The Unforeseeable? Assessing the Relevance of Explanatory Factors in Forecasting Terrorist Threats

von Lampe, Bjarne

THE UNFORESEEABLE?

Assessing the Relevance of Explanatory Factors in Forecasting Terrorist Threats

Abstract

Can terrorist threats be forecasted in a systematic way? Which variables help to do so in the most accurate way? The present study examines the relative importance of features when building forecasting models on terrorist threats. To do so, it draws on both academic literature and publications by counterterrorist practitioners. This study addresses three key gaps in existing research. Specifically, it allows for comparing the utility of different theoretical models to each other, it puts an explicit focus on machine learning-based forecasting with out-of-sample performance metrics, and it explicitly aims to incorporate knowledge from the practitioner sector, which is understandably less open about their work than the academic community but has still produced several insightful publications on the topic of forecasting terrorist threats. The outcomes of the analysis do not confirm the expectation that variables of interest to both academics and practitioners would have the highest predictive power. Rather, it is the population of a country that scores highest, followed by GDP, data on weapon flows into the country, and religious fragmentation in models with no features based on lagged versions of the outcome variable. In models including such variables, the lag of the terrorist attack occurrence consistently scores second highest, and these models consistently out-perform their counterparts missing these variables. The results obtained in this paper are arguably of most use to academic research, in that they add onto a so far relatively limited body of work on out-of-sample forecasting and provide insight into the relative predictive power of existing theoretical models. Practitioners may be more interested in the methodological approach taken in this piece, which can be of use to them when evaluating the priority list of warning indicators to take into consideration when assessing the severity of terrorist threats.

Universiteit Leiden
The Netherlands

Bjarne von Lampe
b.n.von.lampe@umail.leidenuniv.nl

Supervisor: Dr. S. RezaeeDaryakenari
Second reader: Dr. M. DiGiuseppe

Word count: 10,000

Table of Contents

Introduction

Be it in Paris (Balboni, 2016; Hugues, 2015), Munich (Augsburger Allgemeine, 2016), Bamako (Février, 2015), or London (Erickson & Stanley-Becker, 2021): terrorist attacks across the world are often considered "unforeseeable" or "unpredictable" in the aftermath. Nevertheless, it is one of the core tasks of national security institutions to forecast threats and vulnerabilities to society, including this very type of violence. The forecasting of terrorist threats is also of academic interest, but despite various scholarly attempts at this task, these have usually fallen short of expectations (E. Bakker, 2012).

The main objective of the present paper is to evaluate the relative importance of various variables in forecasting such threats. The variables in question are inspired by multiple concurrent theoretical models on the causal relationships relating to the occurrence of terrorism in academic works. One cluster of such variables concerns country-specific data on the political circumstances, such as the level of democratisation and the type of elections, if any. This is supplemented by data focusing on the ability of a regime to consolidate its position of power, such as a measure of state fragility and the durability of the regime in place. Another cluster of variables relates to the economic situation – such as its GDP, inflation rate, and economic inequality – in the country, which is commonly cited in academic research as affecting the threat of terrorist attacks, though not uncontroversially so. Additionally, this paper incorporates factors that were identified as being relevant to counterterrorism practitioners, who are more familiar with the specific task of forecasting threats due to the nature of their work. This includes data on weapon movement into the country and internet usage, which is used to approximate social media usage rates and communication capabilities in the country. Furthermore, ties with the United States, which may be relevant to both the ability to provide security and to a country's ideological attractiveness as a target of terrorist attacks, and time-lagged information on the previous occurrence of terrorist attacks in the country in question are also incorporated based on insights gained about the work of practitioners.

The paper establishes forecasting models for terrorist threats based on three different algorithms. Logistic regression, frequently used in quantitative studies in political science, is used as a baseline, while the principal focus lies on random forest models, which are complemented by gradient boosting machines. The models are evaluated based on their out-of-sample forecasting performance and subsequently analysed on the relative importance of variables in the models to answer the research question:

RQ: Which variables are the most important to forecasting the threat of terrorist attacks?

This piece is structured as follows. First, a literature review provides a brief overview over related publications and the gaps it identifies therein. It then briefly explains how this study will work towards bridging each of these gaps. The theory section then provides a theoretical justification for the many features that are used in this paper, be they inspired by publications in the academic or professional sector or reflected in both bodies of work. The data section lists and explains the various data sources and variable formats used to feed the models developed in this paper. It also provides a brief discussion where multiple data sources are available. The methodological approach used in this paper, including the model types and an explanation of out-of-sample performance and variable importance metrics, is then outlined in the methods section. The results and discussion section provides the outcomes of the analysis carried out and relates some of the most interesting insights back to the pre-existing work on terrorism research. Finally, the conclusion briefly summarizes key points made in the previous sections and provides a discussion of the main contribution of this paper to academic research, as well as potentially to practitioner work. It also suggests avenues for future research that may build onto the methodological design and results of this study.

Literature review

This paper builds onto a large body of terrorism research, in which it identifies, and consequently attempts to bridge, three main gaps. Firstly, there has been ample research into causal factors and drivers of terrorism. This body of work has resulted in various theoretical frameworks differing in the level of analysis- ranging from the individual to the nation- and in the causal factors taken into consideration. The results of this research, however, are often left to stand on their own instead of being compared to competing explanatory models. One of the major fields of research into drivers of terrorism is psychological in nature and often focuses on micro-level processes (at the scale of individuals) (Abbasi & Khatwani, 2018; McCauley & Moskalenko, 2008). Others have focused on the political context in which terrorism does or does not occur (Chenoweth, 2013; George, 2018; Piazza, 2007; Tikuisis, 2009), while yet another cluster of research examines the relationship between terrorism and economic factors such as globalisation (Bergesen & Lizardo, 2004), income and poverty (Enders et al., 2016; Krueger & Malečková, 2003b), inequality (Nabin et al., 2022), and demographics (Coccia, 2018). Other research still has explored the importance of socio-cultural aspects and phenomena such as minority discrimination (Piazza, 2012), religion (Stern, 2009), conspiracy theories (Douglas et al., 2019), and societal characteristics such as the spectrum of individualism and collectivism (Kluch & Vaux, 2017). While many such studies have yielded valuable results, they are rarely used in combination with each other, which is something this piece aims to address by assembling a model based on a variety of features taken from fields of research named above.

Secondly, there is an extensive existing body of quantitative research in terrorism studies. However, forecasting-related studies using non-parametric algorithms and, more importantly, using out-of-sample metrics to assess the forecasting abilities of the developed models, are few in-between the many studies into drivers and causal relationships affecting dynamics of terrorism. The latter is certainly vital work and "does much to shed light on the factors that may make a state more likely to suffer […] terrorist attacks" (Desmarais & Cranmer, 2013, p. 2). Nevertheless, the relative lack of studies focused on out-of-sample performance represents a gap in the existing literature. In one of the existing exceptions to this pattern, Bakker et al. (2014) find the following:

> *Most research tends to overfocus on hypothesis testing to the exclusion of predictive power. That is, while the move away from maximizing measures of in-sample fit (e.g., "explained variance") has been fruitful, a complete lack of attention to how well a model fits the data is lamentable.* (p. 53)

This issue is addressed by constructing models using non-parametric algorithms that have proven useful in forecasting complex social phenomena. As will be explained in more detail in the methodology, the models themselves will be evaluated based exclusively on their out-of-sample forecasting ability, representing an ability to deal with data that has not been seen before, as would be the case in a practical application of such models. There is a growing body of research on out-of-sample forecasting and early warning systems in the related field of conflict studies (see e.g. Hegre et al., 2017; Pinckney & RezaeeDaryakenari, 2022; Uppsala University, n.d.), which in parts serves as a methodological inspiration for the present study.

Lastly, terrorism research projects and publications tend to build onto each other but fail to incorporate insights provided by the professional sector- for example, publications by intelligence agencies and government officials working in counterterrorism. While the intelligence sector is naturally less open about their work and the knowledge produced through it than the academic world, the few publications from this field that do exist should not be discounted and instead incorporated into academic research.

The call for more systematic approaches in forecasting terrorism voiced by Bakker (2012) is shared in the professional sector, as can be seen in Khalsa's (2004) book detailing a methodology for the formidable task of transforming large amounts of information into reliable and actionable threat assessments. While representing a significant improvements over prior approaches as described by the author in terms of rendering the process more systematic and less prone to individual biases, it still relies on analysts' "intuition" (Khalsa, 2004, p. 12) when faced with a body of information that is largely qualitative in nature.

This paper aims to bridge the gap between the academic and professional sectors by incorporating several factors identified by professional publications as being useful in forecasting terrorist threats into the models and comparing their forecasting value to that of the features identified from academic literature. Combining insights from the two benefits both academia and professionals: the academic sector can benefit from incorporating an as-of-yet underutilized body of expertise, especially in terms of forecasting, and the professional sector can make use of theories and models developed in academia and incorporate features empirically identified as being of use when forecasting terrorist threats into their methodologies.

Theory

One of the main clusters of explanatory variables found in academic literature is of political nature, with a focus on the structure of the state in question. A fragile state system, for example, may provide terrorist groups with the opportunity to operate without being countered by an effective security apparatus and therefore be a facilitating factor for terrorist attacks (George, 2018; Tikuisis, 2009). Similarly, low regime durability can be an indicator of a general incapability to carry out key government tasks, including the provision of security, and has been found to have a positive effect on the occurrence of terrorism (Ajide, 2021). Furthermore, the type of electoral system (if any) in a country is also found to impact the likelihood of terrorist attacks being committed, with proportional representation systems being less frequently targeted by international terrorism than other systems (Li, 2005).

While there is dissent over the exact nature of the relationship, academics also frequently identify the level of democracy as a predictor for terrorist threats. Eyerman (1998, as cited in Piazza, 2007) identifies two rivalling schools of thought on the relationship between these two factors, arguing that democracy either increases or decreases the likelihood of terrorist attacks. Scholars of the first group, referred to as the “political access” school (Piazza, 2007, p. 523), argue that democracies provide “multiple avenues by which actors can advance their political agendas [and thus] increase the utility of legal political activity for all political actors, including terrorists” (p. 523). In other words, by creating options for legitimate political participation, democracies provide an alternative to terrorism and thereby reduce the incentives for engaging in the latter. Opposing scholars from the “strategic school” (Piazza, 2007, p. 523) argue that democracies are less capable of performing tasks related to surveillance and counterterrorism due to their obligation to uphold and protect human and civil rights: “these same restrictions of executive and police power that are features of democracy also make democratic countries good hosts for terrorist groups” (p. 523). Furthermore, democracies may also be more appealing targets for terrorism according to this view, as democratic governments would be more inclined to negotiate with terrorist groups due to fearing unpopularity and eventual electoral defeat in case terrorist activity does not cease.

Other scholars yet have suggested that the relationship may not be a simple linear one (Li, 2005), but rather curvilinear (often represented by using the square of the level of democracy as an independent variable) or dependent on multiple sub-factors such as democratic participation, civil and political freedoms, and government restrictions thereon, and that the likelihood for terrorist attacks may be highest for states in the middle of the democratisation process (Chenoweth, 2013). While testing theories derived from the various schools of thought on this subject against one another lies beyond the scope of this paper, the large body of ongoing research into the relationship between democracy and terrorism legitimises the inclusion of features related to the former in the models to follow.

Economic factors may also play a role, as suggested by Krueger and Malečková (2003b). Inequality, for example, is associated with a higher frequency of terrorist attacks and a higher number of casualties (injuries and fatalities) from said attacks, as noted by Fleming et al. (2022), who study the effect of inequality within ethnic groups on these metrics of terrorist activity. Furthermore, high inflation rates and volatility thereof have been found to lead to increased threats of terrorist attacks in Pakistan (Ismail & Amjad, 2014; Shahbaz, 2013) and Africa (Ajide & Alimi, 2023). The importance of economic factors is not uncontroversial within this field, however, with other scholars noting that it may be commonly overestimated (Krueger & Malečková, 2003a; Piazza, 2006).

Furthermore, the effect of education on terrorist activities is a particularly interesting one. Brockhoff et al. (2012) find that the relationship is complex in that education has a positive effect on the occurrence of terrorist attacks in countries with low levels of socioeconomic development, while it has the opposite effect of decreasing terrorist activity in countries with more favourable structural circumstances. Korotayev et al. (2021) find a similar non-linear relationship, where an increase in schooling in countries with low levels of educational development is positively associated with terrorist activity, while an increase in schooling in countries with an already high level of educational development has a significant negative impact on terrorism. A similar, curvilinear relationship is also found by Elbakidze and Jin (2015), who study participation in transnational terrorism.

Finally, demographic aspects can also be related to the occurrence of terrorism. Coccia (2018) finds a significant relationship between population growth and terrorism, writing that high levels of population growth, particularly when combined with poor socioeconomic conditions, is conducive to terrorist activities. Going further, in his attempt to examine the relationship (or lack thereof) between economic factors and terrorism, Piazza (2006) finds that population is one of the most important predictors of both the incidence and the casualty rate of terrorism.

Practitioner publications show some overlap, but also some differences in variables that appear. The different variables will be discussed first, with the overlapping ones following at the end of this section. While previous terrorist attacks are of limited interest to academics due to them not explaining the root causes that lie at the origin of terrorist threats, information on such past events are used by practitioners to inform forecasts (Sinai, 2002).

Khalsa (2004) identifies cooperation with and ties to the United States as being relevant to the forecasting of terrorist threats in that it is estimated to affect the provision of security in that state's territory. While this view is highly US-centric (unsurprisingly so, coming from an American intelligence professional), another factor legitimising the inclusion of this variable is the framing of terrorism as going against supposed US imperialism commonly found in Islamist terrorist groups, particularly those originating in the Middle East, where the United States have a history of interventionism (Paz, 2003). This again singles out the United States as a factor of interest in forecasting terrorism and justifies including ties to this country in this paper.

The US Homeland Security Advisory Council identifies the internet as another indicator to be watched, as it "has become a major facilitator of terrorist activities, especially the spread of jihadist ideology" (2007, p. 4). They also point out the more general effect of access to technology in allowing terrorist groups to build up increased capability to carry out attacks, motivating this variable's inclusion in the models. Relatedly, the Dutch agency for coordinating counter-terrorist activities NCTV (2022) warns that the spread of conspiracy theories plays an increasingly threatening role in driving potential terrorist tendencies. Again, usage of the internet is a facilitator for this mechanism, as conspiracy theories are now frequently spread over social media (Douglas et al., 2019).

Some variables are used in both academic and professional publications. A notable example is the mention of "longer-term changes [and] trends" (Piazza, 2007, p. 523) in terrorist violence on a global scale. Furthermore, foreign military interventions can harden sentiments against a specific country, making them a potential driving factor in transnational terrorism (Chenoweth, 2013; Homeland Security Advisory Council, 2007). Religious fragmentation and grievances can also significantly drive the likelihood of terrorist attacks, as pointed out by practitioners (Homeland Security Advisory Council, 2007) and scholars (Piazza, 2006, 2012) alike. This is perhaps unsurprising given the recent politicisation and securitisation of religion in views of its role in terrorism by groups such as Al-Qaeda and ISIL, but it should be noted that this fragmentation is not necessarily connected with any particular religious group, or that it is religion itself that is identified as a causal factor for terrorism. Rather, it is the potential alienation of minorities that may lead to the grievances mentioned above. Lastly, globalisation is thought to profoundly impact the nature and feasibility of terrorist activities, as it promotes both communication of ideas and the circulation of people and goods, including potentially lethal material (Bergesen & Lizardo, 2004; Homeland Security Advisory Council, 2007).

As to the question on predictive power, the expectation is that predictors reflected in both bodies of literature will tend to have higher explanatory power than those described by only academics or only practitioners.

Data

The dependent variable for this project is the occurrence of terrorist attacks per country-year as a dichotomous variable, coded 1 if one or more attacks occurred and 0 otherwise. To avoid data leakage, the model is created in such a manner that it attempts to predict the occurrence (or lack thereof) of a terrorist attack in the year $t + 1$ based on feature data from year t. Terrorist attacks are defined following the definition used by the University of Maryland in the Global Terrorism Database methodology: "the threatened or actual use of illegal force and violence by a nonstate actor to attain a political, economic, religious, or social goal through fear, coercion, or intimidation" (2021a, p. 11). They further identify three key criteria, of which at least two must be present for an incident to be considered an act of terrorism: "The incident must be intentional […] The incident must entail some level of violence or immediate threat of violence [… and] the perpetrators of the incidents must be sub-national [non-governmental] actors" (p. 12). While there is no single agreed-upon definition of terrorism, these criteria reflect some of the key aspects generally identified and discussed in terrorism literature (Schmid, 2004). This definitional scope is also what motivates the choice of the Global Terrorism Database (START, 2021b) as the source for data on the dependent variable, as opposed to other options. The RAND Database of Worldwide Terrorism Incidents (RAND National Security Research Division, n.d.), for example, only contains data on domestic terrorism starting at the year 1998, as their methodology was only then updated to include these incidents in addition to transnational terrorism. The Center for Systemic Peace (2022) offers another option with their High Casualty Terrorist Bombings dataset, but this would represent a limitation that finds little justification in theoretical arguments. Shootings, stabbings, and vehicle ramming attacks, for example, would be excluded entirely from the study in the case this dataset and its associated operationalisation of terrorist attacks were used.

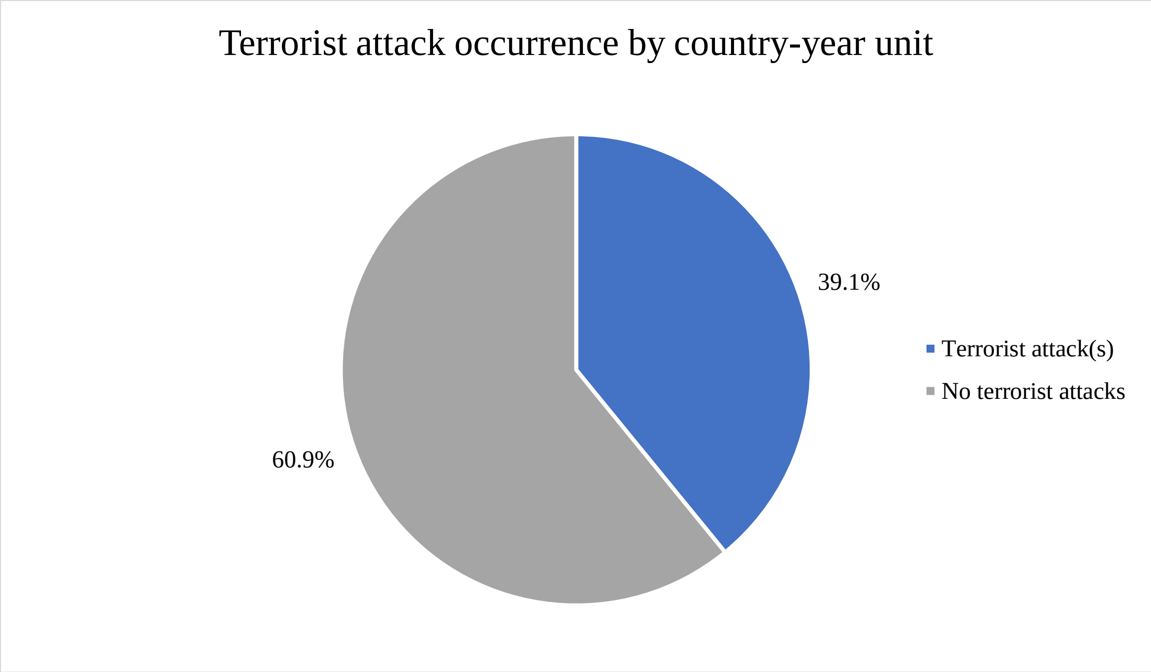

Figure 1: Proportion of country-year units with and without terrorist attacks

The outcome variable is evenly distributed, perhaps surprisingly so: over the entire used dataset, 39% of country-year units experience at least one terrorist attack, as Figure 1 shows. Plotting the yearly distribution of the terrorist attacks (see Figure 2) reveals two important pieces of information. The first is that the idea of long-term trends in terrorism is reflected in this data, with one 'wave' of terrorist attacks beginning in the mid-70s and decreasing again starting in the 90s and a second, much more drastic increase in terrorist attacks beginning shortly before 2010. The second is that in addition to missing data in 1993 (shown by a gap in the graph), there are multiple years with exactly one terrorist attack reported. This is highly unlikely to reflect the real situation in those years and will have to be addressed in building the models, where a set based on the data as is will be compared to an 'improved' version in which those years in which only one attack was registered are replaced by interpolated values.

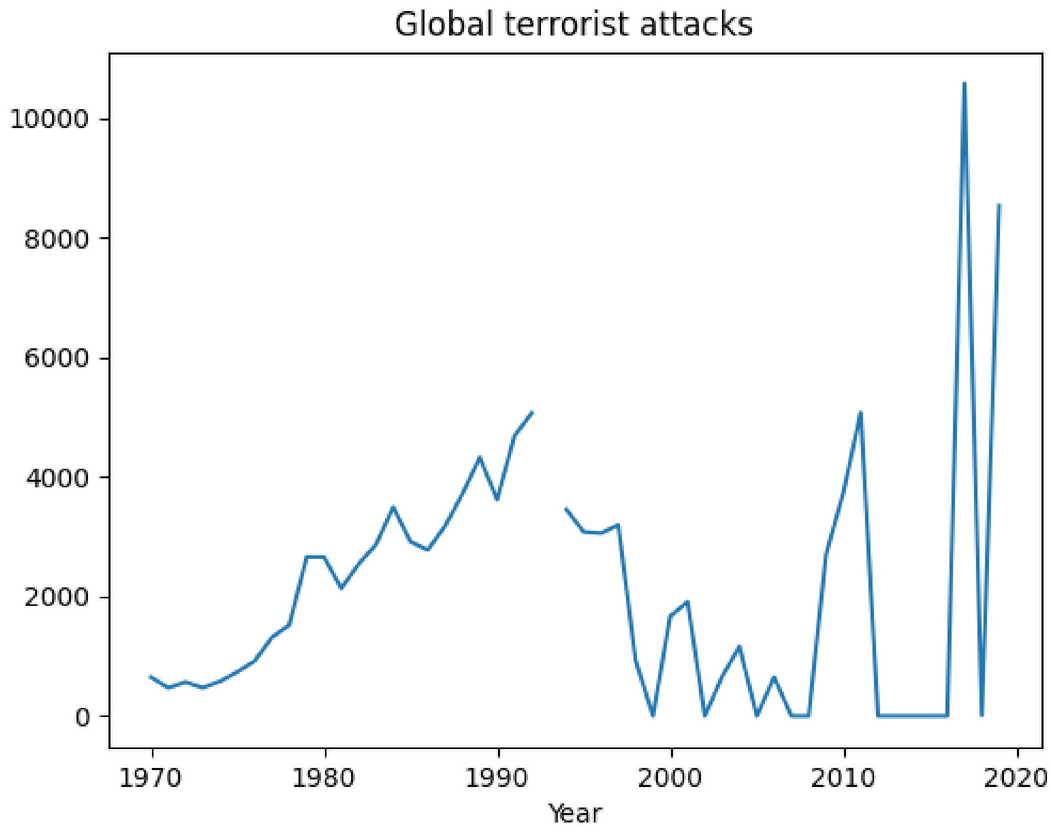

Figure 2: Yearly distribution of global terrorist attack count, as reported in the GTD dataset

The independent variables, or features, for this paper are selected based on both academic and professional publications, as mentioned above. The full list of variables for this study can be found in Table 1 below. As mentioned previously, the features are divided into those derived from academic publications, those derived from professional work, and those found in both bodies of literature.

Table 1: List of variables

Type of variable	Variable	Source
Dependent Variable	Terrorist incident	Global Terrorism Database
Academic Predictor	State fragility	Polity Project
Academic Predictor	Regime durability	Polity Project
Academic Predictor	Electoral system	International IDEA
Academic Predictor	Level of democracy	Polity Project
Academic Predictor	Political rights	Freedom House
Academic Predictor	Civil liberties	Freedom House
Academic Predictor	GDP	Penn World Table
Academic Predictor	Economic inequality	World Bank
Academic Predictor	Poverty rate	World Bank
Academic Predictor	Inflation	World Bank
Academic Predictor	Literacy rate	World Bank
Academic Predictor	Education	Barro-Lee Education Attainment Dataset
Academic Predictor	Population	World Bank
Practitioner Predictor	Terrorist attack (lagged national)	Global Terrorism Database
Practitioner Predictor	Trade volume with the US	Correlates of War Trade Dataset
Practitioner Predictor	Internet usage	World Bank
Practitioner Predictor	Weapon imports	SIPRI Arms Transfers Database
Common Predictor	Terrorist attack count (lagged global)	Global Terrorism Database
Common Predictor	Foreign interventions	PRIF/HSFK Humanitarian Military Interventions Dataset
Common Predictor	Religious fragmentation	Composition of Religious and Ethnic Groups (CREG) Project
Common Predictor	Globalisation	ETH Zürich: KOF Globalization Index

Two features are derived from the same data as the dependent variable, namely the national terrorist attack feature adding information on previous attacks in the country in question and the global lagged terrorist attack variable representing large-scale trends in terrorism identified as a potential predictor in the theory section. The first is simply a lagged version of the dependent variable, meaning it is also coded as a dichotomous country-specific variable, while the second is a count variable representing the number of terrorist attacks registered for a particular calendar year throughout the world. As both variables represent extracting information from previous occurrences of terrorism, neither of them can contribute much to the prediction of the initial onset of this phenomenon. On the other hand, they can hold significant real-world forecasting value (as terrorism does not occur in a vacuum), as evidenced by the use of such data by professionals. To avoid relying too heavily on prior data of what is essentially a variation of the dependent variable, the models for this paper were trained both with and without these two features. A comparison of the two approaches follows in the results section.

Political data is largely taken from datasets offered by the Polity project (Center for Systemic Peace, 2020), which offers extensive data on various political issues with a worldwide country sample and a large time frame. The state fragility index is compiled by the Polity project based on a variety of indicators grouped into effectiveness and legitimacy of the state and consists of a unified scale of 25 points, with high scores associated with high levels of fragility. The durability variable represents the "number of years since the most recent regime change […] or the end of transition period defined by the lack of stable political institutions" (Marshall & Gurr, 2020, p. 17), where a with a regime change is coded zero, and subsequent years without regime change are coded with increments of one. For the level of democracy, the chosen metric is the Polity2 variable from the same dataset, measured as the difference between metrics of democratic and autocratic regime characteristics.

The International Institute for Democracy and Electoral Assistance (IDEA) offers a comprehensive dataset on electoral system design, which is used here to indicate the electoral system family (such as plurality, proportional representation, and mixed systems) of the latest national election in a given country-year unit (International IDEA, 2023). Electoral system families are transformed into dummy variables, with the absence of data (meaning no elections from the beginning of the dataset to a given country-year unit) used as the default case which is later dropped from the dataset.

Additionally, this project makes use of data published by Freedom House (2023), which provide yearly reports on the civil rights and political liberties situation in each country. Both indices are coded on a seven-point scale from 1 to 7, with high numerical values being associated with low levels of rights and liberties and vice-versa.

Data on GDP is taken from the Penn World Table from the University of Groningen (Feenstra et al., 2015). This dataset offers a variety of GDP metrics. This study uses the expenditure-side real GDP metric, which uses units that are constant both across countries and across years and is therefore best-suited for comparisons in both dimensions (Feenstra et al., 2015), which is needed for the models to follow.

The rest of the economic data for this project is taken from databases of the World Bank. Economic inequality (World Bank, 2023d) is measured as the Gini index per country per year, which is probably the most well-known and established metric for inequality. This coefficient measures the extent to which the distribution of income deviates from perfect equality and ranges from 0 (perfect equality) up to 1 (perfect inequality). Poverty data is again taken from the same World Bank dataset, with the chosen metric being the "Poverty gap at \$6.85 a day (2017 PPP) (%)" variable, which is measured as the mean difference from the poverty line of \$6.85 per day, with individuals with income superior to that poverty line being counted as zero difference. This metric has the benefit of combining the extent (number of people) and severity of poverty into a single variable. Inflation data is again taken from the World Bank (2023b), who offer various measures of inflation, among which the Headline Consumer Price Inflation metric combines inflation indicators such as food and energy prices into a unified country-year dataset. It is expressed as a percentage representing yearly inflation.

For education, this study uses data on literacy again taken from the World Bank (2022), who offer yearly country-wise data on the percentage of literate individuals aged 15 or above per country. An alternative measure for education is taken from the Barro-Lee Educational Attainment Dataset (Barro & Lee, 2021), which offers data on the proportion of people in a country with specific levels of education. For this project, the variable used is the share of the entire population having completed primary schooling, expressed as a percentage.

The World Bank (2023c) is also the source for data on country population used in this project, which is simply the population in millions per country-year units. The use of this variable follows two justifications. Firstly, it is used as an independent feature following its theoretical backing described above. Secondly, it is used since features such as the economic data (GDP per country) are not scaled and allows for little international comparisons with no information on the population of a country. It is only with the incorporation of information about the size of a country (in this case, its population) that absolute values of GDP or weapon imports can be contextualised to compare countries with vastly different sizes, for which the same values in these absolute metrics would represent very different situations in terms of GDP per capita or weapon density.

The trade volume of a country with the US is used as a proxy for ties with and possible dependency on the United States. The data hereon is taken from the Correlates of War project's trade dataset (Barbieri & Keshk, 2016), which offers yearly dyadic data on trade flows in constant units (current US dollars at the time of publication of the dataset). Each country's own trade volume with the United States (specifically, the sum of both imports from and exports to the US) is divided by its GDP in the same year to obtain a simple metric that is comparable across both time and countries, as both the GDP and trade data use constant units. Note that as this indicator does not exist for the United States themselves, an alternative set of models excluding this variable is computed that can include the US in the sample.

Internet usage data is another variable on which data is taken from the extensive World Bank (2023a) datasets and measures the share of the population (again expressed as a percentage) using the internet. Specifically, the metric used in the World Bank dataset is individuals who have accessed the internet at least once in the previous three months.

The Stockholm International Peace Research Institute (2023) offers an extensive dataset on weapons imports and exports with the SIPRI Arms Transfer Database. Of particular interest is their measure of trend-indicator values (TIV), which are “intended to represent the transfer of military resources rather than the financial value of the transfer” (*Sources and Methods | SIPRI*, n.d., sec. 2). This provides several key advantages. Firstly, it allows for a single metric when comparing vastly different types of weapon transfers, so that the import of small arms can reasonably be compared to that of heavier equipment, which would not be possible when looking exclusively at count data on pieces of equipment transferred. Second, TIV scores include the depreciation in value of equipment over time, so that used equipment is assigned less value than new items of the same type, though refurbishments are also considered. This adds informative value that would not be present if looking exclusively at the production cost of armament imports. Lastly, the TIV system is designed in a way to remain consistent over time, allowing for cross-temporal analyses such as the one carried out in this project.

Data on military interventions is taken from the PRIF/HSFK Humanitarian Military Interventions Dataset (Peace Research Institute Frankfurt, 2019). This dataset charts all proclaimed humanitarian interventions since the end of the Second World War, be it by individual nation states or by international organisations such as the UN. Due to possible differences in legitimacy between individual interventions and those carried out with the backing of an international organisation, two different features are created from the PRIF data: one for interventions where the primary intervening party is a nation state, and one for interventions for which the primary party is an International Organisation (IO). This data is turned into a dichotomous variable per country per year, where 1 represents an ongoing intervention, either of the country on its own or of an IO that the country is a member of, and 0 represents no such intervention.

Religious fragmentation is calculated based on data stemming from the University of Illinois' (n.d.) Composition of Religious and Ethnic Groups (CREG) Project. The metric is inspired by Pinckney and RezaeeDaryakenari (2022), who calculate their "Ethnic and Religious ethnolinguistic fractionalization" (p. 1009) variable[1] in such a way that it ranges from 0 to 1, with 0 representing a perfectly homogenous population consisting only of a single (in this case, religious) group, and 1 an unreachable infinitely fragmented population.

Finally, the metric used for globalisation is the country-specific KOF Globalisation Index published by the ETH Zürich (Gygli et al., 2019). This data combines indicators on economic (such as trade and cross-border investment), social (such as tourism, migration, internationally recognisable brands, and access to international information sources), and political (such as diplomatic relations, membership in international organisations, and the presence of international non-governmental organisations) aspects of globalisation. As such, it offers itself well to representing the wide-ranging, but ill-defined concept of globalisation.

Methods

The variables listed above are combined into a unified dataset with a country-year unit of analysis. Missing values are, wherever possible, imputed through linear interpolation. Missing values outside of existing data (before the earliest or after the latest existing datapoint)

[1] "$ELF = 1 - \pi_i^2$, where pi group i's population share." (Pinckney & RezaeeDaryakenari, 2022, p. 1009). In this case, groups are determined based on religious affiliation.

are left out[2], and only a maximum of ten consecutive years for any country may be interpolated to avoid erring too much on the side of guessing.

Due to the nature of the model to be developed, this research project takes large-n[3] approach on a global scale, with the country-year as the unit of analysis. Country-years are excluded based only on data availability. The selected dataset for the dependent variable, terrorist attacks, covers the years 1970 to 2020, but the final sample is limited to the years 1995-2013 due to data availability for the other variables.

Several types of models are used and compared to each other in this project. While a logistic regression model is included, this is mostly to create a baseline and to facilitate comparisons with the wider terrorism literature which tends to rely on such parametric models for its large-N studies. This project focuses on other models using Random Forest (RF) and Gradient-Boosting Machine (GBM) algorithms, which are more suitable for modelling and forecasting complex social phenomena such as terrorism thanks to their implicit ability to work with non-linear relationships and interactions between features without having to implement these manually. Furthermore, these algorithms are known for their relative robustness to outliers.

The data is split into distinct training and test sets in such a way that roughly 75%[4] of the data used in the main model lies in the training set. All data up to and including the year 2009 is used in the training set, and data from 2010 onwards is reserved for the final testing. The training data is then used for a five-fold cross-validation process in which iterations of each of the three algorithms are trained on different subsets of the training set and then

[2] Note that this applies only to the interpolation carried out in this project- the KOF Globalisation Index as published on the ETH Zürich website, for example, already contains some interpolated and extrapolated data, including values that were back-carried to before the earliest data point or forward-carried to after the latest available data.
[3] N = 1,271 for the main model
[4] Specifically, 74.7%

evaluated on their out-of-sample performance when predicting values present in a different subset of the training set. For each algorithm, the best-performing iteration in terms of area under the precision-recall curve (AUC-PR, a metric explained below) is kept and used in the final evaluation and comparison of the algorithms by measuring its performance on the dedicated test set mentioned above. Rather than inherent and deterministic metrics such as R-squared values for parametric models, the forecasting models are evaluated by the quality of their predictions for the test set.

Such predictions can be assessed in terms of precision and recall at a specific classification threshold - that is, the minimum predicted probability of an outcome that should be met for the classifier to predict the outcome itself. A more interesting metric to assess forecasts coming from classifiers, however, is the area under the Receiver Operating Characteristic curve (AUC-ROC), which plots the true and false positive rates against each other for a variable prediction threshold and measures the integral of the curve in question. While this classification threshold is generally set at 0.5 by default for most classifiers, there is no particular reason why this should be the value of choice, and the AUC-ROC metric has the benefit of providing a score of the model that summarises the performance at various thresholds in addition to combining information on true and false positives into a single metric. AUC-ROC scores range from 0 to 1, with 0.5 representing no ability to distinguish between outcome classes and higher values representing better performance. Values close to 0 represent an inversion of the classes, meaning the model guesses the wrong classes, but does so in a consistent manner. This is not generally encountered.

Another metric relevant to this paper is the area under the precision-recall curve (AUC-PR), which plots the precision and recall values against each other for a variable classification threshold. While this metric lacks a precise baseline like the value of 0.5 for the AUC-ROC metric, the occurrence rate of the event to be predicted can be used as a general benchmark in the way that AUC-PR scores above this number represent an improvement over chance (Morgan et al., 2019). A key advantage of this metric over AUC-ROC scores is that it is more robust to rare events forecasting. The AUC-ROC score only provides an assessment of the quality of positive predictions, which can inflate it artificially in cases where the event to be predicted is so rare that the models almost never output a positive forecast (Pinckney & RezaeeDaryakenari, 2022). While the classes in this study are not highly skewed, as shown in Figure 1, they are still imbalanced, justifying the inclusion of and focus on this latter metric.

Variable importance is measured as the mean decrease in impurity (MDI), the most common metric for feature importance in decision tree-derived models such as random forests and gradient boosting machines. It measures the extent to which decision nodes based on a given variable can split mixed data into pure outcomes, or in other words, discriminate between the classes to be predicted. This is where the main analytical difference between MDI scores in these models and more commonly seen regression coefficients and p-values in regression models lies. While the latter concern a global relationship between a variable and the outcome and measure the direction and strength of that effect, the former is related purely to the degree to which a variable permits the model to differentiate between outcomes. Expressing the relationships modelled in non-parametric models as a single coefficient is usually unfeasible, as many nodes in a single random forest model can be based on the same variable, resulting in complex relationships which are often conditional on outside factors such as the values of other variables. The MDI measures are scaled so that the most important feature in any model obtains a score of 1, facilitating comparisons between model sets.

Results and Discussion

A simple logistic regression model built only from predictors based on academic literature (i.e., both ‘academic’ and ‘common’ predictors from Table 1) is constructed to provide a baseline, since such regression algorithms form the majority of quantitative analyses on terrorism. Like the more complex, non-parametric models to follow, however, the performance of this first regression is measured by its forecasting ability instead of the more traditional evaluation of p-values and R-squared scores. The AUC-ROC score, described in the methods section, is 0.83 for this model, and the AUC-PR score is 0.88, indicating that the model has significant value in forecasting terrorist attacks, though this is not the case at the default classification threshold, at which the model consistently predicts the presence of a terrorist attack in every country-year unit. This illustrates the necessity for threshold-independent metrics such as the areas under the ROC and PR curves. Unsurprisingly, the Random Forest and Gradient Boosting Machine algorithms both score higher on both metrics, though this is not the case for all variable sets, as will be shown shortly.

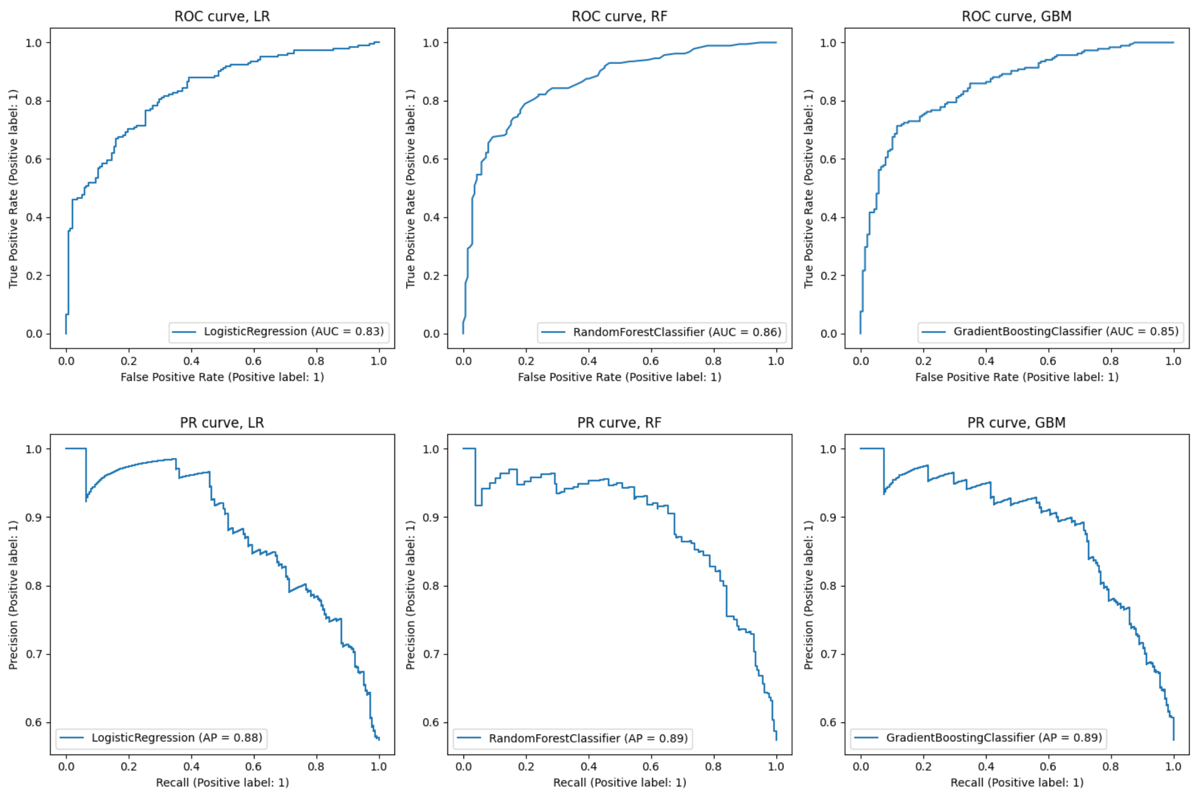

Figure 3: ROC and PR curves per algorithm for the main model (using all variables)

The AUC-ROC scores for the various algorithms and variable combinations can be seen in Table 2, which reports this metric for each type of model (in the columns) and variable set (in the rows). Academic and professional refer to the variable sets using only features identified from either body of literature, as described above. *Combined* refers to the two sets of features together, while *All* is the same set with the two GTD-derived features added as well (which are not included in any of the first three sets).

Table 2: Model performance (AUC-ROC) by variable set and model type, rounded

	LR	RF	GBM
Academic	0.83	0.86	0.85
Professional	0.83	0.84	0.80
Combined	0.83	0.86	0.82
All	0.83	0.87	0.85

While the distribution of classes is not highly imbalanced for these models, as seen in Figure 1, an inclusion of the AUC-PR scores is still relevant. As can be seen in Table 3, the general patterns identified in a comparison of the AUC-ROC scores are also reflected in the results relating to this metric.

Table 3: Model performance (AUC-PR) by variable set and model type, rounded

	LR	RF	GBM
Academic	0.88	0.89	0.89
Professional	0.88	0.87	0.85
Combined	0.88	0.89	0.87
All	0.88	0.89	0.89

An interesting observation is that while the academic-only set allows for better scores than the professional-only set for the non-parametric models, the difference is insignificant for the logistic regression, despite this being a method much more commonly used by political scientists. Additionally, the two features derived from the Global Terrorism Database (the country-specific lagged terrorist attack variable and the lagged global terrorist attack count) indeed add predictive power to the non-parametric models, though they fail to do so in a significant way for the logistic regression.

Furthermore, while the results show the added value in combining variables from both academic and professional work, it is interesting to note that the differences in performance between the various groups are limited. In fact, all four sets of features result in models that provide significant forecasting abilities, with differences being approximately as strong between algorithms as between variable sets.

Another representation of these scores can be found in Figure 4, which plots the former score over the latter for the three algorithms per variable set. As both axes have their origin in the lower left, the best-performing models are those that are closest to the top right corner. The random forest algorithm generally scores best, with the exception for the *professional* variable set, in which the logistic regression out-performs it in terms of AUC-PR.

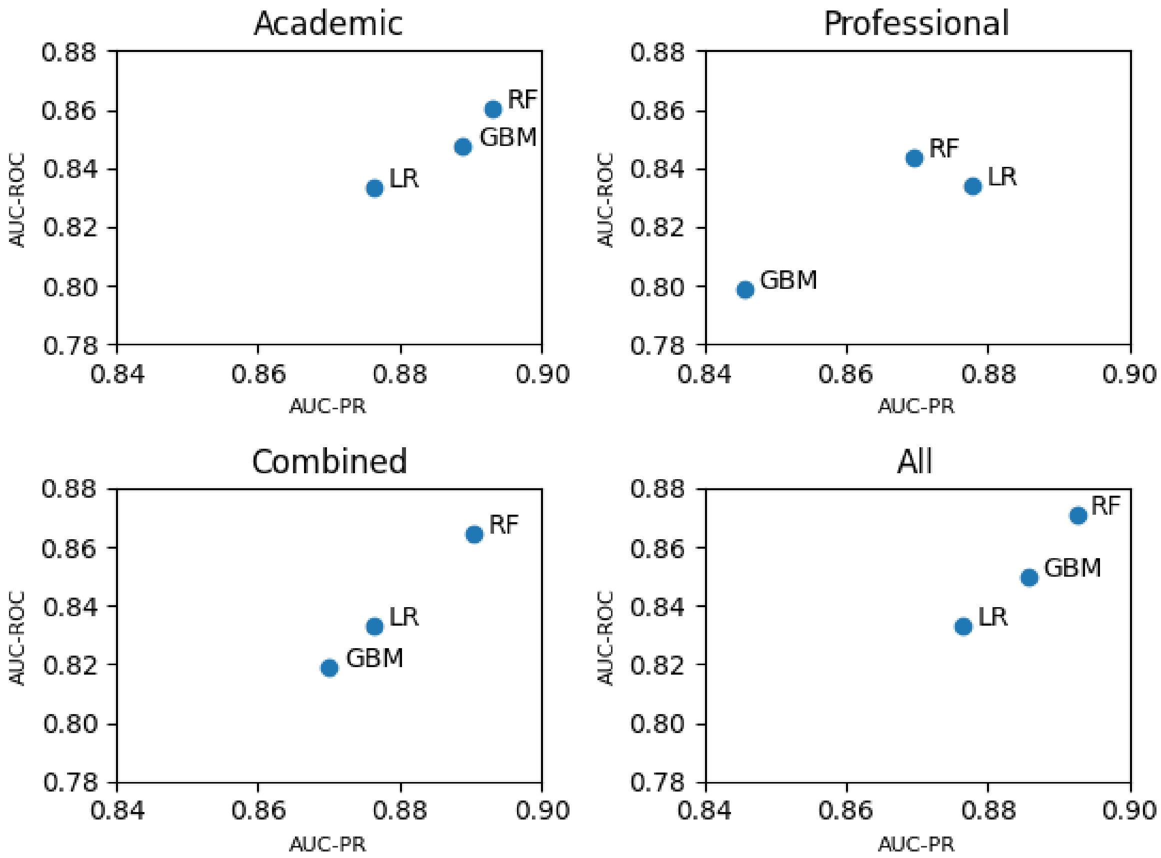

Figure 4: AUC-ROC score over AUC-PR score per algorithm for each of the main variable sets

Let us now turn to the relative importance of the features. As the random forest models tend to perform the best, it is this algorithm with the *all* feature set that will be considered the main model for the analysis to follow. The result of the scaled MDI computation, described in the methods section, can be found in Figure 5 below.

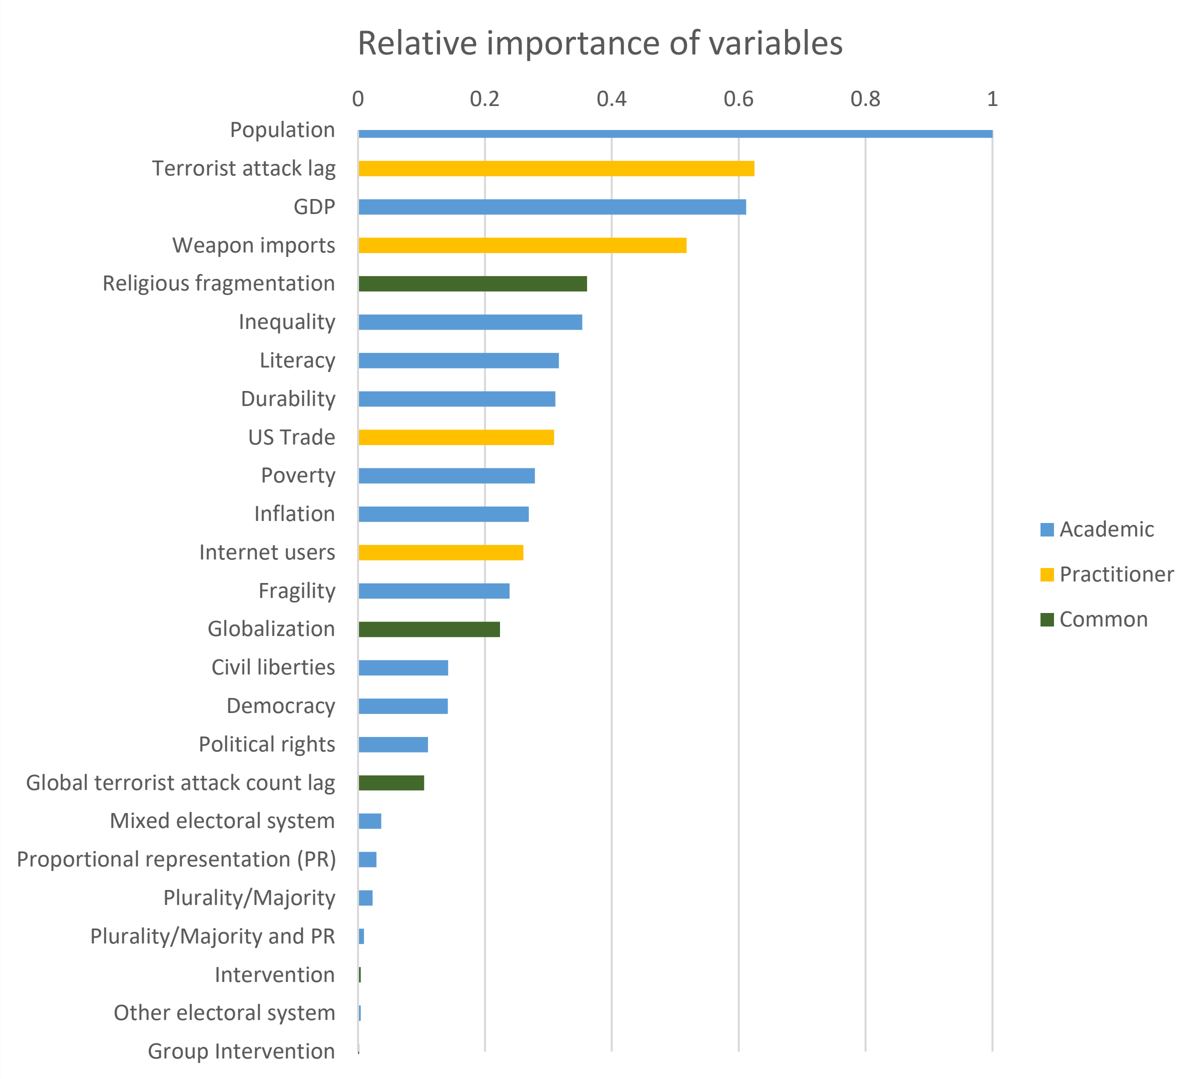

Figure 5: Scaled relative feature importance for the all-encompassing RF model

This figure shows that the most important predictor for terrorist attack occurrence (with a considerable gap) for this model is the population of a country. To some extent, this result is unsurprising, as the outcome variable for this model is dichotomous, and the likelihood of zero terrorist attacks happening in a country for a full year decreases as the population increases. This assumption is reflected in the data, as can be seen in Figure 6.

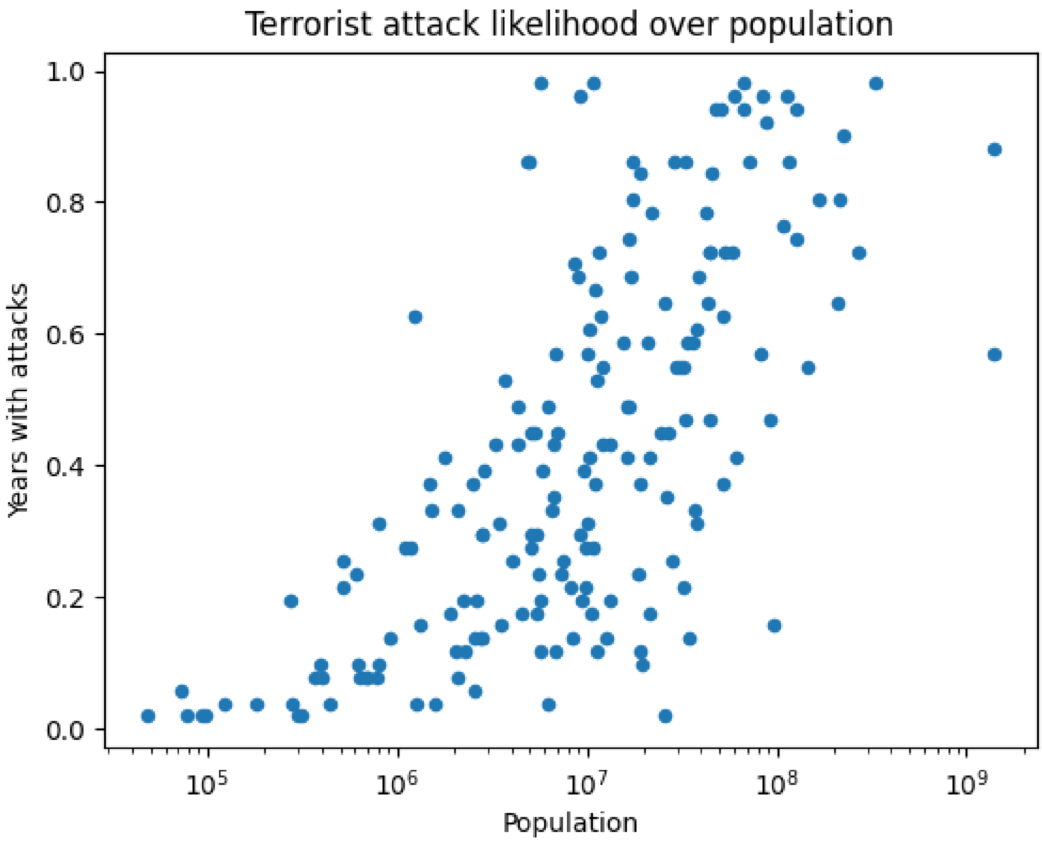

Figure 6: Proportion of years with at least one terrorist attack plotted per country over its 2020 population

Nevertheless, there may be more to this result. While the correlation between population and the binary outcome variable helps explain it, all other features, too, are included due to expectations of significant relationships based on previous work, be it professional or academic. Furthermore, previous studies, too, have found population (Piazza, 2006) or population growth (Coccia, 2018) to be a significant predictor of terrorism and, in a related field of study, civil war (Fearon & Laitin, 2003; Hegre & Sambanis, 2006). Despite these repeated and robust findings, the exact mechanisms at play behind these relationships have seen relatively restricted academic interest as of now, and this body of work is nowhere near as extensive as that on the effect of various political metrics on terrorism. Still, variables such as those reflecting the election system and those taken from Freedom House appear to add little to no value to the model, despite being much more frequently identified in academic research as topics of interest.

Unsurprisingly, the added predictive power of models containing the lagged derivates of the outcome variable is also reflected in Figure 5. Interestingly, it appears it is almost exclusively the country-specific terrorist attack variable that adds to the predictive power, as the global terrorist attack count scores low in relation to the other features.

The expectation that common variables would have the highest predictive value cannot be confirmed. The difference in interests, objectives, and methods between terrorism researchers and counter-terrorism practitioners means that the overlap in information used may reflect data that is useful for a particular variety of purposes, rather than data that is particularly useful for one specific task.

Four variants of the models are created for the sake of comparison and robustness tests, which will be addressed below. The detailed results in terms of performance and feature importance of the main models and of these variants can be found in the appendix. First, the literacy data from the World Bank is replaced by the education attainment data from the Barro-Lee dataset. This results in a moderate decrease in overall performance, but the variable has a somewhat higher relative MDI-score than the literacy feature in the previous configuration. More interestingly, this is the only configuration in which the country-specific terrorist attack lag overtakes population in relative importance.

Second, a new set of models is proposed where the data on global terrorist attack counts (used as an independent variable in the *All* configuration) has been replaced by interpolated values for those years in which the GTD reported exactly one terrorist attack worldwide, as described above. This has no significant impact on performance for any of the three algorithms, though the relative importance of the feature increases slightly.

Third, the US trade variable is dropped from the dataset to create a model that does not exclude the United States from its sample. This again has no significant impact on forecasting ability for any of the three algorithms, indicating that dropping the United States from the main sample to add the feature on ties to the US was likely unproblematic.

A last model is created dismissing the GDP and weapon imports variables from the feature list to examine the effect on the relative importance of the population variable. This causes a drop in out-of-sample performance[5], except for the GBM model, which performs slightly better in terms of the area under the ROC curve. More importantly, the population feature retains its position as having the highest MDI score, and the gap to the second-ranked feature grows from the latter having a scaled relative importance score of 0.62 to one of 0.59. This confirms that the forecasting value of the population feature is not restricted to providing context for other variables, but that it independently provides relevant information. As discussed above, this is partly due to the construction of the dependent variable, but also confirms the indications by previous research that there are as-of-yet under-researched causal mechanisms at play.

Conclusion

This paper has aimed to provide a systematic assessment of the value of various factors identified from academic and practitioner publications when it comes to predicting the threat of terrorist attacks. To do so, three machine learning algorithms were trained on data on these factors, with random forests consistently outperforming the more conventional logistic regression that is commonly encountered in quantitative political science research. An investigation into the relative importance of the variables given to the models, measured as the mean decrease in impurity (MDI), reveals multiple points of interest. First, that the population

[5] For the random forest model, from an AUC-PR of 0.89 to 0.87

of a country appears to bring the highest predictive value across almost all variations of the model, irrespective of whether it is needed to provide context for economic data. Second, that the lagged terrorist attack variable, as expected, has a high predictive power, though the lagged global terrorist attack count variable does not to the same extent. Still, adding these two variables to the model provides for a boost in forecasting ability. Lastly, the expectation that data used by both academics and practitioners would have the highest MDI scores did not hold true.

This paper, to my knowledge, provides the first comprehensive overview of the value of variables inspired by various theoretical models and practitioner publications in forecasting terrorism. This is of interest to academics, as it provides a new way of evaluating the usefulness of existing theoretical frameworks and potential inspiration for the development of new ones.

Furthermore, the methodological approach used here can be of particular use to practitioners. In the methodology for intelligence analysis described by Khalsa (2004), one of the principal objectives and merits of the new approach to analysis lies in rendering it more systematic. Among the yearly tasks in the methodology, the two tasks "Identify/validate [warning] indicators" and even more so "Determine/validate priorities of [warning] indicators" (Khalsa, 2004, p. 9) resemble what has been done in this paper and consequently could benefit from evaluating these "priorities" by means of an empirical investigation into the relevance of indicators to forecasting terrorist threats, as has been done here. Such an analysis would of course be carried out with the indicators that are already being used, the very nature of some of which is left out of Khalsa's (2004) book for security reasons, but this would not be a problem for the intelligence organisation itself, which would have access to this data. Using such an empirical tool would potentially provide great help to the "leading counterterrorism experts" (Khalsa, 2004, p. 10), who meet annually to carry out the two tasks mentioned above in this methodology.

A potential avenue for further research would be to develop similar models on a subnational level, potentially also on a more disaggregated time scale. This would get closer to the work of counterterrorism practitioners, though the necessary data may be harder to obtain. A comparable project is the Violence Early-Warning System (ViEWS) of Uppsala University (n.d.) and the Peace Research Institute Oslo, which provides monthly forecasts for state-based conflict “for each country and 55x55 km location in Africa and the Middle East” (para. 1). Examples of data such models may include are short-term socio-political events, on which data is available and has been used (though in a yearly aggregated form) to forecast the onset of political dissent (Pinckney & RezaeeDaryakenari, 2022), as well as near real-time political sentiment data extracted from Twitter or social media through some form of natural language processing. The vastly superior quantities of data, as well as the increased difficulty in ensuring cross-country comparability, put such projects far beyond the scope of this study, but this does not take away from the potential value in such research in so far as that this piece has determined the relative feature importance only when looking at long-term country-level developments in terrorist threats, which may be of limited use to practitioners.

Another pathway to building onto this study would be to add new types of algorithms to the models. In addition to other machine learning algorithms suitable for single-year analyses – Pinckney and RezaeeDaryakenari (2022) use a battery of 13 algorithms for such a study – this may also allow for expanding the scope of research. For example, the models in this study use only the feature values of year t to predict the occurrence or lack of occurrence of terrorist attacks in year $t + 1$. Whether the use of accumulative data on past values over the years, perhaps even by means of a weighted average where recent values are given more importance or by rates of change, would have resulted in more predictive power and a higher relative importance remains an open question. Indeed, inflation rates (Ajide & Alimi, 2023) and population (Coccia, 2018), to name just two, have both been identified as variables in which the change over time can affect the onset of terrorism. The use of time-accumulated data would transform this study into a multivariate time series analysis, for which a long short-term memory (LSTM) neural network, for example, may be a good fit.

Turning to building onto the results of this study, there is a clear necessity for a detailed investigation into the role of a country's population as a predictor for terrorism, as this study joins an existing set of studies that found indicators for the existence of such a relationship. Given the repeated and robust findings in this direction, it is surprising to see the lack of investigation into the mechanisms at play when compared to economic or political factors, on which there is an extensive existing body of high-quality research.

Bibliography

Abbasi, I., & Khatwani, M. (2018). A review of psycho-social theories of terrorism. *Grassroots*, *51*(2), 319–333.

Ajide, K. B. (2021). Democracy, Regime Durability and Terrorism in Africa. *Defence and Peace Economics*, *32*(5), 550–571. https://doi.org/10.1080/10242694.2019.1671090

Ajide, K. B., & Alimi, O. Y. (2023). Inflation, inflation volatility and terrorism in Africa. *International Journal of Finance & Economics*, *28*(1), 493–509. https://doi.org/10.1002/ijfe.2432

Augsburger Allgemeine. (2016, July 24). Amok und Terrorismus: Was ist der Unterschied? *Augsburger Allgemeine*. https://www.augsburger-allgemeine.de/panorama/Amoklauf-in-Muenchen-Amok-und-Terrorismus-Was-ist-der-Unterschied-id38588662.html

Bakker, E. (2012). Forecasting Terrorism: The Need for a More Systematic Approach. *Journal of Strategic Security*, *5*(4), 69–84.

Bakker, R., Hill, D. W., & Moore, W. H. (2014). Modeling Terror Attacks: A Cross-National, Out-of-Sample Study. In R. Caruso & A. Locatelli (Eds.), *Understanding Terrorism: A Socio-Economic Perspective* (Vol. 22, pp. 51–68). Emerald Group Publishing. https://doi.org/10.1108/S1572-8323(2014)0000022008

Balboni, J. (2016, January 15). Attentats de Paris: Chakib Akrouh, un inconnu devenu terroriste. *DHnet*. https://www.dhnet.be/actu/monde/2016/01/15/attentats-de-paris-chakib-akrouh-un-inconnu-devenu-terroriste-T6OBAF5DXBDCRLSCUXAOW2EPQM/

Barbieri, K., & Keshk, O. M. G. (2016). *Correlates of War Project Trade Data Set* (4.0) [Data set]. https://correlatesofwar.org/data-sets/bilateral-trade/

Barro, R. J., & Lee, J.-W. (2021). *Barro-Lee Educational Attainment Dataset* [Data set]. http://www.barrolee.com/

Bergesen, A. J., & Lizardo, O. (2004). International Terrorism and the World-System. *Sociological Theory*, *22*(1), 38–52. https://doi.org/10.1111/j.1467-9558.2004.00203.x

Brockhoff, S., Krieger, T., & Meierrieks, D. (2012). Great expectations and hard times: The (nontrivial) impact of education on domestic terrorism. *CESifo Working Paper*, *3817*.

Center for Systemic Peace. (2020). *Polity5 Dataset* [Data set]. https://www.systemicpeace.org/inscrdata.html

Center for Systemic Peace. (2022). *INSCR Data Page*. https://www.systemicpeace.org/inscrdata.html

Chenoweth, E. (2013). Terrorism and Democracy. *Annual Review of Political Science*, *16*(1), 355–378. https://doi.org/10.1146/annurev-polisci-032211-221825

Coccia, M. (2018). The relation between terrorism and high population growth. *Journal of Economics and Political Economy*, *5*(1), 84–104.

Desmarais, B. A., & Cranmer, S. J. (2013). Forecasting the locational dynamics of transnational terrorism: A network analytic approach. *Security Informatics*, *2*(1), 8. https://doi.org/10.1186/2190-8532-2-8

Douglas, K. M., Uscinski, J. E., Sutton, R. M., Cichocka, A., Nefes, T., Ang, C. S., & Deravi, F. (2019). Understanding Conspiracy Theories. *Political Psychology*, *40*, 3–35. https://doi.org/10.1111/pops.12568

Elbakidze, L., & Jin, Y. H. (2015). Are Economic Development and Education Improvement Associated with Participation in Transnational Terrorism? *Risk Analysis*, *35*(8), 1520–1535. https://doi.org/10.1111/risa.12378

Enders, W., Hoover, G. A., & Sandler, T. (2016). The Changing Nonlinear Relationship between Income and Terrorism. *Journal of Conflict Resolution*, *60*(2), 195–225. https://doi.org/10.1177/0022002714535252

Erickson, A., & Stanley-Becker, I. (2021, December 1). How ramming cars into crowds became a major terror tactic. *Washington Post*. https://www.washingtonpost.com/news/worldviews/wp/2017/03/22/as-the-london-attacks-show-the-sedate-sedan-can-be-turned-into-a-deadly-unpredictable-weapon/

Fearon, J. D., & Laitin, D. D. (2003). Ethnicity, Insurgency, and Civil War. *American Political Science Review*, *97*(1), 75–90. https://doi.org/10.1017/S0003055403000534

Feenstra, R. C., Inklaar, R., & Timmer, M. P. (2015). The Next Generation of the Penn World Table. *American Economic Review*, *105*(10), 3150–3182.

Février, R. (2015, March 7). Attaque d’un restaurant à Bamako: ‘Un attentat parfaitement imprévisible’. *L’OBS*. https://www.nouvelobs.com/monde/20150307.OBS4099/attaque-d-un-restaurant-a-bamako-un-attentat-parfaitement-imprevisible.html

Fleming, C. M., Manning, M., Pham, H.-T., & Vorsina, M. (2022). Ethnic Economic Inequality and Fatalities From Terrorism. *Journal of Interpersonal Violence*, *37*(11–12), NP9089–NP9114. https://doi.org/10.1177/0886260520976226

Freedom House. (2023). *Country and Territory Ratings and Statuses, 1973-2023* [Data set]. https://freedomhouse.org/reports/publication-archives

George, J. (2018). State Failure and Transnational Terrorism: An Empirical Analysis. *Journal of Conflict Resolution*, *62*(3), 471–495. https://doi.org/10.1177/0022002716660587

Gygli, S., Haelg, F., Potrafke, N., & Sturm, J.-E. (2019). The KOF Globalisation Index – Revisited. *Review of International Organizations*, *14*(3), 543-574. https://doi.org/10.1007/s11558-019-09344-2

Hegre, H., Metternich, N. W., Nygård, H. M., & Wucherpfennig, J. (2017). Introduction: Forecasting in peace research. *Journal of Peace Research*, *54*(2), 113–124. https://doi.org/10.1177/0022343317691330

Hegre, H., & Sambanis, N. (2006). Sensitivity Analysis of Empirical Results on Civil War Onset. *Journal of Conflict Resolution*, *50*(4), 508–535. https://doi.org/10.1177/0022002706289303

Homeland Security Advisory Council. (2007). *Report of the Future of Terrorism Task Force*. U.S. Department of Homeland Security. https://www.dhs.gov/xlibrary/assets/hsac-future-terrorism-010107.pdf

Hugues, P. (2015, November 22). Nach dem Terror von Paris: Wenn das Alltägliche plötzlich entstellt ist. *Der Tagesspiegel Online*. https://www.tagesspiegel.de/politik/wenn-das-alltagliche-plotzlich-entstellt-ist-5651845.html

International IDEA. (2023). *Electoral system design database* [Data set]. https://www.idea.int/data-tools/data/electoral-system-design

Ismail, A., & Amjad, S. (2014). Determinants of terrorism in Pakistan: An empirical investigation. *Economic Modelling*, *37*, 320–331. https://doi.org/10.1016/j.econmod.2013.11.012

Khalsa, S. K. (2004). *Forecasting Terrorism: Indicators and Proven Analytic Techniques*. Scarecrow Press. https://rowman.com/ISBN/9780810850170/Forecasting-Terrorism-Indicators-and-Proven-Analytic-Techniques

Kluch, S. P., & Vaux, A. (2017). Culture and Terrorism: The Role of Cultural Factors in Worldwide Terrorism (1970–2013). *Terrorism and Political Violence*, *29*(2), 323–341. https://doi.org/10.1080/09546553.2015.1038523

Korotayev, A., Vaskin, I., & Tsirel, S. (2021). Economic Growth, Education, and Terrorism: A Re-Analysis. *Terrorism and Political Violence, 33*(3), 572–595. https://doi.org/10.1080/09546553.2018.1559835

Krueger, A. B., & Malečková, J. (2003a). Education, Poverty and Terrorism: Is There a Causal Connection? *Journal of Economic Perspectives*, *17*(4), 119–144. https://doi.org/10.1257/089533003772034925

Krueger, A. B., & Malečková, J. (2003b). Seeking the Roots of Terrorism. *The Chronicle*, *49*(39), 6.

Li, Q. (2005). Does Democracy Promote or Reduce Transnational Terrorist Incidents? *Journal of Conflict Resolution*, *49*(2), 278–297. https://doi.org/10.1177/0022002704272830

Marshall, M. G., & Gurr, T. R. (2020). Political Regime Characteristics and Transitions, 1800-2018: Dataset Users' Manual. *Center for Systemic Peace*. https://www.systemicpeace.org/inscr/p5manualv2018.pdf

McCauley, C., & Moskalenko, S. (2008). Mechanisms of Political Radicalization: Pathways Toward Terrorism. *Terrorism and Political Violence*, *20*(3), 415–433. https://doi.org/10.1080/09546550802073367

Morgan, R., Beger, A., & Glynn, A. (2019). Varieties of Forecasts: Predicting Adverse Regime Transitions. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.3389194

Nabin, M. H., Chowdhury, M. T. H., Bhattacharya, S., & Sgro, P. M. (2022). Terrorism and uneven economic development. *Economic Modelling*, *113*, 105870. https://doi.org/10.1016/j.econmod.2022.105870

NCTV. (2022). *Dreigingsbeeld NCTV* [Nieuwsbericht]. Ministry of Justice and Security. https://www.nctv.nl/actueel/nieuws/2022/11/07/dreigingsbeeld-nctv-dreiging-in-en-tegen-nederland-veelzijdiger-en-meer-diffuus

Paz, R. (2003). Islamists and Anti-Americanism. *Middle East Review of International Affairs*, *7*(4), 53–61.

Peace Research Institute Frankfurt. (2019). *PRIF Dataset on Humanitarian Military Interventions* (1.14) [Data set]. https://www.humanitarian-military-interventions.com/

Piazza, J. A. (2006). Rooted in Poverty?: Terrorism, Poor Economic Development, and Social Cleavages. *Terrorism and Political Violence*, *18*(1), 159–177. https://doi.org/10.1080/095465590944578

Piazza, J. A. (2007). Draining the Swamp: Democracy Promotion, State Failure, and Terrorism in 19 Middle Eastern Countries. *Studies in Conflict & Terrorism*, *30*(6), 521–539. https://doi.org/10.1080/10576100701329576

Piazza, J. A. (2012). Types of Minority Discrimination and Terrorism. *Conflict Management and Peace Science*, *29*(5), 521–546. https://doi.org/10.1177/0738894212456940

Pinckney, J., & RezaeeDaryakenari, B. (2022). When the levee breaks: A forecasting model of violent and nonviolent dissent. *International Interactions*, *48*(5), 997–1026. https://doi.org/10.1080/03050629.2022.2090933

RAND National Security Research Division. (n.d.). *About RDWTI*. Retrieved 15 March 2023, from https://www.rand.org/nsrd/projects/terrorism-incidents/about.html

Schmid, A. (2004). Terrorism—The Definitional Problem. *Case W. Res. J. Int'l L*, *36*, 375.

Shahbaz, M. (2013). Linkages between inflation, economic growth and terrorism in Pakistan. *Economic Modelling*, *32*, 496–506. https://doi.org/10.1016/j.econmod.2013.02.014

Sinai, J. (2002). *How to Forecast the Next Waves of Catastrophic Terrorism | Office of Justice Programs*. https://www.ojp.gov/ncjrs/virtual-library/abstracts/how-forecast-next-waves-catastrophic-terrorism

Sources and methods | SIPRI. (n.d.). Stockholm International Peace Research Institute. Retrieved 25 May 2023, from https://www.sipri.org/databases/armstransfers/sources-and-methods

START. (2021a). *Global Terrorism Database codebook: Methodology, inclusion criteria, and variables*. University of Maryland. https://www.start.umd.edu/gtd/downloads/Codebook.pdf

START. (2021b). *Global Terrorism Database (GTD)*. [Data set]. https://www.start.umd.edu/gtd

Stern, J. (2009). *Terror in the Name of God: Why Religious Militants Kill*. Harper Collins.

Stockholm International Peace Research Institute. (2023). *SIPRI Arms Transfers Database* [Data set]. https://www.sipri.org/databases/armstransfers

Tikuisis, P. (2009). On the relationship between weak states and terrorism. *Behavioral Sciences of Terrorism and Political Aggression*, *1*(1), 66–79. https://doi.org/10.1080/19434470802482175

University of Illinois. (n.d.). *Composition of Religious and Ethnic Groups (CREG) Project* [Data set]. https://clinecenter.illinois.edu/project/Religious-Ethnic-Identity/composition-religious-and-ethnic-groups-creg-project

Uppsala University. (n.d.). *ViEWS - Department of Peace and Conflict Research—Uppsala University, Sweden*. Uppsala University, Sweden. Retrieved 11 June 2023, from https://www.pcr.uu.se/research/views

World Bank. (2022). *Literacy rate, adult total (% of people ages 15 and above)* [Data set]. https://data.worldbank.org/indicator/SE.ADT.LITR.ZS

World Bank. (2023a). *Individuals using the internet (% of population)* [Data set]. https://databank.worldbank.org/source/world-development-indicators/Series/IT.NET.USER.ZS

World Bank. (2023b). *Inflation, consumer prices (annual %)* [Data set]. https://data.worldbank.org/indicator/FP.CPI.TOTL.ZG

World Bank. (2023c). *Population, total* [Data set]. https://data.worldbank.org/indicator/SP.POP.TOTL

World Bank. (2023d). *Poverty* [Data set]. https://data.worldbank.org/topic/poverty

Appendix A: Detailed main model results

Academic variables

	LR	RF	GBM
Fragility		0.272346	0.111838
Durability		0.368759	0.128892
Democracy		0.173166	0.019982
FH_pol		0.145559	0.022761
FH_civ		0.173686	0.08142
Inequality		0.387193	0.165646
Poverty		0.376623	0.062535
Inflation		0.374403	0.12648
Literacy		0.396118	0.137601
Intervention		0.001962	0
Group Intervention		0.003151	0
Religious fragmentation		0.482574	0.239073
Globalization		0.274682	0.011785
GDP		0.799361	0.105381
Population		1	1
elecsys_Mixed		0.037144	0.006831
elecsys_Other		0.002777	0
elecsys_PR		0.03646	0.000249
elecsys_Plurality/Majority		0.029009	0.001189
elecsys_Plurality/Majority and PR		0.002825	0
Accuracy	0.574534	0.78882	0.754658
Precision	0.574534	0.909091	0.844156
Recall	1	0.702703	0.702703
ROC-AUC	0.833143	0.860229	0.847524
PR-AUC	0.876464	0.893104	0.888874

Professional variables

	LR	RF	GBM
Internet users		0.39518	0.133921
Intervention		0.004382	0
Group Intervention		0.00448	0
Religious fragmentation		0.503869	0.330728
Globalization		0.35199	0.131952
Population		1	1
US Trade		0.417158	0.162814
Weapon imports		0.55907	0.240204

Accuracy	0.574534	0.729814	0.714286
Precision	0.574534	0.901639	0.849624
Recall	1	0.594595	0.610811
ROC-AUC	0.83405	0.843717	0.798619
PR-AUC	0.87782	0.869492	0.84556

Combined variables (without GTD-derived features)

	LR	RF	GBM
Fragility		0.264008	0.126042
Durability		0.354663	0.112037
Democracy		0.133865	0.03578
FH_pol		0.120419	0.015392
FH_civ		0.177517	0.073486
Inequality		0.358087	0.186951
Poverty		0.272904	0.079155
Inflation		0.312976	0.063856
Literacy		0.317699	0.108647
Internet users		0.277125	0.054277
Intervention		0.001128	0
Group Intervention		0.005152	0
Religious fragmentation		0.409677	0.165393
Globalization		0.240312	0.03236
GDP		0.670226	0.059758
Population		1	1
US Trade		0.268008	0.0725
Weapon imports		0.515508	0.166543
elecsys_Mixed		0.027507	0.007633
elecsys_Other		0.002497	0
elecsys_PR		0.038426	0.002332
elecsys_Plurality/Majority		0.027664	0.002252
elecsys_Plurality/Majority and PR		0.00352	0
Accuracy	0.574534	0.736025	0.732919
Precision	0.574534	0.890625	0.851064
Recall	1	0.616216	0.648649
ROC-AUC	0.833143	0.864766	0.819235
PR-AUC	0.876464	0.890333	0.870035

All variables

	LR	RF	GBM
Terrorist attack		0.624978	0.886852

Fragility		0.2388	0.116801
Durability		0.311194	0.107429
Democracy		0.141474	0.048376
FH_pol		0.110173	0.037567
FH_civ		0.141954	0.059589
Inequality		0.353431	0.190948
Poverty		0.278824	0.063771
Inflation		0.269254	0.110531
Literacy		0.316492	0.137104
Internet users		0.2608	0.053307
Intervention		0.00453	0
Group Intervention		0.001885	0
Religious fragmentation		0.36119	0.179203
Globalization		0.223549	0.031799
GDP		0.611733	0.087259
Population		1	1
US Trade		0.30922	0.06664
Weapon imports		0.517821	0.189521
Global terrorist attacks		0.104336	0.033085
elecsys_Mixed		0.036812	0.008326
elecsys_Other		0.004287	0
elecsys_PR		0.029239	0
elecsys_Plurality/Majority		0.023115	0
elecsys_Plurality/Majority and PR		0.009384	0
Accuracy	0.574534	0.76087	0.770186
Precision	0.574534	0.875	0.888112
Recall	1	0.681081	0.686486
ROC-AUC	0.833182	0.870744	0.849891
PR-AUC	0.876496	0.892556	0.885725

Appendix B: Detailed model variant results

Barro-Lee education attainment data instead of literacy

	LR	RF	GBM
Terrorist attack		1	1
Fragility		0.232361	0.016777
Durability		0.416295	0.140231
Democracy		0.140212	0.023031
FH_pol		0.106637	0.006865
FH_civ		0.121295	0.006085
Inequality		0.413336	0.104036
Poverty		0.323233	0.064534
Inflation		0.360438	0.064801
Internet users		0.416077	0.115987
Intervention		0.003796	0
Group Intervention		0.050453	0.011986
Religious fragmentation		0.473945	0.14203
Globalization		0.30544	0.029608
Education		0.420478	0.125542
GDP		0.60759	0.089166
Population		0.994161	0.460385
US Trade		0.34376	0.099649
Weapon imports		0.535377	0.127785
Global terrorist attacks		0.188904	0.047191
elecsys_Mixed		0.040705	0.01402
elecsys_Other		0.006223	0
elecsys_PR		0.031387	0
elecsys_Plurality/Majority		0.047246	0.015975
elecsys_Plurality/Majority and PR		0.000803	0
Accuracy	0.479592	0.734694	0.72449
Precision	0.479592	0.744186	0.794118
Recall	1	0.680851	0.574468
ROC-AUC	0.739257	0.819358	0.851481
PR-AUC	0.688574	0.813556	0.841828

Fully interpolated global terrorist attack counts

	LR	RF	GBM
Terrorist attack		0.703117	0.897584
Fragility		0.300194	0.11913
Durability		0.340663	0.116939
Democracy		0.165121	0.045618

FH_pol		0.127051	0.039251
FH_civ		0.123289	0.058615
Inequality		0.393611	0.196125
Poverty		0.318268	0.076312
Inflation		0.297922	0.122058
Literacy		0.339605	0.135971
Internet users		0.283108	0.050213
Intervention		0.006578	0
Group Intervention		0.002912	0
Religious fragmentation		0.408972	0.172315
Globalization		0.238533	0.026123
GDP		0.698274	0.106358
Population		1	1
US Trade		0.315843	0.075424
Weapon imports		0.589208	0.194909
Global terrorist attacks		0.167166	0.03074
elecsys_Mixed		0.036047	0.008331
elecsys_Other		0.003251	0
elecsys_PR		0.04273	0.000115
elecsys_Plurality/Majority		0.030694	0.002066
elecsys_Plurality/Majority and PR		0.007559	0.003825
Accuracy	0.574534	0.773292	0.757764
Precision	0.574534	0.878378	0.902256
Recall	1	0.702703	0.648649
ROC-AUC	0.833182	0.873979	0.848372
PR-AUC	0.876471	0.899338	0.886804

No US trade

	LR	RF	GBM
Terrorist attack		0.741563	0.8591
Fragility		0.319964	0.089375
Durability		0.370561	0.103061
Democracy		0.163678	0.023765
FH_pol		0.152673	0.035861
FH_civ		0.15365	0.054902
Inequality		0.417736	0.196864
Poverty		0.367102	0.100503
Inflation		0.346834	0.121949
Literacy		0.411615	0.152846
Internet users		0.300355	0.054416
Intervention		0.0071	0
Group Intervention		0.002259	0

Religious fragmentation		0.444927	0.185031
Globalization		0.296181	0.023631
GDP		0.717009	0.109971
Population		1	1
Weapon imports		0.558883	0.187278
Global terrorist attacks		0.123669	0.012811
elecsys_Mixed		0.038293	0.013394
elecsys_Other		0.006558	0
elecsys_PR		0.042675	0
elecsys_Plurality/Majority		0.044878	0.004177
elecsys_Plurality/Majority and PR		0.007676	0.005019
Accuracy	0.570988	0.787037	0.777778
Precision	0.570988	0.886667	0.895105
Recall	1	0.718919	0.691892
ROC-AUC	0.833871	0.872623	0.854482
PR-AUC	0.875893	0.889312	0.892869

No GDP or weapon imports

	LR	RF	GBM
Terrorist attack		0.586108	0.839536
Fragility		0.232088	0.123955
Durability		0.316143	0.105436
Democracy		0.171903	0.049815
FH_pol		0.14069	0.019061
FH_civ		0.134875	0.03095
Inequality		0.358264	0.26624
Poverty		0.300742	0.09279
Inflation		0.29503	0.133944
Literacy		0.341362	0.145098
Internet users		0.285779	0.072165
Intervention		0.007352	0
Group Intervention		0.002544	0
Religious fragmentation		0.398009	0.167504
Globalization		0.272181	0.062755
Population		1	1
US Trade		0.318234	0.082804
Global terrorist attacks		0.119155	0.053943
elecsys_Mixed		0.038661	0.020367
elecsys_Other		0.005736	0
elecsys_PR		0.032856	0.001036
elecsys_Plurality/Majority		0.035424	0

elecsys_Plurality/Majority and PR		0.00971	0.002719
Accuracy	0.582133	0.783862	0.766571
Precision	0.63388	0.874214	0.879195
Recall	0.597938	0.716495	0.675258
ROC-AUC	0.715922	0.87169	0.860067
PR-AUC	0.800074	0.874022	0.889647