0149)
4th Year Post-conflict Peace	0.00770	0.0121	0.0000339
	(0.00523)	(0.00773)	(0.00962)
5th Year Post-conflict Peace	-0.00370	-0.00199	-0.00126
	(0.00624)	(0.00700)	(0.00677)
6th Year Post-conflict Peace	-0.000231	-0.000573	0.00384
	(0.00484)	(0.00594)	(0.00878)
7th Year Post-conflict Peace	0.00178	0.00156	0.0000347
	(0.00536)	(0.00665)	(0.0104)
Observations	5730	5730	5730
Country FE	Yes	Yes	Yes
Time FE	Yes	Yes	Yes
(Within country) R^2	0.123	0.131	0.136

Notes: The table shows the regression results from equation 2 used to estimate the growth vector for different conflict definitions, each associated with surpassing different deciles of violence. The analysis is performed using the 8th decile, 9th decile, and 10th decile on the aggregate sample. A country fixed effects model is employed where GDP per capita growth is regressed on a set of dummies representing the states of the model. The base category is stable peace, defined as having more than 7 consecutive years of post-conflict peace. Year fixed effects were also included as control variables. Conflict is defined as having more than 21.71 (8th decile), 51.77 (9th decile), and 155.56 (10th decile) battle-related deaths per million inhabitants during a year. Robust standard errors are in parentheses. $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.
Source: Authors' calculations based on GDP (constant 2015 US$) and population data from WB/WDI, and number of battle-related fatalities from UCDP/GED.

Table B.8: Estimated transition probabilities when $\tau = 4$

	Stay in Conflict	Unstable Peace to Conflict				Stable Peace to Conflict
	$\hat{\pi}_0$	$\hat{\pi}_1$	$\hat{\pi}_2$	$\hat{\pi}_3$	$\hat{\pi}_4$	$\hat{\pi}_5$
AS	0.76	0.22	0.16	0.09	0.03	0.02
PS	0.71	0.20	0.12	0.07	0.03	0.01
CS	0.82	0.25	0.23	0.12	0.03	0.05

Notes: The table shows the estimated transition probabilities when $\tau = 4$ for each specification: the aggregate sample (AS), the peaceful sample (PS), and the conflict sample (CS). They are estimated using the proportion of transitions observed in the data. $\hat{\pi}_i$ denotes the probability of transitioning to conflict from state i, where $i = 0$ is conflict, $i = k$ s.t. $k \in [1, 4]$ is the kth number of consecutive years in post-conflict peace, and $i = 5$ is stable peace. Conflict is defined as surpassing the threshold of violence intensity that significantly harms the economy, which corresponds to having more than 9.35 battle-related deaths per million inhabitants during a year.
Source: Authors' calculations based on population data from WB/WDI, and number of battle-related fatalities from UCDP/GED. The methodology for sample partitioning is described in Section 2.3.2

Table B.9: Estimation of the growth vector when $\tau = 4$

	Aggregate Sample	Peaceful Sample	Conflict Sample
	GDP per capita Growth		
Conflict	-0.0323***	-0.0333**	-0.0327**
	(0.00813)	(0.0110)	(0.0112)
1st Year Post-conflict Peace	0.00452	0.00436	0.00484
	(0.00734)	(0.0104)	(0.00900)
2nd Year Post-conflict Peace	0.00147	0.00510	-0.00374
	(0.00677)	(0.00984)	(0.00622)
3rd Year Post-conflict Peace	0.000628	0.000801	-0.00201
	(0.00594)	(0.00752)	(0.00989)
4th Year Post-conflict Peace	0.00579	0.00511	0.00245
	(0.00523)	(0.00656)	(0.00920)
Observations	5730	4676	1054
Country FE	Yes	Yes	Yes
Time FE	Yes	Yes	Yes
(Within country) R^2	0.122	0.175	0.0811

Notes: The table shows the regression results when $\tau = 4$ from equation 2 used to estimate the growth vector for each specification: the aggregate sample, the peaceful sample, and the conflict sample. A country fixed effects model is employed where GDP per capita growth is regressed on a set of dummies representing the states of the model. The base category is stable peace, defined as having more than 4 consecutive years of post-conflict peace. Year fixed effects were also included as control variables. Conflict is defined as surpassing the threshold of violence intensity that significantly harms the economy, which corresponds to having more than 9.35 battle-related deaths per million inhabitants during a year. Robust standard errors are in parentheses. $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.
Source: Authors' calculations based on GDP (constant 2015 US$) and population data from WB/WDI, and number of battle-related fatalities from UCDP/GED. The methodology for sample partitioning is described in Section 2.3.2.

Table B.10: Estimation of the growth vector with country-specific time trends

	Aggregate Sample	Peaceful Sample	Conflict Sample
	GDP per capita Growth		
Conflict	-0.0371***	-0.0351**	-0.0372**
	(0.00851)	(0.0119)	(0.0117)
1st Year Post-conflict Peace	0.00542	0.00958	0.00162
	(0.00770)	(0.0108)	(0.0108)
2nd Year Post-conflict Peace	0.00154	0.00667	-0.00497
	(0.00708)	(0.00995)	(0.00963)
3rd Year Post-conflict Peace	0.000591	0.00206	-0.00265
	(0.00677)	(0.00811)	(0.0139)
4th Year Post-conflict Peace	0.00655	0.00779	0.00412
	(0.00587)	(0.00682)	(0.0123)
5th Year Post-conflict Peace	0.000269	0.000837	-0.000189
	(0.00502)	(0.00588)	(0.0104)
6th Year Post-conflict Peace	0.00173	-0.00131	0.00603
	(0.00519)	(0.00575)	(0.0104)
7th Year Post-conflict Peace	0.00103	-0.00477	0.0155
	(0.00464)	(0.00508)	(0.00992)
Observations	5730	4707	1023
Country FE	Yes	Yes	Yes
Time FE	Yes	Yes	Yes
Country Time Trends	Yes	Yes	Yes
(Within country) R^2	0.197	0.249	0.161

Notes: The table shows the regression results from equation 2, estimating the growth vector for each specification: the aggregate sample, the peaceful sample, and the conflict sample. A country fixed effects model is employed where GDP per capita growth is regressed on a set of dummies representing the states of the model, including country-specific time trends. The base category is stable peace, defined as having more than 7 consecutive years of post-conflict peace. Year fixed effects were also included as control variables. Conflict is defined as surpassing the threshold of violence intensity that significantly harms the economy, which corresponds to having more than 9.35 battle-related deaths per million inhabitants during a year. Robust standard errors are in parentheses. $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.
Source: Authors' calculations based on GDP (constant 2015 US$) and population data from WB/WDI, and number of battle-related fatalities from UCDP/GED. The methodology for sample partitioning is described in Section 2.3.2.

Appendix C

Appendix for Chapter 3

C.1 Omitted proofs

C.1.1 Proof of Proposition 3.1

To proof existence and uniqueness I follow Angeletos et al. (2007) methodology.

Using the equilibrium Equation 3.3 and 3.6, we can create a function named $U(\hat{\theta}_k)$ s.t.:

$$U(\hat{\theta}_k) = \hat{\theta}_k - S_k \cdot \Phi\left(\frac{(\alpha-1)(\hat{\theta}_k - \mu_k) - \beta\Phi^{-1}\left(\frac{1-t}{1-p}\right)}{\sigma_x}\right). \tag{C.1}$$

To proof the function is monotonic in $\hat{\theta}_k$ and hence, the FP exist and ! we need that

$$\frac{\partial U(\cdot)}{\partial \hat{\theta}_k} = 1 - S_k \frac{1}{\sigma_x}\phi(\cdot)(\alpha-1) > 0 \tag{C.2}$$

To proof this derivative is positive, consider the maximum values of $\phi(\cdot) = 1/\sqrt{2\pi}$ and $S_k = 1$, then

$$\frac{\partial U(\cdot)}{\partial \hat{\theta}_1} > 0 \Rightarrow 1 - \frac{1}{\sigma_x}\frac{1}{\sqrt{2\pi}}\frac{\sigma_x^2}{\sigma_\mu^2} > 0 \Rightarrow \sigma_\mu^2 > \frac{\sigma_x}{\sqrt{2\pi}} \tag{C.3}$$

Therefore the last condition is both necessary and sufficient for $U(\cdot)$ to be monotonic on $\hat{\theta}_k$, in which case the monotone equilibrium is unique. Finally, to prove that this equilibrium is the only one surviving iterated deletion of strictly dominated strategies, see Morris and Shin (2001). ■

C.1.2 Proof of Corollary 3.1

To proof the second point we need to use implicit differentiation in Equation C.1 ,

$$\frac{\partial \hat{\theta}_k}{\partial \hat{x}_k} = -\frac{\partial U/\partial \hat{x}_k}{\partial U/\partial \hat{\theta}_k} = \frac{S_k \frac{\phi(\cdot)}{\sigma_x}}{1 - S_k \phi_1(\cdot)\frac{\alpha-1}{\sigma_x}} > 0 \tag{C.4}$$

The latter inequality comes from the condition for existence and uniqueness. Using Equation 3.6,

$$\frac{\partial \hat{x}_k^*}{\partial \hat{\theta}_k} = \alpha > 0 \tag{C.5}$$

Regarding the second point, using implicit differentiation in Equation C.1 ,

$$\frac{\partial \hat{\theta}_k}{\partial S_k} = -\frac{\partial U/\partial S_k}{\partial U/\partial \hat{\theta}_k} = \frac{\Phi(\cdot)}{1 - S_k \phi_1(\cdot)\frac{\alpha-1}{\sigma_x}} > 0 \tag{C.6}$$

The latter inequality comes from the condition for existence and uniqueness. Then, using Equation 3.6,

$$\frac{\partial \hat{x}_k^*}{\partial S_k} = \alpha \cdot \frac{\Phi(\cdot)}{1 - S_k \phi_1(\cdot)\frac{\alpha-1}{\sigma_x}} > 0 \tag{C.7}$$

Regarding the different tax rates,

$$\frac{\partial \hat{\theta}_k}{\partial t} = -\frac{\partial U/\partial t}{\partial U/\partial \hat{\theta}_k} = -\frac{S_k \frac{\phi(\cdot)}{\sigma_x} \beta \frac{\partial \Phi^{-1}(\cdot)}{\partial t}}{1 - S_k \phi_1(\cdot) \frac{\alpha - 1}{\sigma_x}} > 0 \quad \text{given that} \quad \frac{\partial \Phi^{-1}(\cdot)}{\partial t} < 0, \tag{C.8}$$

then

$$\frac{\partial \hat{x}_k}{\partial t} = \alpha \frac{\partial \hat{\theta}_k}{\partial t} - \beta \frac{\partial \Phi^{-1}(\cdot)}{\partial t} > 0. \tag{C.9}$$

W.r.t. p

$$\frac{\partial \hat{\theta}_k}{\partial p} = -\frac{\partial U/\partial p}{\partial U/\partial \hat{\theta}_k} = -\frac{S_k \frac{\phi(\cdot)}{\sigma_x} \beta \frac{\partial \Phi^{-1}(\cdot)}{\partial p}}{1 - S_k \phi_1(\cdot) \frac{\alpha - 1}{\sigma_x}} < 0 \quad \text{given that} \quad \frac{\partial \Phi^{-1}(\cdot)}{\partial p} > 0, \tag{C.10}$$

then

$$\frac{\partial \hat{x}_k}{\partial p} = \alpha \frac{\partial \hat{\theta}_k}{\partial p} - \beta \frac{\partial \Phi^{-1}(\cdot)}{\partial p} < 0. \tag{C.11}$$

Finally, regarding μ_k,

$$\frac{\partial \hat{\theta}_k}{\partial \mu_k} = -\frac{\partial U/\partial \mu_k}{\partial U/\partial \hat{\theta}_k} = -\frac{S_k \phi(\cdot) \frac{\alpha - 1}{\sigma_x}}{1 - S_k \phi_1(\cdot) \frac{\alpha - 1}{\sigma_x}} < 0, \tag{C.12}$$

then

$$\frac{\partial \hat{x}_k}{\partial \mu_k} = \alpha \frac{\partial \hat{\theta}_k}{\partial \mu_k} - (\alpha - 1) = -\alpha \frac{S_k \phi(\cdot) \frac{\alpha - 1}{\sigma_x}}{1 - S_k \phi(\cdot) \frac{\alpha - 1}{\sigma_x}} - (\alpha - 1) < 0. \tag{C.13}$$

■

C.1.3 Proof of Lemma 3.1

We need to show that $\partial \Pi_j(S_j|\mathcal{I}^1)/\partial S_j > 0$, i.e.,

$$\frac{\partial}{\partial S_j}(1-p)\Pr(x_{ij} \leq \hat{x}_j^*(S_j), \theta_j \leq \hat{\theta}_j(S_j)|\mathcal{I}^1) + \frac{\partial}{\partial S_j}(1-t)\Pr(x_{ij} > \hat{x}_j^*(S_j)|\mathcal{I}^1) > 0 \tag{C.14}$$

Rewriting the probability term,

$$\frac{\partial}{\partial S_j}(1-p)\Pr(x_{ij} \leq \hat{x}_j^*(S_j), \theta_j \leq \hat{\theta}_j(S_j)|\mathcal{I}^1) + \frac{\partial}{\partial S_j}(1-t)\left[1 - \Pr(x_{ij} \leq \hat{x}_j^*|\mathcal{I}^1)\right] > 0 \tag{C.15}$$

which simplifies further to:

$$(1-p)\frac{\partial}{\partial S_j}\Pr(x_{ij} \leq \hat{x}_j^*(S_j), \theta_j \leq \hat{\theta}_j(S_j)|\mathcal{I}^1) - (1-t)\frac{\partial}{\partial S_j}\Pr(x_{ij} \leq \hat{x}_j^*(S_j)|\mathcal{I}^1) > 0 \tag{C.16}$$

The joint distribution of x_{ij} and θ_j follows a bivariate normal distribution:

$$f_{x_{ij},\theta_j}(x_{ij}, \theta_j) \sim \text{Bivariate Normal with } \boldsymbol{\mu} = \begin{pmatrix} \mu_j \\ \mu_j \end{pmatrix}, \boldsymbol{\Sigma} = \begin{pmatrix} \sigma_x^2 + \sigma_\mu^2 & \sigma_\mu^2 \\ \sigma_\mu^2 & \sigma_\mu^2 \end{pmatrix}, \rho = \frac{\sigma_\mu}{\sqrt{\sigma_x^2 + \sigma_\mu^2}}; \tag{C.17}$$

The corresponding cumulative distribution function (CDF) w.r.t to the threshold $\hat{x}_j$ and $\hat{\theta}_j$ is

$$\begin{aligned} F_{x_{ij},\theta_j}(\hat{x}_j,\hat{\theta}_j) &= \int_{-\infty}^{\hat{x}_j}\int_{-\infty}^{\hat{\theta}_j} \frac{1}{2\pi\sigma_\mu\sqrt{\sigma_x^2+\sigma_\mu^2}\sqrt{1-\rho^2}} e^{-\frac{z}{2(1-\rho^2)}} d\theta_1 dx_{ij} \\ z &= \left(\frac{x_{ij}-\mu_j}{\sqrt{\sigma_x^2+\sigma_\mu^2}}\right)^2 - \frac{2\rho(x_{ij}-\mu_j)(\theta_j-\mu_j)}{\sigma_\mu\sqrt{\sigma_x^2+\sigma_\mu^2}} + \left(\frac{\theta_j-\mu_j}{\sigma_\mu}\right)^2 \end{aligned} \tag{C.18}$$

Applying Leibniz's rule, the derivative of $\Pr(x_{ij} \leq \hat{x}_j^*(S_j)|\mathcal{I}^1)$ is:

$$\frac{1}{\sqrt{2\pi(\sigma_x^2+\sigma_\mu^2)}} e^{-\frac{1}{2}\left(\frac{\hat{x}_j^*-\mu_j}{\sqrt{\sigma_x^2+\sigma_\mu^2}}\right)^2} \frac{\partial \hat{x}_j^*}{\partial S_j} \tag{C.19}$$

Similarly, the derivative of $\Pr(x_{ij} \leq \hat{x}_j^*(S_j), \theta_j \leq \hat{\theta}_j(S_j)|\mathcal{I}^1)$

$$\int_{-\infty}^{\hat{\theta}_j} f_{x_{ij},\theta_j}(\hat{x}_j,\theta_j) d\theta_j \frac{\partial \hat{x}_j^*}{\partial S_j} + \int_{-\infty}^{\hat{x}_j} f_{x_{ij},\theta_j}(x_{ij},\hat{\theta}_j) dx_{ij} \frac{\partial \hat{\theta}_j}{\partial S_j} \tag{C.20}$$

Substituting the bivariate normal distribution, we obtain:

$$\begin{aligned} &\int_{-\infty}^{\hat{\theta}_j} \frac{1}{2\pi\sigma_\mu\sqrt{\sigma_x^2+\sigma_\mu^2}\sqrt{1-\rho^2}} e^{\left(-\frac{1}{2(1-\rho^2)}\left(\left(\frac{\hat{x}_j-\mu_j}{\sqrt{\sigma_x^2+\sigma_\mu^2}}\right)^2 - \frac{2\rho(\hat{x}_j-\mu_j)(\theta_j-\mu_j)}{\sigma_\mu\sqrt{\sigma_x^2+\sigma_\mu^2}} + \left(\frac{\theta_j-\mu_j}{\sigma_\mu}\right)^2\right)\right)} d\theta_j \frac{\partial \hat{x}_j^*}{\partial S_j} \\ &+ \int_{-\infty}^{\hat{x}_j} \frac{1}{2\pi\sigma_\mu\sqrt{\sigma_x^2+\sigma_\mu^2}\sqrt{1-\rho^2}} e^{\left(-\frac{1}{2(1-\rho^2)}\left(\left(\frac{x_{ij}-\mu_j}{\sqrt{\sigma_x^2+\sigma_\mu^2}}\right)^2 - \frac{2\rho(x_{ij}-\mu_j)(\hat{\theta}_j-\mu_j)}{\sigma_\mu\sqrt{\sigma_x^2+\sigma_\mu^2}} + \left(\frac{\hat{\theta}_j-\mu_j}{\sigma_\mu}\right)^2\right)\right)} dx_{ij} \frac{\partial \hat{\theta}_j}{\partial S_j} \end{aligned} \tag{C.21}$$

Then, Equation C.16 becomes

$$\begin{aligned} &(1-p)\Bigg[\int_{-\infty}^{\hat{\theta}_j} \frac{1}{2\pi\sigma_\mu\sqrt{\sigma_x^2+\sigma_\mu^2}\sqrt{1-\rho^2}} e^{\left(-\frac{1}{2(1-\rho^2)}\left(\left(\frac{\hat{x}_j-\mu_j}{\sqrt{\sigma_x^2+\sigma_\mu^2}}\right)^2 - \frac{2\rho(\hat{x}_j-\mu_j)(\theta_j-\mu_j)}{\sigma_\mu\sqrt{\sigma_x^2+\sigma_\mu^2}} + \left(\frac{\theta_j-\mu_j}{\sigma_\mu}\right)^2\right)\right)} d\theta_j \frac{\partial \hat{x}_j^*}{\partial S_j} \\ &+ \int_{-\infty}^{\hat{x}_j} \frac{1}{2\pi\sigma_\mu\sqrt{\sigma_x^2+\sigma_\mu^2}\sqrt{1-\rho^2}} e^{\left(-\frac{1}{2(1-\rho^2)}\left(\left(\frac{x_{ij}-\mu_j}{\sqrt{\sigma_x^2+\sigma_\mu^2}}\right)^2 - \frac{2\rho(x_{ij}-\mu_j)(\hat{\theta}_j-\mu_j)}{\sigma_\mu\sqrt{\sigma_x^2+\sigma_\mu^2}} + \left(\frac{\hat{\theta}_j-\mu_j}{\sigma_\mu}\right)^2\right)\right)} dx_{ij} \frac{\partial \hat{\theta}_j}{\partial S_j}\Bigg] \\ &- (1-t)\frac{1}{\sqrt{2\pi(\sigma_x^2+\sigma_\mu^2)}} e^{-\frac{1}{2}\left(\frac{\hat{x}_j^*-\mu_j}{\sqrt{\sigma_x^2+\sigma_\mu^2}}\right)^2} \frac{\partial \hat{x}_j^*}{\partial S_j} > 0 \end{aligned} \tag{C.22}$$

Considering $\hat{x}_j^* \propto \alpha\hat{\theta}_j$, and eliminating common terms ($\frac{1}{\sqrt{\sigma_x^2+\sigma^2\mu}}$), we obtain:

$$\begin{aligned}
&(1-p)\frac{1}{2\pi\sigma_\mu\sqrt{1-\rho^2}}\Bigg[\alpha\int_{-\infty}^{\hat{\theta}_j} e^{\left(-\frac{1}{2(1-\rho^2)}\left(\left(\frac{\hat{x}_j-\mu_j}{\sqrt{\sigma_x^2+\sigma_\mu^2}}\right)^2-\frac{2\rho(\hat{x}_j-\mu_j)(\theta_j-\mu_j)}{\sigma_\mu\sqrt{\sigma_x^2+\sigma_\mu^2}}+\left(\frac{\theta_j-\mu_j}{\sigma_\mu}\right)^2\right)\right)}d\theta_j \\
&+\int_{-\infty}^{\hat{x}_j} e^{\left(-\frac{1}{2(1-\rho^2)}\left(\left(\frac{x_{ij}-\mu_j}{\sqrt{\sigma_x^2+\sigma_\mu^2}}\right)^2-\frac{2\rho(x_{ij}-\mu_j)(\hat{\theta}_j-\mu_j)}{\sigma_\mu\sqrt{\sigma_x^2+\sigma_\mu^2}}+\left(\frac{\hat{\theta}_j-\mu_j}{\sigma_\mu}\right)^2\right)\right)}dx_{ij}\Bigg] \\
&-(1-t)\frac{1}{\sqrt{2\pi}}\alpha e^{-\frac{1}{2}\left(\frac{\hat{x}_j^*-\mu_j}{\sqrt{\sigma_x^2+\sigma_\mu^2}}\right)^2}>0
\end{aligned} \tag{C.23}$$

Substituting ρ and multiplying both sides by $\sqrt{2\pi}$, we get:

$$\begin{aligned}
&(1-p)\frac{\sqrt{\sigma_x^2+\sigma_\mu^2}}{\sqrt{2\pi}\sigma_\mu\sigma_x}\Bigg[\alpha\int_{-\infty}^{\hat{\theta}_j} e^{\left(-\frac{1}{2(1-\rho^2)}\left(\left(\frac{\hat{x}_j-\mu_j}{\sqrt{\sigma_x^2+\sigma_\mu^2}}\right)^2-\frac{2\rho(\hat{x}_j-\mu_j)(\theta_j-\mu_j)}{\sigma_\mu\sqrt{\sigma_x^2+\sigma_\mu^2}}+\left(\frac{\theta_j-\mu_j}{\sigma_\mu}\right)^2\right)\right)}d\theta_j \\
&+\int_{-\infty}^{\hat{x}_j} e^{\left(-\frac{1}{2(1-\rho^2)}\left(\left(\frac{x_{ij}-\mu_j}{\sqrt{\sigma_x^2+\sigma_\mu^2}}\right)^2-\frac{2\rho(x_{ij}-\mu_j)(\hat{\theta}_j-\mu_j)}{\sigma_\mu\sqrt{\sigma_x^2+\sigma_\mu^2}}+\left(\frac{\hat{\theta}_j-\mu_j}{\sigma_\mu}\right)^2\right)\right)}dx_{ij}\Bigg] \\
&-(1-t)\alpha e^{-\frac{1}{2}\left(\frac{\hat{x}_j^*-\mu_j}{\sqrt{\sigma_x^2+\sigma_\mu^2}}\right)^2}>0
\end{aligned} \tag{C.24}$$

Using $\sqrt{\alpha}=\frac{\sqrt{\sigma_x^2+\sigma_\mu^2}}{\sigma_\mu}$, we rewrite:

$$\begin{aligned}
&(1-p)\frac{1}{\sqrt{2\pi}\sqrt{\alpha}\sigma_x}\Bigg[\alpha\int_{-\infty}^{\hat{\theta}_j} e^{\left(-\frac{1}{2(1-\rho^2)}\left(\left(\frac{\hat{x}_j-\mu_j}{\sqrt{\sigma_x^2+\sigma_\mu^2}}\right)^2-\frac{2\rho(\hat{x}_j-\mu_j)(\theta_j-\mu_j)}{\sigma_\mu\sqrt{\sigma_x^2+\sigma_\mu^2}}+\left(\frac{\theta_j-\mu_j}{\sigma_\mu}\right)^2\right)\right)}d\theta_j \\
&+\int_{-\infty}^{\hat{x}_j} e^{\left(-\frac{1}{2(1-\rho^2)}\left(\left(\frac{x_{ij}-\mu_j}{\sqrt{\sigma_x^2+\sigma_\mu^2}}\right)^2-\frac{2\rho(x_{ij}-\mu_j)(\hat{\theta}_j-\mu_j)}{\sigma_\mu\sqrt{\sigma_x^2+\sigma_\mu^2}}+\left(\frac{\hat{\theta}_j-\mu_j}{\sigma_\mu}\right)^2\right)\right)}dx_{ij}\Bigg] \\
&-(1-t)e^{-\frac{1}{2}\left(\frac{\hat{x}_j^*-\mu_j}{\sqrt{\sigma_x^2+\sigma_\mu^2}}\right)^2}>0
\end{aligned} \tag{C.25}$$

Since all terms are positive, we simplify by disregarding the second integral, adopting a sufficient condition. The second integral "helps" in proving the statement. Extracting the leading exponent term further refines the condition to:

$$\begin{aligned}
&(1-p)\frac{\sqrt{\sigma_x^2+\sigma_\mu^2}}{\sqrt{2\pi}\sigma_\mu\sigma_x}e^{\left(-\frac{1}{2(1-\rho^2)}\right)}\int_{-\infty}^{\hat{\theta}_j} e^{\left(-\frac{1}{2(1-\rho^2)}\left(-\frac{2\rho(\hat{x}_j-\mu_j)(\theta_j-\mu_j)}{\sigma_\mu\sqrt{\sigma_x^2+\sigma_\mu^2}}+\left(\frac{\theta_j-\mu_j}{\sigma_\mu}\right)^2\right)\right)}d\theta_j \\
&-(1-t)>0
\end{aligned} \tag{C.26}$$

Since the coefficient of $(1-p)$ is positive, the condition holds when the difference between p and t is sufficiently large. For instance, setting $t=1$ ensures the condition is met for all $p>0$. ■

C.1.4 Proof of Proposition 3.2

To prove that these strategies constitute a Perfect Bayesian Nash Equilibrium (PBNE), we must verify two conditions: Stage 2 optimality and Stage 1 optimality. Stage 2 optimality has already been established in Proposition 1. We now focus on Stage 1.

Since all strategies considered are symmetric, we analyze the case where all agents follow the same strategy. Rather than expressing expected payoffs as functions of the thresholds—which themselves depend on specialization proportions and public signals—it is more convenient to define expected payoffs directly as a function of $(S_1, S_2, \mu_1, \mu_2) : \Pi_{ij}(S_j, \mu_j \mid \mathcal{I}^1) = \Pi_{ij}(\hat{x}_j(S_j, \mu_j), \hat{\theta}_j(S_j, \mu_j) \mid \mathcal{I}^1)$.

Assume all agents follow Strategy 1, which prescribes specializing in the tax haven with the lower public signal, with ties broken in favor of TH_1. Then:

- If $\mu_1 < \mu_2$, all agents specialize in TH_1, so $(S_1 = 1, S_2 = 0)$. By Lemma 1, $\Pi_{i1}(1, \mu_1 \mid \mathcal{I}^1) > \Pi_{i2}(0, \mu_2 \mid \mathcal{I}^1), \forall i$. No agent has an incentive to deviate and $(S_1 = 1, S_2 = 0)$ is consistent.
- If $\mu_1 = \mu_2$, tie-breaking leads all agents to specialize in TH_1, again yielding $(S_1 = 1, S_2 = 0)$. By Lemma 1, $\Pi_{i1}(1, \mu_1 \mid \mathcal{I}^1) > \Pi_{i2}(0, \mu_1 \mid \mathcal{I}^1), \forall i$. No agent has an incentive to deviate and $(S_1 = 1, S_2 = 0)$ is consistent.
- If $\mu_1 > \mu_2$, all agents specialize in TH_2, so $(S_1 = 0, S_2 = 1)$. By Lemma 1, $\Pi_{i1}(0, \mu_1 \mid \mathcal{I}^1) < \Pi_{i2}(1, \mu_2 \mid \mathcal{I}^1), \forall i$. No agent has an incentive to deviate and $(S_1 = 0, S_2 = 1)$ is consistent.

Thus, Strategy 1 constitutes a Perfect Bayesian Nash Equilibrium.

Assume all agents follow Strategy 2, which prescribes specializing in the tax haven with the lower public signal, and mixing uniformly when the signals are equal. Let $\sigma_i(j)$ denote the probability of agent i specializing in tax haven j.

- If $\mu_1 < \mu_2$, all agents specialize in TH_1, so $(S_1 = 1, S_2 = 0)$. By Lemma 1, $\Pi_{i1}(1, \mu_1 \mid \mathcal{I}^1) > \Pi_{i2}(0, \mu_2 \mid \mathcal{I}^1), \forall i$. No agent has an incentive to deviate and $(S_1 = 1, S_2 = 0)$ is consistent.
- If $\mu_1 > \mu_2$, all agents specialize in TH_2, so $(S_1 = 0, S_2 = 1)$. By Lemma 1, $\Pi_{i1}(0, \mu_1 \mid \mathcal{I}^1) < \Pi_{i2}(1, \mu_2 \mid \mathcal{I}^1), \forall i$. No agent has an incentive to deviate and $(S_1 = 0, S_2 = 1)$ is consistent.
- If $\mu_1 = \mu_2$, all agents mix uniformly: $\sigma_i(1) = \sigma_i(2) = 0.5$, resulting in $(S_1 = 0.5, S_2 = 0.5)$. By symmetry,
$$\Pi_{i1}(0.5, \mu_1 \mid \mathcal{I}^1) = \Pi_{i2}(0.5, \mu_2 \mid \mathcal{I}^1), \quad \forall i.$$
Agents are indifferent and have no incentive to deviate and $(S_1 = 0.5, S_2 = 0.5)$ is consistent.

Thus, Strategy 2 also constitutes a Perfect Bayesian Nash Equilibrium.

■

C.1.5 Proof of Proposition 3.4

We analyze the effect of increasing μ_1 on equilibrium evasion $\tilde{A}$. By symmetry, the same logic applies to increases in μ_2 when μ_1 is held constant.

According to Equation 3.14, the effect of increasing μ_1 varies across three domains:

- **Case 1:** $\mu_1 < \mu_2$
  In this region, all agents specialize in TH_1, so $S_1^* = 1$ and $S_2^* = 0$. Since specialization remains fixed as μ_1 increases reducing evasion. This proves (i).

- **Case 2:** $\mu_1 = \mu_2$
  Under Strategy 1, agents coordinate fully on TH_1, resulting in:

$$\tilde{A}_{\text{Strat 1}} = \Phi\left(\frac{\hat{x}_1^*(1) - \mu}{\sqrt{\sigma_x^2 + \sigma_\mu^2}}\right) \tag{C.27}$$

  Under Strategy 2, agents split evenly across the two havens:

$$\tilde{A}_{\text{Strat 2}} = 0.5 \cdot \Phi\left(\frac{\hat{x}_1^*(0.5) - \mu}{\sqrt{\sigma_x^2 + \sigma_\mu^2}}\right) + 0.5 \cdot \Phi\left(\frac{\hat{x}_2^*(0.5) - \mu}{\sqrt{\sigma_x^2 + \sigma_\mu^2}}\right) \tag{C.28}$$

  By symmetry, $\hat{x}_1^*(0.5) = \hat{x}_2^*(0.5)$, implying:

$$\tilde{A}_{\text{Strat 2}} = \Phi\left(\frac{\hat{x}_1^*(0.5) - \mu}{\sqrt{\sigma_x^2 + \sigma_\mu^2}}\right) \tag{C.29}$$

  Since $\hat{x}_1^*(1) > \hat{x}_1^*(0.5)$ (by Lemma 1), we conclude that $\tilde{A}_{\text{Strat 1}} > \tilde{A}_{\text{Strat 2}}$, proving part (ii). However, note that a marginal increase in μ_1 beyond the symmetry point immediately triggers a switch back to full specialization under Strategy 2. As a result, evasion jumps back to the level of Strategy 1, reversing the previous decline. This discontinuous increase confirms part (iii).

- **Case 3:** $\mu_1 > \mu_2$
  Once all agents are specialized in TH_2, further increases in μ_1 have no effect on specialization or evasion, as no one responds to TH_1's signal anymore. This proves part (iv).

■

Bibliography

Angeletos, George-Marios, Christian Hellwig, and Alessandro Pavan. 2007. "Dynamic global games of regime change: Learning, multiplicity, and the timing of attacks." *Econometrica* 75 (3): 711–756.

Montalvo, José G, and Marta Reynal-Querol. 2005. "Ethnic polarization, potential conflict, and civil wars." *American Economic Review* 95 (3): 796–816. Data retrieved from http://www.econ.upf.edu/~reynal/data_web.htm.

Morris, Stephen, and Hyun Song Shin. 2001. "Global Games: Theory and Applications." Cowles Foundation Discussion Paper No. 1275R, no. 1275R (August).

Nunn, Nathan, and Diego Puga. 2012. "Ruggedness: The blessing of bad geography in Africa." *Review of Economics and Statistics* 94 (1): 20–36. Data retrieved from http://diegopuga.org/data/rugged/.