ERAMUS UNIVERSITY ROTTERDAM
ERASMUS SCHOOL OF ECONOMICS
International BSc Economics and Business Economics
Bachelor International Economics

On the impact of terrorism on exchange rate returns and volatility in OECD countries

Abstract

This paper aims to answer the question if terrorist attacks affect exchange rate returns and exchange rate volatility. It does this by ARMA-GARCH modeling 507 different terrorist events in 8 different currencies using daily exchange rate data. It finds no conclusive evidence on the link between terrorism and the level of the return of the local exchange rate. It does find that in times of terrorist activity, volatility appears to be significantly lower in the investigated countries. Lastly, it finds no difference for the effects when attacks occur in regions that regularly experience violence.

Author: Carlijn van der Linden
Student number: 415347
Thesis Supervisor: dr. L.C.G. Pozzi
Finish Date: 09-07-2019

1. Introduction

The shape and form of large-scale violence has changed in the past era. Since the second world war there has been decline in warring nations. Given the newfound interconnectedness of the world, nations find that the costs of going to war are higher, and the potential gains are smaller than ever. Through the decreased relative importance of owning land, increased importance of human capital and a general consensus that peace may be better than war, the major players on the world stage appear to have set their views on trade and attempts at cooperation rather than violence. These factors allowed for the rise of organizations such as the United Nations to further promote the decrease of international violence.

Though war is on the decline, civilians around the world now increasingly fear another type of danger. Transnational terrorism has been on the rise. With occurrences such as the very recent Sri-Lanka Easter attacks, the Paris attacks and, most notably, the Twin Towers attack of 9/11, we appear to have entered a new era of violence. One where danger is not a matter of a state of war or not, but an ever-looming danger from an unknown enemy. The link between these attacks is their unexpected nature, shock to the people and, in most cases, a religious or political motivation. In Europe, most nations have experienced one or more significant attacks. There are many different aspects to the impact of terrorism, which used to make up the largest of share of scientific knowledge on exchange rates. These are the political, sociological and historical factors, which used to be more commonly the subject of investigation in terrorism, than economic factors. More recently, the focus has shifted more towards the economic impact. The shock to a country can be enormous in the emotional sense. However, there can also be very serious economical implications. Attacks could disrupt markets, production and consumption. This paper aims to add to the existing knowledge of economic consequences of terrorist attacks by examining the effects on the local exchange rates. Exchange rates effectively price the assets in a country and, therefore, if exchange rates are affected by events like these, there needs to be scientific consensus about that. This paper hopes to achieve this by using financial models to find the effects that attacks have on the movements of exchange rates. Specifically, it will use event study ARMA-GARCH modeling to find a possible link between terrorist attacks and exchange rate movements.

Though violence still occurs world-wide, this paper will focus on attacks in OECD nations, mostly on the North American, European and Oceanian continent. The reason for this is twofold and lies both in the unattainability of covering every nation in the world, as well as with nature of the investigation and the availability of data. The increase in this particular type of violence, with widespread fear and media coverage, has been most apparent and in the developed world. Since large-scale violence is uncommon in these nations, one could expect that the effect will be greater. The second reason for choosing to research OECD nations is the relative abundance of data. Many undeveloped countries face poor data management, illiquidity and high volatility to their currencies. Gathering reliable daily data is not feasible if the currency is not traded in adequate volumes. Therefore, we again select westernized nations with liquid currencies.

Different effects to the movements of exchange rates could be observed in times of terrorism. This paper specifically investigates returns and volatility in exchange rates. As can be read in the literature review, the absolute level of exchange rates is not commonly found to be changed due to changing macro-economic variables. Therefore, this variable will not be investigated. In contrast, since they are standardized and comparable among currencies, returns to local exchange rates will be investigated. The other aspect to the exchange rate that might change is the volatility. Even though this is somewhat linked to returns and absolute levels, volatility is a separate concept that could produce very different results. Therefore, it will be examined as well.

Now, we formulate three null hypotheses that we will aim to disprove based on the notions outlined above. The hypotheses attempt to test the impact of terrorism on the exchange rate of the targeted nations:

1: There is no significant effect of terrorist attacks on the returns to local exchange rate data

2: There is no significant effect of terrorist attacks on volatility of local exchange rates.

The recurrence of attacks and the possible desensitization towards the effects is another of the main questions in this paper. One could hypothesize that recurrent attacks may lessen the possible effect on the exchange rate. Since we are interested in how these effects shift if a region more regularly experiences terrorist violence, we formulate the third hypothesis:

3: The effect of terrorism on the returns and volatility of local exchange rates are significantly lower in regions that are regularly targeted.

To test these hypotheses, daily exchange rate data (1970-2018) from 8 different currencies is used and paired with 507 individual terrorist events. Firstly, an event-by-event approach is applied, which treats every event as a separate ARMA-GARCH model that aims to measure the effect of the occurrence of terrorism when compared to periods that experience no terrorist attacks. We will collectively investigate the significance of the 507 separately generated coefficients using a test-statistic. Secondly, the effect on the 8 different currencies are tested. This is done by aligning all daily exchange rates from 1970 until 2018 and assigning all days that saw a terrorist attacks in the past 14 days a dummy value of 1, and testing for significance of the dummy variable. The third hypothesis will be tested by running analysis on three different subgroups of the tested population in the United Kingdom. In the conflict in Northern Ireland, during the so-called troubles, terrorist attacks regularly occurred. In testing the third hypothesis, attacks in Northern Ireland, the rest of the United Kingdom and the United Kingdom as a whole are tested as different populations and their outcomes compared to find a possible effect. We are curious to see if Northern Irish attacks are less impactful on movements in exchange rates than attacks in other parts of the country due to desensitization. The currency-based data alignment will be used in testing the third hypothesis.

Due to the large number of parameters and the way the data is constructed, some events had to be dropped from the investigation. In these cases, estimating the GARCH coefficients was not possible. This lowered the amount of observations in some of the estimations. Running the methodology with the lowered amount of observations, the paper finds that returns are ambiguously affected by terrorism. In the event-based run they are negative and significant. In the currency-based run, the effects are insignificant. Therefore, no conclusion can be drawn about returns to the local exchange rate. In regard to the second hypothesis, volatility was found to be significantly lower in both the runs of data. This effect is presumed to be because of tendencies within the data, and not an actual calming effect of terrorist attacks. This is outlined in the results section. For the third hypothesis, the effects on desensitization on the matter of Northern Ireland are inconclusive.

In the next sections the reader will find the body of the research. In section 2 we will investigate existing literature on the effects that terrorism can have on economic systems. In section 3 there will be a more elaborate justification of the data and the selection of the different subgroups. In section 4 there will be an overview of the methodology that is used in the research. In section 5 there is a detailed outline of the results that were generated by the methodology. All sections will be summarized and concluded upon in the sixth and final section, the conclusion.

2. Literature review

In the three headers below, an overview of the currently existing knowledge on the topic will be drawn. First, we will discuss the existing knowledge on the links between terrorist attacks and economic variables. Secondly, we will briefly touch on the subject of exchange rates and the earlier scientific modeling about how they tend to move them in general. Lastly, we will combine the findings under the first two headers and, combined with some existing literature on terrorism and exchange rates, attempt to make a prediction of the possible outcome of our statistical study as described under methodology.

Existing knowledge on impact of economic variables

Firstly, we need a clear picture of the known impact of terrorist attacks. Many variables tend to move together in somewhat predictable ways. By mapping out an array of different economic variables, we will able to make better predictions on the outcome of our analysis. This part will first elaborate on the impact of terrorism on long-term macro-economic variables. After that some financial variables that move in the short-term as well are discussed. Literature directly linking exchange rates and unrest can be found under section 3 of this literature review.

Marco-economic variables

The first variable that will be examined is economic growth. Since 9/11 is one of the most impactful attacks modern history, it seems a good place to start discussing the impact on growth. In the short-term one of the first notable effects was the decline in consumer confidence. After the attacks, GDP forecasters corrected their predictions by 0.5 percent point for 2001 and 1.2 percent point in 2002. The net GDP effect in the short run has been calculated to be 500 billion USD. Industries that were hit hardest were related industries which are discussed in the next section. (OECD, 2002)

Another paper that elaborates on growth in under terrorist threat was written by Abadie and Gardeazabal. It is one of the largest studies into terrorism and is a case study into the conflict in the Basque region in Spain from 1970 until early 2000. In this study, researchers were able to compare the effect of non-economically motivated attacks by the ETA, a local terrorist organization, to other Spanish regions instead of other nations. This approach together with the broad time horizon allowed the team to better isolate the conflict-specific effects. The aim of the study was to capture the differences in GDP growth, by creating a synthetic Basque region that, if designed correctly, would behave as the Basque region would in the absence of violence.

Researchers found a 10 percent difference between GDP per capita of the synthetic region, presumably without the effect of terrorism, and the real Basque region. This negative effect on the Basque region is hypothesized to be a result of diverted foreign investment because of economic uncertainty. However, some factors such as relatively low population growth in specifically the Basque region, may be unrelated to the terrorist threat but still present in any of the observations. The paper concludes that it is not necessarily the permanent existence of a terrorist threat that is the best predictor of loss of potential growth, but rather the level of intensity of the terrorist threat in a given year. (Abadie & Gardeazabal, 1999).

Inflation is another part of the macro-economic framework. It is only very rarely covered in existing literature in the context of terrorism. In one study, Shahbaz investigates the inflation rates in Pakistan in relation to the domestic terrorism in this nation. Using causal networks, thus allowing output, terrorism and inflation to influence each other, the paper found that terrorism causes inflation. Additionally, it also finds cointegration between output, inflation and terrorism. The increased output and inflation appear to increase terrorist activity. A first policy suggestion that is made is to limit the inflation as it appears to cause increased terrorism. As a second solution in terms of policy the paper suggests increasing employment.

Unemployment as a variable is used more often in this debate. Unemployment is found as a significant determinant of terrorist activity, unlike other factors such a relative deprivation and tertiary education. (Richardson, 2011). More studies like these have used unemployment as an explaining factor in the debate of terrorism. In a large study on the effects of terrorism on the economy in Palestine, it was found that a successful terrorist attack caused a fall of between 0.52 and 0.66 in employment on the national level. Locally, the effects are even greater and are a minimum of 5.3%. (Benmelech, Berrebi, & Klor, 2010). There are also papers that explain other macro-economic variables in the context of social instability. One is (Gupta & Venieris, 1986) underlines the effect on savings and investment.

Concluding from the papers above, it is safe to say that there is a significant and important role in the macro-economic variables when talking about the impact of terrorism.

Financial variables and businesses

A large body of research regarding terrorism exists in favor of investors and deals with the impact of financial markets. This means that there is a lot of knowledge into the reaction of stocks in the presence of terrorism. Chesney, Reshetar and Karaman (2011) provide us with many interesting conclusions regarding the effects on global financial markets. After investigating different stock markets worldwide, they found that around 70% of terrorist attacks significantly affect at least one of the measured financial markets. There is a notable difference in the susceptibleness to shocks among the different markets. This may be due to either overall resilience of their respective financial system, the degree of international activity of the companies listed on the local stock exchanges or the size of the economy (Chesney, Reshetar, & Karaman, 2011).

Additionally, for the Basque region, instances like the cease-fire in the period of 1998-1999 created a natural experiment which allowed researchers to observe the economy in relative absence of danger. If the attacks were having a short-term negative effect on financial variables like stock prices, one could assume that these effects would diminish during a period of cease-fire. For this study, using the daily stock prices of different Basque firms were selected and cumulative abnormal returns were calculated. Using dummy variables for periods of good news and bad news, the effects of periods after good and bad events was measured. This yielded no significant effect. However, this approach simply assumed that, until a new 'event' took place, the sentiment on the market would remain the same until another event would take place. When narrowing the event window, to the trading dates after the events took place, both the positive and negative news was found to have a significant effect on the stock prices of Basque companies. The market model, constant mean return model and the Fama-French three-factor model all drew the same conclusion. (Abadie and Gardeazabal, 2003)

Another paper investigating the stock market was written by Eldor and Melnick. It focuses only on the domestic Israeli market. As well as (Chesney, Reshetar, & Karaman, 2011), it finds significant effects of effects on stock markets. This paper also aimed to investigate the attack-specific factors that were significant to the magnitude of the effect. The number of casualties and the amount of attacks on a day are both significant for the degree of depreciation in the Shekel. It also finds that suicide bombings have the biggest impact and that this effect is permanent. Also, because the long timeframe of the investigation, the paper can also conclude on a possible apparent desensitization towards terrorism. It finds that this does not occur. The effect of the recurring attacks on stock markets remains the same throughout the sampled period. This is an interesting result towards the third hypothesis. (Eldor & Melnick, 2004)

Some industries are intrinsically more susceptible to the consequences of terrorism. The most notable of which, next to finance and banking, are the transport business, such as airline and traveling industries, insurance and oil and gas. By measuring indices for all these industries, it finds that the negative effects to insurance and airline stocks are the largest, and banking and finance are least affected. This is makes sense intuitively. For airlines and travel agencies, large-scale insecurity and fear are bad news, since it may stop a consumer from booking a flight to a recently struck country, causing their stock to fall. As for insurance, the nature of their business is to pay out when disaster strikes. A large-scale disaster like an intended terrorist attack may diminish profits, and therefore stock returns. (Chesney, Reshetar, & Karaman, 2011) The same was found by the OECD 9/11 report. Among the affected industries were aviation and airlines, the tourist sector and insurance companies. However, the Federal reserve and congress actions were swift and rigorous, limiting the damage that could have arisen to the economic system as a whole. (OECD, 2002). These findings are also backed by (Abadie and Gardeazabal, 2003).

Chesney, Reshetar and Karaman investigate commodities as well. Most notably, since gold is regarded as a safe haven in times of insecurity, the price jumps in the short-term after an attack. All consequences mentioned above are measured in the short-term. The paper sets an event window of six days in which it hopes to capture the effects of the event. It does, however, also differentiate between the effect on the short-term, as well as the immediate effects on the event date itself. Some markets, such as bond markets, appear to react on the event date itself, whereas the gold price only reacted several days later. Stocks seem to move significantly on both the event date and the period thereafter (Chesney, Reshetar, & Karaman, 2011). As for 9/11, the short-term financial variables were affected as well. After the attacks, a brief period of panic caused investors to flee from equity and products that were regarded as riskier. Spreads on corporate and government bonds and US and emerging markets bonds widened.

There is also insight towards the second hypotheses, which is about the exchange rate volatility. This is paper that investigates volatility rather than simply the returns of stocks. The researchers define terrorist volatility as a movement in the index of terrorist activity. The paper tests the links between volatility in the stock market and the terrorist index using a VAR-GARCH framework and daily stock and interest rate data. The first conclusion it draws is that volatility in the terrorist index corresponds to a negative significant effect in returns. This conclusion is similar to the conclusions drawn before. The second test, which attempts to explain stock price volatility by terrorist index volatility, extends earlier conclusions. It appears to affect stock return negatively. The sign is irrelevant in this case, but it tested to be significant. Here, again, notable differences between different countries were found. European nations tended to not react as strongly as did the other nations. (Peren Arin, Ciferri, & Spagnolo, 2008)

Framework on exchange rates

This section will provide a framework to explain levels movements in exchange rates. It will start by making a basic introduction on research on the levels of exchange rates. After that it will elaborate on the small body of research that does attempt to draw parallels between exchange rates and the real economy.

A good place to start is by asking the question how an exchange rate is determined. The most basic model for explaining exchange rate levels is the monetary model. It assumes that today's exchange rate between two countries is determined by underlying economic variables. Most notably, it uses relative money supply, national income and relative price levels to explain the difference. (Copeland, 2014) When the model is combined with market predictions it creates a logical framework to explain exchange rate variation. However, since the impact of the different economic variables, as well as the value of the underlying variables, can be hard to accurately observe. When empirically tested, the model lacks the fundaments needed to provide a solid basis for conclusion. (Frenkel, 1976) As an extension, both the Dornbusch overshooting model and the portfolio balance model were created to solve this lack of explanatory power. However, neither solved the problem of empirical weakness.

If these fundamental underlying economic variables fail to help the explanation of levels of exchange rates, but they are still very closely linked to them, then the changes that occur to the level of these fundamentals may explain movements in the levels of exchange rates. This is the central idea to the news model. Since markets always have expectations, the news model uses the difference between the expected and real change in the fundamental variables as the explanatory variables in exchange rate fluctuations. When tested empirically, the news-component appears to explain somewhere between 5-20% of the otherwise unexplained variation. This outcome is somewhat unsatisfactory result is due to the sheer volatility (Copeland, 2014). Another problem to emphasize again is the possible unobservability of the fundamentals. Either we might not be able to measure them accurately, or we might not know which variables are underlying to exchange rates. Both problems would pose serious issues to the power of the news model. An alternative view is the determination of exchange rates through market sentiment. This framework does not eliminate the idea of market fundamentals, but rather uses them as a basis for the valuation. The equilibrium price, however, is not intrinsic in the fundamentals in this view, but rather a product of evaluation of all buyers and sellers in a currency market. The buyers and sellers in this market also use fundamentals in their evaluation. Even though the level of exchange rates is notoriously hard to model, the volatility aspect does yield better power of predictability. (Hopper, 1997)

Another assessment of the fundamentals and their underlying effect on exchange rates was written by Samo & Schmeling (2011). They note the fact that some papers before them found weak links between fundamentals and exchange rates and use this notion. They aimed to test if current exchange rates affect future fundamentals and which fundamentals are the most important. It achieves this by creating portfolios of countries based on their lagged spot rate change against the USD that are updated every period and use these different portfolios to find patterns in their movements. They found that, portfolios that depreciated significantly against the dollar, experienced significant inflation as well as a significantly lower output growth differential. Overall, they conclude that spot exchange rates have predictive power on future fundamentals. In addition, future fundamentals seem to have significant effect on current spot exchange rates. These conclusions are limited to nominal macroaggregates. Real macroaggregates do not adhere to these predictions. If anything, this paper suggests a link between expectations of future fundamentals and spot exchange rates. In other words, market expectations matter on the subject of nominal exchange rate determination. (Sarno & Schmeling, 2011).

One last crucial factor in the ability of an exchange rate to take on a value that is representative is the exchange rate regime that is put in place by the central bank of a nation. After World War II the Bretton Woods system was administered for two reason. Firstly, the rebuilding of Europe had to be facilitated and, secondly, to prevent the aggressive appreciations and depreciations that occurred in the past. This meant there were controlled, fixed rates that were readjusted when necessary and the United States Dollar adhered to the Gold Standard. This means that they would always be willing to trade an ounce of gold for a set amount in dollars. Other countries fixed their currencies to the dollar. Starting from 1973, the Bretton Woods fixed rates were let go one by one and switched to floating rates. This allowed every currency to freely depreciate and appreciate to supply and demand. (Copeland, 2014) An important question to ask here is if the exchange rates regime play a role in the adaptability to shocks, such as a terrorist attack. According to Meade, for open economies, shocks are incorporated in their nominal exchange rates. If these nominal exchange rates are fixed, the shocks are redirected to nominal prices and domestic wages. (Meade, 1951) As a more recent justification, Edwards and Yeyati (2004), tested the conventional wisdom of flexible exchange rates being a better regime for taking in shocks. They too found that floating rates have a better capacity of reducing the impact of real shocks in terms of economic output. This shock absorption is through the adjustment of the nominal exchange rate, which is one of the variables of interest in this paper. The evidence suggests that the regime selection is an important factor to consider in analyzing the effect of a shock, such as terrorist activity. It also predicts that the largest movements in nominal exchange rates would be visible under freely floating exchange rate regimes (Edward & Yeyati, 2004).

In conclusion, attempting to model the level of exchange rates through macro-economic modeling is challenging. However, these are the best-known models of explaining exchange rate levels that we have. Therefore, no conclusions will be drawn solely on the basis of either of the above-mentioned models, but it does allow us to put the events that may explain exchange rate movements in a perspective. In answering our first hypothesis this paper will contribute to this body of knowledge. Researching volatility rather than absolute levels appears to be promising, Hopper found, which is investigated in our second hypothesis.

Terrorism and exchange rates

The description of terrorist events in the context of movements in the financial markets provide us with a framework to make some basic predictions. We will investigate some literature that deals with the terrorism and exchange rates specifically and combine this with the knowledge gathered before to make a prediction.

In a paper by Eldor and Melnick (2004), the focus is solely on the Israeli financial market. As well as the implications for the stock market described above, it also tests the effect of terrorism on national level in terms of currency exchange rates. The currency exchange today is daily and in large enough volumes to immediately reflect any changes in exchange rate fluctuations due to shocks. In the past there have been some different regimes that the paper takes into consideration. In the regression analysis the paper makes, there is indeed a significant negative effect of a terrorist attack on exchange rates. There is, however, a very low R^2. This is supposedly due to the use of monthly data in some part of the equation. In a news-model regression, there is no permanent significant effect on the exchange rate. (Eldor & Melnick, 2004) The short-term focus of our paper seems suitable to investigate this matter.

One particular paper is very closely related in their research question. Melvin & Tan (1996) investigated the bid-ask spread of the South-African Rand versus the United States dollar over a period of 1242 observations. What they aimed to find out was if there exists a possible link between outings of violence and the spread on this local currency. Since bid-ask spreads are a measure of liquidity of a currency, meaning constant and relatively equal numbers of supply and demand, a shock in money demand on the markets, maybe due to uncertainty and fear, may increase uncertainty about the future and therefore affect the conditional variance. Similar to our approach, Melvin & Tan modeled exchange rate data using GARCH modeling, including a dummy for violence in the volatility term. South-Africa, in those days, regularly encountered violent attacks. Around a quarter of the days that were measured in total, saw one of the types of violent behavior. They divided these attacks in four groups and measured among other effects, their severity per type. These groups were riots, violent demonstrations, political strikes and armed attacks. They also tested death count as an explanatory variable for a possible increase in the spread to account for the magnitude of unrest. The conclusion they drew is that every type of attack, as well as the death count are significant as explanatory variables of the exchange rate spread. In other words, sociopolitical disturbance significantly affected the foreign exchange rate markets. This research was carried out in one nation over a longer period of time. To increase the validity of their conclusion, they did a cross-section examination of 36 countries. They used risk indices to measure the amount of social unrest in all nations between 1987-1990. The research shows that the bid-ask spread, again, was significantly higher in nations with higher risk indices. This result does not directly link to exchange rate volatility, but it does underline the fact that social unrest can influence exchange rates in some meaningful way (Melvin & Tan, 1996).

Expectations from the existing literature combined with literature review

The main macro-economic variable researched was GDP growth. According to the fundamentals theorem of exchange rate modeling, this is one of the fundamentals to exchange rates. Consequently, any change in the future expectation of GDP growth should, by the news model, reflect for around 5%-20% in the exchange rate. Since terrorism causes a shock like this, we might expect, through an expectation of slower future growth, a depreciation in the currency. This is of consequence to our first hypothesis. Also, we have learned that there a link can be proven between future market fundamentals an current exchange rates, further strengthening this assertion. (Sarno & Schmeling, 2011)

As for stock prices, we saw earlier that terrorism tends to decrease stock returns. We can read later on that stocks returns and exchange rates are indeed correlated. First, as we saw earlier, Peren Arin, Cifferi and Spagnolo (2008) found that increased variance in terrorist activity also increased variance in stock prices. Smith established a model to estimate exchange rates with stock prices as an explanatory variable. He found a significant relationship between equity returns and exchange rates, but not vice versa (Smith, 1992). In his paper, Kanas investigated the volatility spillovers from exchange rates to stock returns and vice versa. The aim was to extend the previous results that found a relationship between stock returns and exchange rates. The second moment, or in other words, the variance of the variables was tested for granger causality. Kanas found significant results in 5 of the 6 countries examined. There is a significant volatility spillover from volatility in stock returns, which, like the mean of the returns works from, only works from stock returns to exchange rates and not vice versa. The spillovers are symmetric. However, this does again give us an extra hint towards answering our second hypothesis. If terrorism causes stock prices to become more volatile, and under normal market circumstances, volatility from stocks tends to spill over into exchange rate volatility, we would not be surprised if exchange rate volatility tends to increase in times of terrorist threat (Kanas, 2003).

Lastly, simply as a piece of historic data, we again consider the source from the 9/11 chapter. In this particularly impactful event, the exchange rate volatility that we aim to measure in this paper, has indeed increased significantly, along with the volatility of other variables.

From (Eldor & Melnick, 2004) and (Melvin & Tan, 1996) as well as from some other papers, there are signs that our hypotheses may be disproven. From our research on exchange rates economics, we know that our first hypothesis on returns will be hard to disprove, since finding any fundamental explanatory variable in exchange rates levels has been challenging. As for the second hypothesis, we learn that volatility in exchange rates has produced results before and is therefore one of the most promising of our hypotheses. Additionally, many of the significant relationships described above, were on the topic of volatility, rather than absolute values. For the third hypothesis, which examines the difference in effects for a region that regularly experiences violence, very little literature was available. The literature that has been written is not in favor of desensitization towards shocks due to terrorism. This also may be a hard hypothesis to disprove.

3. Data

The aim of this investigation is to find if there is an effects to the returns and volatility in local exchange rates in times of terrorist activity in different countries. To this end there is a need for both reliable data on terrorist activity, as well as good exchange rate data. This section will deal with the selection of the event dates for the terrorist activity. As described in the introduction the aim of this research is on terrorist activity in OECD nations. The selections decided upon are specified and justified in this section.

Data on terrorism

The dataset that will be used for the retrieval of event dates is the Global Terrorism Database, or GTD for short. It is an extensive database with over 180000 instances of violence. It tries to collect many different aspects of the attacks, such as geographical location, nature and the number of fatalities. The literal definition the dataset attributes to the 1970-1997 terrorist attacks are as follows: "the threatened or actual use of illegal force and violence by a non-state actor to attain a political, economic, religious, or social goal through fear, coercion, or intimidation." (University of Maryland, 1970-2018). Like many controversial topics, many aspects of what constitutes a terrorist attack are hotly debated and starting from 1998, GTD decided to include three more criteria.

1. The violent act was aimed at attaining a political, economic, religious, or social goal;

2. The violent act included evidence of an intention to coerce, intimidate, or convey some other message to a larger audience (or audiences) other than the immediate victims; and
3. The violent act was outside the precepts of International Humanitarian Law.

These criteria make the data in the set useful for answering the central question of this thesis. As mentioned in the introduction, we chose to focus on non-state actions so in that sense the data fit well with our wishes. Another benefit of this database is the sourcing of the data. Several respected institutions have collected data in several stages of the database's existence. Additionally, it sources from other archives on terrorist data. Whenever new information about a historical event surfaces, the file is updated to display the most recent knowledge.

The vast dataset provides us with enough observations to make a good selection of events that are of interest in answering our research question. The scope of worldwide investigation, with all countries on earth included in the research, is far beyond the scope of this effort. Collecting and processing all currencies and event dates would not be practical nor feasible. In the present, with the internet and many other means of direct communication at our disposal, many developing countries still have problems with regularly and liquidly updating their exchange rates on their official websites. Expecting to find reliable exchange rate data, going back as far as 1970, is not a realistic ambition. For this reason, as well as the reasons outlined earlier, undeveloped countries will not be included in the research. The aim is to take a broad and usable dataset in the first run of our statistical analysis, and to subsequently cut out observations to better cater to the intended scope of OECD and developed countries. Using the full GTD set, a step by step description of the selection is described below.

Firstly, we limit ourselves to instances with at least two fatalities. This choice is only made to specify a base level of what a it means to have a significant impact. In this case, 2 is chosen to differentiate a terrorist attack from a murder. Secondly, the first geographical cut to a sample will be to select a set of relatively developed regions. These will be restricted to the continents of Northern America, Western Europe and Oceania. The reason for this cut is mainly due to the availability of data. Most OECD nations have trustworthy supply of exchange rate data. These include pre-euro European nations and other nations that may be of interest in this research. However, many of the data are weekly or monthly, making it incompatible with the short-term focus of this investigation. The countries that are left for research are the United states, United Kingdom, Switzerland, Australia, Canada, Norway, Sweden as well as a dozen of European Countries after 2000 that use the euro. These OECD countries have daily and well-documented exchange rate data without any significant holes or problems. This set is compiled of 507 observations to be tested. The final transformation to be made is the removal of the 1970 occurrences from the data set since the exchange rate data starts in 1970. Since we also need exchange rate dates before the events, these observations cannot be used in this research. The total for this set is 507.

The full set of 507 observations will be the set used in order to test the main hypotheses in this paper, the effect on the absolute level of exchange rates and the volatility of exchange rates. Note that this selection includes some regions that are known for violent attacks, such as the Northern Ireland. In the first set, this region is included to potentially observe a difference in the including or excluding of these regions. In testing the third hypothesis, Northern Ireland, which experiences attacks more regularly, will be tested separately from the rest of the United Kingdom. The intuition here is that the exclusion of these regions may introduce a clearer shock to the exchange rate variables. Regions that tend to experience attacks more often, may be less likely to experience high volatility in their currencies due to the attacks, since the attack is less of a shock and effects of any previous effect still may be visible.

Exchange rate data

The exchange rate data is drawn from the website of the federal reserve of the United States of America and runs from 1970 to 2018. Generating and publicizing data and reports is one of their main purposes and as one of the largest monetary institutions in the world, their data can be regarded as reliable. Since the window of this research is short-term, monthly data is not useful. All nations that experienced terrorism in the full set that were eligible for investigation were reviewed for availability of daily local exchange rate data. All those for which it was were put in the full sample as described before.

If any observations were missing for a small subsequent amount of days, the value of the previous business day was used. An omitted value could be due to technical failure or a holiday. For our intents and purposes, this is not expected to have a significant impact. Also, weekends have been ignored in this paper and treated as if Monday directly followed Friday.

For the examination we will use returns to measure performance of the currency measured. The benefit of using returns when compared to just using absolute values is that it allows for better comparison between points in time and different currencies. It does this by transforming the data into measures that are not affected by the initial price level. For our intents and purposes, we transform the above formula into exchange rate data. The returns per currency are defined as follows:

$$R_t = \frac{c_t - c_{t-1}}{c_{t-1}} \, x \, 100\%$$

With c_t being the nominal exchange rate at time t. The set of countries that will be investigated has been selected on the availability of daily exchange rate data.

The reference currency, or the currency as to which the measured currency is expressed in, is important because it could potentially influence the measurements. If a terrorist event has international exposure, the change in exchange rate could be enhanced or diminished if the event would affect the reference country as well, thus skewing results. This issue is hard to circumvent completely. In this case, we chose consciously for USD. USD is used as a reference currency because of its role in international finance as the reserve currency of the world (Flandreau & Eichengreen, 2009). The United States are big and can be assumed to not be affected significantly by violence abroad. The USD events cannot be measured in terms of itself. We need another currency that is liquid as well. Yen is also a liquid currency with daily available data and was preferred because it was not used elsewhere in the research.

Note that every currency that belongs to the countries in our dataset as specified above, have a freely floating exchange rate regime. The only exception to this rule is Switzerland.

Datasets

In order to test both hypotheses, we will use two different alignments of the data. The first type of composition is the event-based alignment. For this data, a panel data approach is applied. All 507 events are included in the set separately and identified by an ID (1-507). For every event, the respective exchange rate of the 250 days before and the 250 days after the occurrence of the event are included. This makes for a total amount of observations of 500 days per event. For clarity's sake, this means that 507 times 500 rows of data are used in this composition. On the basis of these daily data points, the returns, as specified above, are estimated. For every event in the 507 occurrences, a separate statistical test, as specified in the methodology, is run. As for the testing of the influence of terrorism, the 14 days after the event date are paired with a dummy taking on a value of 1, while the other dates take on a value of 0. These 14 days are called the event window.

The second composition is in terms of currencies and is thus referred to as currency-based data. Every country has its own exchange rate data for every day from the observed period, which for most currencies is 1971 until 2018. This leaves us with 8 sets, one for each of the currencies listed in the data section, that have the exchange rate data from 1971 to 2018 in one column. These allow us to calculate returns. The event dates are then matched with the dates in the exchange rate data. In another column the dummy takes on a value of 1 if an event date occurred within 14 priors. This means that every date in the time frame investigated is paired with either a

0 or a 1 in terms of occurrence of terrorism. Here, the event-window thus works similarly to the event-based alignment. The dummy variable for terrorism thus allows to test the dependent variables in the hypotheses, returns and volatility.

Both alignments test a 14-day ex-post period for suspected influence of terrorism. Notably, since we assume that terrorist events are not suspected by the financial system because they are surprise attacks, we do not test dates before the event. The 14-day length is another matter. Since we want to investigate the shock—effect it makes sense to select relatively small number of days following the event. However, GARCH calculations run into multicollinearity issues when the dummy variable is 0 for a too large fraction of the dataset. This means that, for instance, taking a window of 2 days will not work with the methodology selected. To strike a balance, 14 days was chosen.

Both sets of data will be used in different runs and approaches of the methodology. The statistical operations that are applied to this data are outlined in the next section.

4. Methodology

Now follows a brief description of the methods that will be applied in order to test the three hypotheses. Great progress has been made in the past decades in the field of modeling financial variables. Early models tended to rely on the OLS regression, such as the Market Model as introduced by Sharpe (1963). This model attempts to explain the return of a stock portfolio i, by fitting the return of a market portfolio with a parameter beta, specific to the stock portfolio. This beta measures the inherent risk of the stock portfolio in terms of the market risk. In this model, the stock portfolio return is assumed to be determined by the heteroskedastic scalar that magnifies the return of the market portfolio symmetrically. The market model assumes constant variance throughout the time series, otherwise known as homoskedasticity. in later research, this assumption was disproven. The distribution of returns of stocks turned out to be a heteroskedastic process (Giaccotto & Ali, 1982). To solve the matter of heteroskedasticity, methods that allowed for a change in variance throughout the time series needed to be introduced. That same year, the autoregressive conditional heteroskedasticity, or ARCH model was created. Here, and in the rest of this section y_t is used as an example and is unrelated to any of the expressions our research. (Engle, 1982) created a model where an exceptional observation in y_{t-1} can change the conditional variance y_t. This is done through a contemporaneous estimation of the parameters in two expressions, one normal OLS regression and one expression that models the residuals of this expression. Basically, it applies the Box-Jenkins (1970) ARMA approach to the volatility of the residuals. (Brooks, 2014) Thus, it solves the problem of volatility clustering in economic data. A further extension was introduced a few years later. This time the new modeling allowed for the previous observation of σ_{t-1} to be incorporated into the conditional variance. This method by (Bollerslev, 1986) was called generalized autoregressive conditional heteroskedasticity, or GARCH. Since GARCH is more the parsimonious model (Brooks, 2014) in solving the issue of volatility clustering, we will be applying this method to our hypotheses. However, since this research does not focus on stocks, some adjustments need to be made in order to fit the hypotheses. What is needed is a transformation of the market model to accommodate for exchange rates and an influence of terrorism.

Event-based data

As outlined in the data section, the second type of data is on the event level and takes on a panel data approach. As mentioned earlier, due to the mathematical difficulty in running GARCH on panel levels, a different approach is applied. In testing the first three hypotheses on the level, the following AR-GARCH will be run the event-based dataset:

$$R_{i,t} = p_i + \zeta_i R_{i,t-1} + \gamma_i D_{i,t} + \varepsilon_{i,t} \qquad (1)$$
$$h_{i,t} = a_{i,0} + a_{i,1}\varepsilon_{i,t-1}^2 + \lambda_i h_{i,t-1}$$
$$\text{With } h_{i,t} = E_{t-1}(\varepsilon_{i,t})$$

For this expression, parameters are estimated through the maximum likelihood function. The dependent variable here is defined in the data section. Their first lags are included in the regression and the explanatory effect measured by their respective parameters. As mentioned, the initial dataset that is going to be used consists of observations which are ordered by 507 events i. As for the other parameters, p_i is a constant specific to the event, ε is a stochastic shock term that is assumed to be normally distributed with a zero mean and a conditional volatility of $h_{i,t}$. Volatility $h_{i,t}$ in turn is derived from the main regression and modeled by ε and the ARMA(1,1) terms $h_{i,t-1}$ and $\varepsilon_{i,t-1}^2$. The AR(1) term in the OLS regression is used because the research does not aim to find any factor that may affect the returns, other than terrorism. Instead, it builds on the random walk aspects that exchange rate data tends to follow, the best guess of γ_{t+1} is γ_t. Since we are not looking for any specific explanatory variable outside the terrorist threat, all other factors are included in the AR(1) term.

The first of the two equations above model a basic Ordinary Least Squares regression (OLS). In this type of regression, the dependent variables and independent variables are mapped on a n-dimensional plane, with n being the number of variables in the estimation. Then, minimizing the vertical squared distances from all the mapped coordinates, the optimal fit line is drawn through the points. The vector that describes this line is the set of estimated parameters. Multiplying by transposed vector of variables leaves us with the estimation above. The estimated parameters tell us approximately, the correlation between the independent variable and the dependent variable. Every parameter is tested for significance and so, we have a good estimation of the effects between variables. (Brooks, 2014). The second term in the contemporaneously estimated regression models the conditional volatility. It models the expected volatility for any observation t, based on the past observations in $\varepsilon_{i,t-1}$ and $h_{i,t-1}$. In doing so, it succeeds in capturing the volatility clustering effects as described earlier. In other words, it allows for changing volatility, whereas an ordinary OLS would not. The above expression will be investigated for every observation or event i. This will help us answer our first hypothesis. For the second hypothesis, we need an expression that can capture the effect of a shock in the dummy variable for terrorism in the volatility term of the GARCH expression. This expression is formulated as follows:

$$R_{i,t} = c_i + \beta_i R_{i,t-1} + \varepsilon_{i,t} \qquad (2)$$

$$h_{i,t} = a_{i,0} + a_{i,1}\varepsilon_{i,t-1}^2 + \lambda_i h_{i,t-1} + \delta_i D_{i,t}$$

$$\text{With } h_{i,t} = E_{t-1}(\varepsilon_{i,t})$$

The dummy variable and its parameter follow a process called multiplicative heteroskedasticity. This method allows for explanatory variables in the volatility term of the GARCH estimation. The parameter δ_i will capture the effect of $D_{i,t}$ on the volatility $h_{i,t}$. Which, we recall, in turn is the modeled conditional volatility to $R_{i,t}$ where ε is a stochastic shock term that is assumed to be normally distributed with a zero mean and a conditional volatility of $h_{i,t}$.

The first run of the data will run the ARMA-GARCH expressions separately for every observation or event i. Therefore, there will be 507 separate estimations of the parameters. Since these runs will not produce a readily interpretable result, a standardized test statistic is required to test for significance. The estimated parameters of all observations will be processed through the following test statistic, which follows a t-distribution. $\hat{y}_i$ and $\widehat{\delta_i}$ are estimated running the data. (Balaban & Constantinou, 2006)

$$Q_1(\hat{\gamma}) = (\frac{\sum_{i=1}^{n}\hat{\gamma}_i}{n})/\sqrt{\left\{\left(\frac{1}{n(n-1)}\right)\sum_{i=1}^{n}[\hat{\gamma}_i - \sum_{i=1}^{n}\hat{\gamma}_i/n]^2\right\}} \quad (3)$$

$$Q_2(\hat{\delta}) = (\frac{\sum_{i=1}^{n}\hat{\delta}_i}{n})/\sqrt{\left\{\left(\frac{1}{n(n-1)}\right)\sum_{i=1}^{n}[\hat{\delta}_i - \sum_{i=1}^{n}\hat{\delta}_i/n]^2\right\}} \quad (4)$$

This test statistic tests the average coefficient by weighing it by the average variance. This serves our goal to produce a readily interpretable result which will be outlined in the results section. Even though we have standardized data, we still need to be cautious, since events with outlying volatility will tend to dominate the effects on this statistic. This will be described in the results section.

Currency-based data

Recall that the other type of data was structured in terms of currencies. This gives us multiple events per currency. It also allows us with a longer timeframe per event, namely, either starting from 1971 or 1999 for Euro until the present. This increased timeframe and the increase of observations it produces allow for estimations with more parameters. The same GARCH methodology as applied earlier will be used in finding event induced changes. The expressions below show us the effects of the estimations:

$$R_{k,t} = p_k + \zeta_k R_{k,t-1} + \delta_k D_{k,t} + \varepsilon_{k,t} \quad (5)$$
$$h_{k,t} = a_{k,0} + a_{k,1}\varepsilon_{k,t-1}^2 + \lambda_k h_{k,t-1}$$
$$\text{With } h_{k,t} = E_{t-1}(\varepsilon_{k,t})$$

And to estimate the effect on conditional volatility:

$$R_{k,t} = p_k + \zeta_k R_{k,t-1} + \varepsilon_{k,t} \quad (6)$$
$$h_{k,t} = a_{k,0} + a_{k,1}\varepsilon_{k,t-1}^2 + \lambda_k h_{k,t-1} + \delta_k D_{k,t}$$
$$\text{With } h_{k,t} = E_{t-1}(\varepsilon_{k,t})$$

There will also be an attempt to estimate the following expression:

$$R_{k,t} = p_k + \zeta_k R_{k,t-1} + \delta_k D_{k,t} + \varepsilon_{k,t} \quad (7)$$
$$h_{k,t} = a_{k,0} + a_{k,1}\varepsilon^2_{k,t-1} + \lambda_k h_{k,t-1} + \delta_k D_{k,t}$$
$$\text{With } h_{k,t} = E_{t-1}(\varepsilon_{k,t})$$

However, since this estimation has a very large number of parameters to estimate, the amount of observations for some currencies may not suffice. This will be examined in the results section. The intuition is described above under the expressions used in event-based data. The difference here is the dependent variables and the indicator for different sets. Instead of event i, the different runs are for the 7 different currencies k. Here again, the first lags of the dependent variables are regressed and estimated by their respective parameters. Estimating these parameters will provide insight in the effect of terrorist threat on the returns, abnormal returns and cumulative abnormal returns**.**

As for the third hypothesis, this provides us with additional interesting tests. Recall that the aim is to remove Northern Ireland from the full set and test these regions and the full set with this region omitted separately to create insight in the difference. In this type of data, all relevant observations are in the GBP. Therefore, regression (7) will be run on all subgroups within the United Kingdom. This means, to be clear, the United Kingdom with and without Northern Ireland and Northern Ireland by itself. This will be outlined in the results section.

5. Results

The next section is dedicated to showcasing the results that were produced by the methodology described earlier. All regressions were run, and the most notable results are outlined in the next few pages. These all serve the purpose of answering our three hypotheses: firstly, does terrorism affect the level of returns of local exchange rates, secondly does terrorism affect the volatility of returns of local exchange rates and finally does these results differ for regions that more regularly experience violence. This section is divided by the different types of data. Firstly, the methodology that handles the data event by event will be examined. Next, the data will be regarded that is treated on a currency level. This section is dedicated to regarding all results in interesting for answering our hypotheses. The conclusion section is devoted to answering the hypotheses.

Event-based data

Recall that the event-based data consists of 507 separate events. Every event data is paired with the exchange rates and calculated returns of the 250 days before and after the event, leaving us with 500 observations per event. This data is processed by running two different kinds of GARCH specifications. That means that 507 separate events will be run twice, once by equation (1) and once by equation (2). The event-by-event approach is finalized with the test statistic that was mentioned in the methodology as well. The test-statistic will allow us to collapse the 507 distinct events to one interpretable statistic. This part will also include a discussion of the observations per GARCH estimate to give us extra sense in the sign, magnitude and significance of the coefficients.

Here we will discuss the findings after the estimation of the parameters in equation (1). As can be found in the methodology section, the first GARCH model had a terrorism term in the OLS leg of the GARCH model. This means that it tests for an effect in the returns of the currency by testing the dummy variable for significance. Table 1 displays a summary of the estimated coefficients of the dummy variable for the events. Note that here, the number of events decreased from 507 to 432. This is due to the large number of parameters that had to be estimated

Table 1

GARCH 1	Negative		Positive		Total	
	Count	%	Count	%	Count	%
Significant	55	12.7%	15	3.5%	94	16.2%
Insignificant	176	40.7%	186	43.1%	338	83.8%
Total	231	53.4%	201	46.6%	432	100%

Table 1: Summary of the observations after GARCH estimation (1) in the methodology section. The value of interest to our hypotheses is the fitted value of the coefficient to the dummy that takes on a value of 1 after a terrorist event. All 432 coefficients, their signs and significance are summarized in this table.

in the equation. This caused some of the GARCH-estimations to not be feasible. Events which could not be estimated were dropped from the data. What stands out is that the coefficients are slightly more negative than positive. The only indicator that this table provides on the magnitude of the coefficients. In this case, we see that negative coefficients tend to be a lot more significant than positive coefficients. For extra insight, all individual coefficients are graphically displayed in figure 1.

This gives us a hint to the direction we could expect the t-statistic to tend to, but in order to start drawing any conclusions we need to run the t-statistic itself.

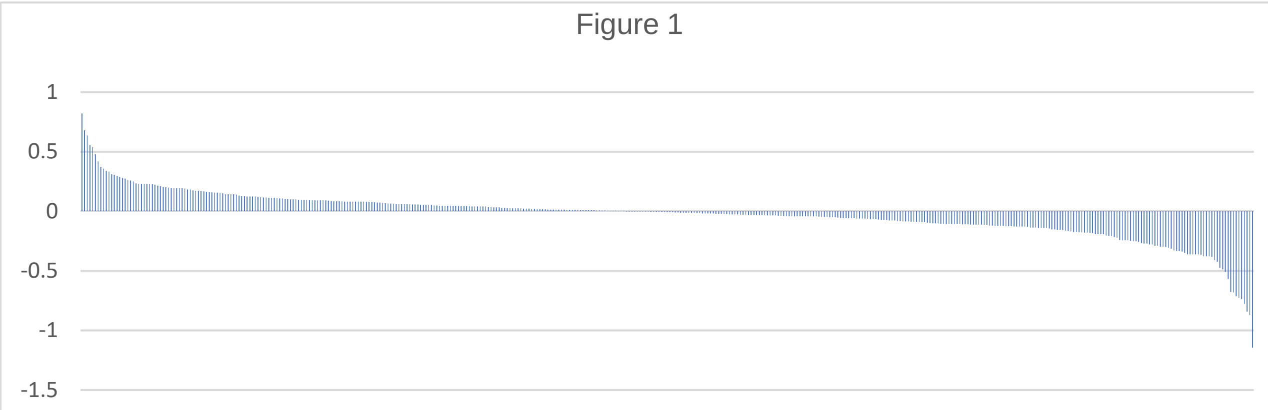

Figure 1: The 432 estimated coefficients to the dummy variable for terrorism in equation (1). This graph provides a clearer picture to the output of the equation. Every observation on the x-axis is a different coefficient to a GARCH estimation. This makes for 432 observations on the x-axis. These output coefficients are used as input for test statistics (3) and (4)

Table 2 summarizes the outcome of the GARCH estimated as calculated by the test statistic. Firstly, there in some information on the fit of this model to the question we want to answer. We included $Return_{t-1}$ to evaluate its inclusion in the regression. The value of the return, once lagged to t-1, appeared to be an excellent predictor of the value at t. It is highly significant and slightly positive, leaving us to conclude that periods of higher return may tend to be subsequent. As for

Table 2

N= 432	Mean	Standard deviation	DF	T-value	P-value α=0.05
$Return_{t-1}$	0.041651561	0.009808	432	4.246778	0.000013***
$Dummy$	-0.027636727	0.009759	432	-2.8319	0.002422***

Table 2: Test-statistics (3) and (4) calculated by the first GARCH approximation of equation (1) where the terrorism dummy is included in the regression term of the GARCH approximation. Mean is the nominator of the test statistic that is used to calculate the T value, and standard deviation is the denominator.

the appropriateness of the volatility modeling, GARCH and ARCH effect were present and highly significant in 99% of observations as was expected in using financial data. The constant was not calculated via the t-statistic since it will not be interpreted for this purpose.

Secondly, we attempt to interpret the result in table 2 for the dummy variable. We see that on days that experience terrorism or the 14 days after, tend to experience a slight decrease in the return of the exchange rate. This result is significant at the 1% level and very interesting to answering our first hypothesis. It is worth noting, however, that the approach of testing by t-statistic is sensible to outliers. There is also an inherent risk in our data that. Since we use different currencies, currencies that tend to be more volatile, and may be more susceptible to terrorist activity, may have a larger effect on the t-statistic than relatively stable currencies. We need to keep this in mind when drawing conclusions. Another risk is the risk of manual errors. Since all the individual regressions need to be copied and pasted by hand, and we rely on the functions in software to gather the right coefficients, there always looms a danger of error that is larger than when the researcher would readily interpret a single regression. However, much effort was taken to eliminate the possibility of errors in the calculation of the statistics.

Table 3

GARCH 2	Negative		Positive		Total	
	Count	%	Count	%	Count	%
Significant	93	30,2%	58	19,0%	151	49,2%
Insignificant	24	7,8%	132	43,0%	166	50,8%
Total	117	38,0%	190	62,0%	307	100%

Table 3: Summary of the observations in the second GARCH estimation by means of equation (2). Every observation is a coefficient as estimated by the GARCH.

The same approach was applied to the second GARCH. Recall that here the dummy variable was estimated in the conditional volatility leg. In table 3, the observations from these estimations are summarized. Notice there that, again, some observations needed to be dropped because of the problems in estimating the parameters, as encountered before.

First, we notice that 62% of the coefficients is positive. In this case the interpretation of a positive coefficient is that volatility tends to be higher in times of terrorist activity. The interpretation of a negative coefficient is exactly the other way around. This would be in line with the expectation of higher volatility in times of fear. If, however, we look at the significance of the observations, the negative observations are significant a lot more often. No conclusions can be drawn based on table three, but it does hint us in the direction that the dummy for terrorism may lead to some exceptional movements in the exchange rates.

For the second GARCH, the t-statistic is calculated in table 4. As was the case in the OLS dummy GARCH, $Return_{t-1}$ appears to be a good predictor of Return. Also, ARCH and GARCH effects were present in all the individual regressions. The results in table 4 are interesting towards answering the second hypothesis that was posed. It appears to be, interestingly, the case that the variance of the return tends to be lower on days of terrorist activity. This would mean that volatility of the exchange rate returns tends to be lower in the fact of terrorism. This surely is an interesting result and is not in line with the expectation that fear, and uncertainty may lead to increased volatility. The coefficient is highly significant at a level of 1%. However, this regression suffers the same risks to internal validity as the earlier regressions do. Namely, the increased risk of manual errors in gathering the data and the biases that may exist in the dataset. These downfalls will be incorporated when drawing conclusions.

Table 4

N= 307	Mean	Standard deviation	DF	T-value	P-value α=0.05
$Return_{t-1}$	5.68374E-10	2.38406E-5	306	17.20771	0.008***
$Dummy$	-7.37688	3.023378102	306	-2.43364541	0.002422***

Table 4: Test-statistics (3) and (4) calculated by the second GARCH approximation of equation (2), where the terrorism dummy is included in the variance term of the GARCH approximation. Mean is the nominator of the test statistic that is used to calculate the T-value, and standard deviation is the denominator. Notice that the information in this table if different from table 1 because in this regression, the dummy is in the conditional variance term of the GARCH.

Ideally, one would wish to include dummy variables in both the regression term and the conditional volatility at the same time. This proved, however, unfeasible with the experiment setup. In attempting to run the GARCH models, two problems arose. The first problem is that the limited number of observations per event caused the number of parameters to be too high to estimate the coefficients. The second problem was the fact that the dummy variable has a value of 0 in most of the data and a very short window where it is not equal to 0, making it almost a constant and very hard for the GARCH to estimate its effect. This means that, sadly, in the event-based section we cannot include these regressions

Currency-based data

The second division of our data that may help us gain insight into our hypotheses is the currency-based data. Recall that in this set-up, the full set of dates serves as the time variable throughout the time series. The events that occur cause the dummy variable to take on a value of 1. Contrary to the event-based run of the data, this approach evidently allows for multiple events per currency. The methodology that was described before will be applied to the currencies. Our data consists of 8 different currencies. These are summarized below in Table 5. It is worth noting that there are very large differences in the sample sizes per currency. When concluding on any visible effects in the results, this shortcoming to the data needs to be acknowledged and incorporated.

The first investigation is a simple OLS regression without any volatility clustering elements. Such a regression can highlight any early signs of tendencies in the data. These provided no interesting results and can be found under appendix A.

Next, the aim is to run the GARCH regression with both dummy variables included in both the OLS term and the conditional volatility. For all currencies outside the United States and the United Kingdom there were significantly more issues. The dummy variables take on a value of 0 in such a large percentage of the observations that the variable becomes highly multicollinear. This means that due to linear dependency, the matrix that is used in the calculation of the fitted parameters, can no longer be inverted. Since one of the terms in the OLS calculation requires the inverted form of the matrix, the whole calculation is impossible. This limits the currencies that can be used in the investigation to two. The fact that only two currencies can be used hinders the external validity of this research. However, there are still interesting results for these nations.

Table 5

Country	Currency	Number of terrorist events	Timeframe
United states	USD	81	04-01-1971 until 31-12-2017
Australia	AUD	5	04-01-1971 until 31-12-2017
Canada	CAD	6	04-01-1971 until 31-12-2017
Europe*	EUR	41	04-01-1999 until 31-12-2017
United Kingdom**	GBP	365	04-01-1971 until 31-12-2017
Norway	NOK	2	04-01-1971 until 31-12-2017
Sweden	SEK	2	04-01-1971 until 31-12-2017
Switzerland	CHF	4	04-01-1971 until 31-12-2017

Table 5: For added clarity: a summary of the countries and currency unions that make up our data. * Note that Europe is comprised of all European nations that experienced terrorist violence in the observation period which are specified in the data section. ** United Kingdom is comprised of all observations in their four constitute countries: England, Scotland, Northern Ireland and Wales

In table 6 we first look at the dummy variables in the OLS term. GARCH and ARCH effects, again, were present in both the first dummy, as well as the second one discussed later. Here we see, contrary to the results under the event-based test, that the dummy variable in this run of our data, is insignificant in explaining the return of the local currency. This surprising result may be due to the problems with the test-statistic mentioned earlier. It may also be because the observations that were left out in earlier stages had a significant influence on the calculation of the test-statistic. In both cases, the fact that this second run has a more traditional and well-proven approach. Therefore, we prefer these results in our conclusion.

Table 6

δ_1	Coefficient	Standard deviation	T-value	P-value
GBP	-.0069621	.0088017	-0.79	0.429
USD	-0.002008	0.016264	-0.12	0.902
δ_2	Coefficient	Standard deviation	T-value	P-value
GBP	-1.528203	.1119424	-13.65	0.000***
USD	-0.513904	0.067550	-7.61	0.000***

Table 6: The results of the two-dummy GARCH estimations of equations (5) and (6) for the United Kingdom and the United states

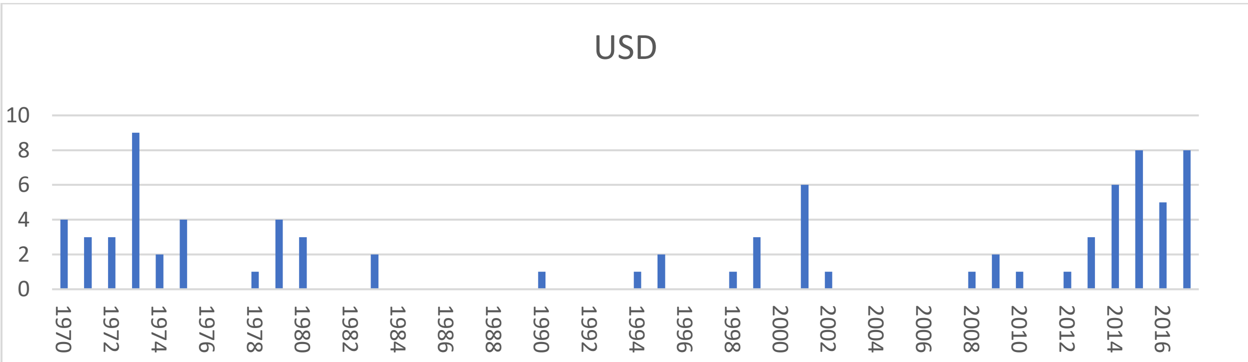

Figure 2: The timeline of the attacks in the United states

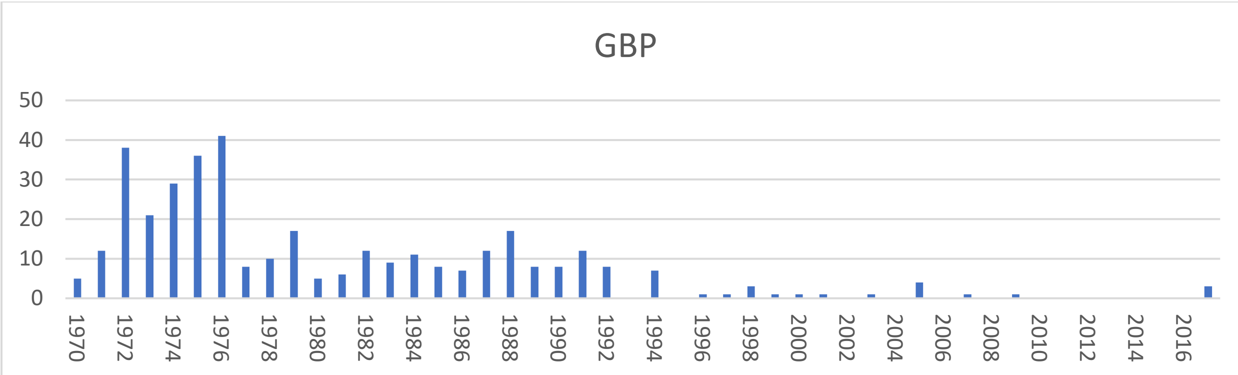

Figure 3: The timeline of the attacks in the United Kingdom

As for the second dummy, we see that both currencies show a negative estimated coefficient and a highly significant p-value This is in accordance with the event-based approach that we calculated earlier. The interpretation here, again, is that in times of terrorist activity, volatility tends to be lower in times of terrorist activity. This result is, again, counterintuitive. Why would terrorism correlate with significant periods of lower volatility? Since we have no hypothesis for such an assertion, we again look for another explanation. One possible explanation could be that terrorist events are clustered themselves. If the homoscedasticity is not fulfilled throughout the data, that means that different periods in the time-series experience different levels of volatility. If clustered periods of terrorist events tend to fall in periods of lower volatility, then it is only natural that there are significant effects. If we look at figure 2 and 3, we see a timeline and the number of attacks in both countries that were able to estimate all parameters. In the United States data, we see, indeed, that the data tend to have some consequent years that see more attacks, and long periods that do not see any attack whatsoever. In the case of the UK, we do see clustering over a very long period. This will be elaborated upon below. As for now, the only conclusion we can draw is that there appears to be a tendency to have lower volatility of exchange rates in the time periods that we investigated. It does appear that this is related to terrorism, but there might also be an omitted variable that we have not incorporated in this paper. As a suggestion for further research, this is a very interesting topic to be further investigated.

The United Kingdom and Northern Ireland, a comparison

The last part of this results section aims to answer the third hypothesis, which aims to find out if there is any change to the effect if a region experiences violence more regularly. This hypothesis can be answered by taking a country that experiences regular attacks and looking how the attacks in regions that do, and do not experience regular attacks differ in their effect on the returns.

The country that was chosen to investigate this question, as explained earlier, is the United Kingdom. As can be seen in Table 5, the United Kingdom has the most observation. Also, it has a region known for violence. This region is Northern Ireland. If we look at figure 2, which we examined in the previous part of this results section, we see basically two very different periods. One before 1998, where there were a lot of attacks in Northern Ireland, and the period thereafter. The year 1998 marked the good Friday agreement between the United Kingdom and the Republic of Ireland. This ended the hard border between Northern Ireland and the Republic of Ireland and has recently been a subject of much conversation in context of the Brexit negotiations. As outlined in the data and methodology, the Northern Irish and non-Northern Irish attacks were separated, and their results summarized in table 7.

Table 7

δ_1	N	Coefficient	Standard deviation	T-value	P-value
Northern Ireland Included	365	-0.0071134	0.0078929	-0.9	0.367
Northern Ireland excluded	33	-0.015397	0.0290413	-0.53	0.596
Solely Northern Ireland	332	-0.0041933	0.007888	0.75	0.451
δ_2	N	Coefficient	Standard deviation	T-value	P-value
Northern Ireland Included	365	-1.508753	0.0972819	-15.51	0.000***
Northern Ireland excluded	33	-0.2137096	0.2999539	-0.71	0.476
Solely Northern Ireland	332	-1.549259	0.101133	-15.32	0.000***

Table 7**:** Comparing the effect of the regressions in the three country level subgroups for the United Kingdom. The models estimated were the same as before with δ_1 estimating the effect of terrorism on returns and δ_2 showing the effect of terrorism on volatility. Estimations of equations (5) and (6)

There are some interesting aspects to this table. Firstly, we see that for all subgroups the tendencies, in terms of sign, that we have found before are the same. If we at the number of observations, we notice that only a very small percentage of observations occurred outside of Northern Ireland. The fact that we do see effects for the second dummy variable in Northern Ireland, but not for the rest of the United Kingdom, may be due solely to the smaller number of events. Even if these few events would be highly significant, they still might not have been able to produce a collectively significant result. This leaves the data we have in order to answer the third hypothesis incomplete. The only thing we can say is that in our sample, Northern Ireland was significantly negative in its effect of terrorism on volatility, whereas this effect was not present in the other parts of the United Kingdom. Another region that has a more balanced set of event dates there may be possibilities to further investigate this question, and this is suggested as further research. However, in the set that we have chosen to use, the results are inconclusive.

Internal validity

There needs to be one final note on the interpretation of the results under all three sets of events. In order to draw this results in the second part of this result section, an assumption of normality had to be made to retrieve the results. Testing for normality, however, finds that the data is not normally distributed. Therefore, these results are not robust to non-normality. However, rejecting the assumption of normality is outside the scope of this paper. We leave this as a suggestion for further research.

6. Concluding remarks

With recently heightened tensions in terms of transnational terrorism in OECD countries, we established a need for better understanding of economic consequences of terrorist attacks. The aim was to answer this call by executing research to find links between terrorist attacks and movements in the local exchange rates. To this end we formulated three different hypotheses. The first aimed to establish a link between the returns of the local currency and the occurrence of a terrorist event, the second sought a link with the volatility of the local currency and the third one aimed to discover if the effects are similar for countries that experience regular episodes of violence. In order to test these, we used a total of 507 observations from 15 different OECD nations with a total of 8 different currencies. The full dataset was treated event by event and per currency.

In order to test these, we used an ARMA-GARCH model that allowed for special features commonly noted in financial data. The models were fit and adjusted to be applicable to all hypotheses.

The first hypothesis finds that through the use of the collective t-statistic, there is significant evidence of an effect of a decrease to the returns in times of terrorist activity. However, in the currency-based run, this effect can be seen in neither of the tested currencies. The results leave too much to be desired to draw a definitive conclusion. We therefore say that we cannot reject out initial hypothesis.

As for the second hypothesis, existing literature led us to believe that there could be visible signs of increased volatility in times of terrorist uncertainty. In both the event-based and currency-based data alignment, we found highly significant, negative effects to volatility in times during the occurrence of a terrorist event. Or in other words, significantly lower volatility in the 14 days during and after an attack. This result is surprising and not in line with the expectations. Some possible explanations have been included in the paper, but as a suggestion to further research, this link may potentially be very interesting.

The third hypothesis could not be disproven. There were very few differences between regularly attacked regions and regions that were rarely attacked. The only difference that was found was most likely due to a lower number of observations. To conclude, we found no evidence to support a rejection of the hypothesis.

Suggestions for further research were coined throughout this paper, but the most important one remains the further extension of volatility modeling in exchange rates. It appears, both from the data, as well as from this research, that a lot of interesting results remain in this variable of interest. As a final remark, a deeper understanding and more detailed application of panel data in this context may also yield interesting results.

References

Abadie, A., & Gardeazabal, J. (1999). *The Economic Costs of Conflict:.*

Balaban, E., & Constantinou, C. (2006). Volatility clustering and event-induced volatility: Evidence from UK mergers and Acquisistions. *The European Journal of Finance*.

Benmelech, E., Berrebi, C., & Klor, E. F. (2010). The Economic Cost of Harboring Terrorism. *Journal of Conflict Resolution*.

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*.

Brooks, C. (2014). *Introductory Econometrics for Finance.*

Chesney, M., Reshetar, G., & Karaman, M. (2011). The impact of terrorism on financial markets: An empirical study. *Journal of Banking & Finance*.

Copeland, L. (2014). *Exchange rates and international finance.*

Edward, S., & Yeyati, E. L. (2004). Flexible exchange rates as shock absorbers. *European economic review*.

Eldor, R., & Melnick, R. (2004). Financial markets and terrorism. *European Journal of Political Economy*.

Engle, R. F. (1982). Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United. *Econometrica*.

Flandreau, M., & Eichengreen, B. (2009). The rise and fall of the dollar (or when did the dollar replace sterling as the leading reserve currency?). *European Review of Economic History*.

Frenkel, J. A. (1976). Monetary Approach to the Exchange Rate: Doctrinal Aspects and Empirical Evidence. *Journal of economics*.

Giaccotto, C., & Ali, M. M. (1982). Optimum Distribution-Free Tests and Further Evidence of Heteroscedasticity in the Market Model. *The Journal of Finance*.

Gupta, D., & Venieris, Y. (1986). Income Distribution and Sociopolitical Instability as Determinants of Savings: A CrossSectional Model. *Journal of Political Economy*.

Hopper, G. (1997). *What Determines the Exchange Rate: Economic Factors or Market sentiment?* Federal Reserve Bank of Philadelphia.

Kanas, A. (2003). Volatility Spillovers Between Stock Returns and Exchange Rate Changes: International Evidence. *Journal of Business, Finance & Accounting*.

Meade, J. E. (1951). The removal of trade barriers: a regional versus universal approach. *Economica*.

Melvin, M., & Tan, K.-H. (1996). Foreign exchange market bid-ask spreads and the market price of social unrest.

OECD. (2002). *Economic consequences of terrorism.*

Peren Arin, K., Ciferri, D., & Spagnolo, N. (2008). The price of terror: The effects of terrorism on stock market returns and volatility. *Economics Letters*.

Richardson, C. (2011). *Relative Deprivation Theory in Terrorism: A study of higher education and unemployment as predictors of terrorism.* New York: New York University.

Sarno, L., & Schmeling, M. (2011). Which Fundamentals drive Exchange Rates? A cross-sectional perspective. *Journal of Money, Credit and Banking*.

Shahbaz, M. (2013). Linkages between inflation, economic growth and terrorism in Pakistan. *Economic Modelling*.

Smith, C. E. (1992). Stock markets and the exchange rate: A multi-country approach. *Journal of Macroeconomics*.

University of Maryland. (1970-2018). *Global Terrorism Database.* Retrieved from https://www.start.umd.edu/gtd/using-gtd/

Appendix A

Table 8

	Variable	Coefficient	Standard deviation	T-value	P-value
GBP	$Return_{t-1}$	0.040827	0.0089375	4.57	0.000***
	$dummy$	-0.0068235	.0122437	0.56	0.577
USD	$Return_{t-1}$	0.0156517	0.0089083	1.76	0.079*
	$dummy$	0.04145	0.0256568	1.62	0.106
EUR	$Return_{t-1}$	0.0237430	0.0346985	0.68	0.494
	$dummy$	0.0074109	0.013755	0.54	0.59

Table 8: Running a basic OLS regression on all different currencies with Returns as a dependent variable and yesterday's returns and the terrorism as explanatory variables.