Saving for a rainy day? Estimating the economic impact of the 2021 Limburg flood.

S. van der Heijden

Student no: 4297741

September 11, 2022

Abstract

In this thesis I use a novel approach to estimate the economic impact of the 2021 Limburg flood using high-frequency transaction data from ABN AMRO bank. High-frequency bank transaction data have previously been proven valuable in accessing the economic impact of the COVID-19 pandemic (Neuteboom et al., 2021). To the best of my knowledge, this type of data has not been used to estimate the economic impacts of a natural catastrophe. I focus on a synthetic difference-in-differences methodology to estimate the impact. I find that the economic impact is 18,045 EUR on average of extra spending per inundated individual in Valkenburg aan de Geul, the most heavily hit area. This is very similar to the damage estimate of the ENW (2021) of 18,713 EUR on average per person for the inundated in Valkenburg aan de Geul. Furthermore, the duration of the economic impact for the inundated is roughly 35 weeks on average. Finally, I did not find a measurable economic impact to uninundated and evacuated individuals. In summary, high-frequency bank transaction data paired with a synthetic difference-in-differences model is a reliable gauge of the economic impact of a flood and should be used to estimate the economic impact of future natural disasters. Additionally, it can be the empirical foundation for calibrating existing damage models.

Supervised by:

Dr. T. van den Bremer, Prof. S.N. Jonkman, Dr. S. Phlippen and Dr. A. Varela Varela

Contents

1 Introduction

1.1 Context

North-west Europe was hit by heavy rainfall and floods in July of 2021 causing extensive damage and loss of life. In particular, the province Limburg in the south-east of The Netherlands, several provinces in west Germany and the north-east of Belgium were hit. Overall losses are estimated to be approximately 46 billion EUR and insured losses are estimated to be approximately 11 billion EUR[1], ranking this event second and third globally in 2021 for overall and insured losses, respectively (MunichRE, 2021). Kreienkamp et al. (2021) find that heavy rainfall was caused by the low-pressure system 'Brend' which is expected to occur once every 400 years in the current climate. Additionally, they find that historical climate change potentially contributed significantly to the likelihood of this event. ENW (2021) find precipitation was predictable between Jul 10 and Jul 11 for the Netherlands, although peak discharge was systematically underestimated. On the 12th, the Veiligheidsregio warned for high water in brooks and the Meuse. Valkenburg was flooded on the 14th of July. Residents from Valkenburg, Roermond and some areas by the Meuse were evacuated on the 15th. After the 17th people were able to return to their homes. Figure 1 shows an overview map of the affected area in The Netherlands.

ENW (2021) find the economic impact is estimated to be in the order of 350 to 600 million EUR, affecting more than 2,500 houses, 5,000 inhabitants and 600 businesses. These damages are estimated with HIS SSSM, first introduced by Vrisou van Eck and Kok (2001). Physical damage, business interruption and damage to infrastructure and crop losses are found to be most significant. Additionally, physical damages to residential and commercial structures are very irregular, meaning damages vary widely between buildings. Different from the floods in 1993 and 1995 which mostly affected the Meuse floodplain, the largest damage occurred in the Geul floodplain, especially in the city of Valkenburg.

The HIS SSSM method relies on assumption of the maximum damage per unit (e.g.

[1]Using a proprietary method, see Sampson et al. (2014) for more information.

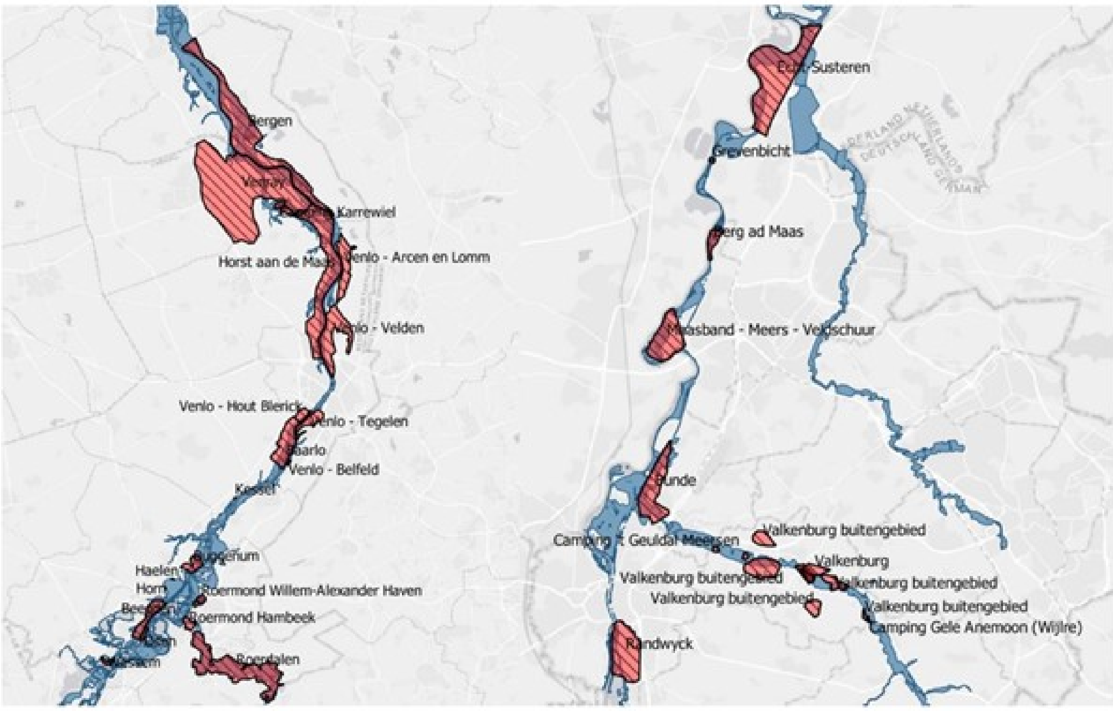

Figure 1: Map of emergency ordinances and evacuations (red) and flood contours (blue) in Limburg and Noord Brabant ENW (2021).

square meter of a house) which can influence the final damage estimate substantially. As one example, maximum damage for a house may be 1,000 EUR per square meter. However, in the aftermath of a flood there may be labour cost inflation causing the replacement value to be substantially higher (Grogan & Angelo, 2005). Additionally, macroeconomic impacts such as recessions or pandemics may inflate or deflate the replacement values.

1.2 Research questions and hypothesises

The aim of this research is to estimate the economic impact caused by the 2021 Limburg flood using geo-located high-frequency transaction data. I define the economic impact as the extra money spend by impacted individuals because of the 2021 Limburg flood. My goal is to answer the following main research questions:

Research question 1: How can high-frequency transaction data be used to estimate the duration of the 2021 Limburg flood impact to individuals?

Research question 2: How can high-frequency transaction data serve as a proxy for the economic impact to individuals caused by the 2021 Limburg flood?

Research question 3: What is the economic impact and duration of the impact caused by the 2021 Limburg flood to individuals using high-frequency transaction data?

I use the following hypothesis to support me in answering the aforementioned research questions:

Hypothesis 1: High-frequency transaction data allow me to compare between flood-affected individuals and similar but unaffected individuals to measure the economic impact caused by the flood in a statistically sound manner.

Hypothesis 2: There is a measurable effect to individuals in transaction data following the 2021 Limburg flood.

Hypothesis 3: The effect is shorter than the length (in time) of the data. Therefore, I can measure the duration of the effect.

1.3 Literature summary

Flood events can lead to a large spectrum of consequences. Jonkman et al. (2008) give an overview of flood damages that are commonplace in the literature such as clean-up costs and damage to residences and vehicles. The economic impact is assumed to be a subset of these flood damage dimensions. Many of the dimensions may be measurable using high-frequency bank transaction data since they are paid for by consumers using bank transfers and cards at point-of-sale (POS) locations (e.g. PIN machine at grocery store). Additionally, the ENW (2021) estimate damages for individuals at the Geul area of the 2021 Limburg flood to be 54-68 million EUR. They use HIS SSSM to estimate these damages which is based on a bottom-up approach that uses flooded surface area, propose of use and hydraulic conditions like inundation depth (Slager & Wagenaar, 2017). Additionally, a damage factor is determined for each building that maps the hydraulic condition onto a factor that measures the percentage of damage done to a particular building. Moreover, total damage is calculated by taking the product of the damage factor and maximum building damage and taking the sum over all affected buildings in the flooded area. Other models exist to estimate the impact of floods and other natural disasters. All of these have limited empirical basis and rely on many assumptions, such as the shape of the damage functions Sampson et al. (2014).

High-frequency transaction data has been used before in different settings. For example, Gelman et al. (2014) and Sobolevsky et al. (2017) use high-frequency transaction data to test consumption smoothing theories and predict socioeconomic indices, respectively. Additionally, Neuteboom et al. (2021) use similar data to infer the impact of the COVID-19 pandemic to consumer spending. To the best of my knowledge, this type of data has not been used to estimate the economic impact of natural disasters.

1.4 Methodology

I investigate a novel approach which combines the fields of Econometrics and Hydraulic Engineering, to evaluate flood damage dimensions using high-frequency transaction and account records data of ABN AMRO bank (AAB) clients. Even though flood damage dimensions cannot be precisely measured using this approach, I may be able to estimate the economic impact using the data as a proxy. A proxy is a variable that is not immediately relevant on its own, but can be used as an indirect measure, when I assume that the observed variable is strongly correlated with the variable of interest. A well-known example is gross domestic product (GDP), the net value of all goods and services created in an economy, which is commonly used as a proxy for well-being. In this case, GDP is the observed variable and well being is the variable of interest.

Given that I am interested in estimating the economic impact flood impacted individuals, I require a methodology that can separate the effects of a flood from all other effects in the transaction data. In particular, I need a method to compare the spending of the inundated residents and inundated and evacuated residents with a control group (unobserved counterfactual) that models the behaviour of the group of impacted individuals (treatment group) as if the flood did not occur. Then, I can subtract the spending from the treatment group from the unobserved counterfactual to obtain the causal effect of the flood. To achieve this several methods are commonly used in literature. Two of the most commonly used methods are difference-in-differences (a.k.a. a fixed-effects panel model) and synthetic control. Both methods rely on strong assumptions that easily invalidate the models. When applying difference-in-differences we have to make a parallel trend assumption. In simple terms, this means that the spending of the treatment group is equal to the spending of the unobserved counterfactual plus a level shift. When applying synthetic control I assume that I can match the spending of the unobserved counterfactual to the treatment cohort using a weighting function (Abadie & Gardeazabal, 2003). Thus I assume that I can find a group that behaves like the impacted individuals without being flooded. Arkhangelsky et al. (2021) propose a combination of the aforementioned methods called synthetic difference-in-differences (SDID) that alleviate these issues to some extend while keeping the benefits of both models.

1.5 Data

To estimate the economic impact of the 2021 Limburg flood, I use flood and evacuation area coordinates from Slager and Wagenaar (2017). These data contain areas of inundation, evacuations, blackouts and sludge blockages in the Amersfoort coordinate system. Additionally, I obtain obtain client information data from ABB. These data contain information such as account numbers, types of clients (individuals/business), names, birthdates, and residential zip-codes. Moreover, I use the Basisregristratie Adressen en Gebouwen (BAG) to infer the residential location of AAB clients in the Amersfoort coordinate system. Furthermore, I obtain high-frequency Single Euro Payments Area (SEPA) and Point-of-sale (POS) transaction data from AAB from Jan 1, 2018 to Apr 3, 2022. SEPA is a system that facilitates cashless transactions in Europa. These flows include items such as online transactions, rent, mortgage & insurance payments, and salaries, but exclude POS transactions. POS transactions include all purchases made using debit cards. Each observation contains a transacted amount and meta data such as a time stamp, account number and textual description.

2 Literature

In the following section I discuss relevant literature which helps me answer the research questions from section 1.2. First, different flood damage dimensions are discussed. Second, commonly used damage estimation methods are discussed. Third, the results from ENW (2021) is summarised. Fourth, an updated version of HIS SSSM, Standaardmethode 2017 Schade en Slachtoffers, a bottom-up method for flood damage estimation, by Slager and Wagenaar (2017) is discussed. Finally, an overview of economic impact estimation using surveys and transaction data is given.

2.1 Flood damage dimensions

In this section I give a general overview of flood damage dimensions. Flood events can lead to a large spectrum of consequences ranging from tangible damages to intangible damages. Jonkman et al. (2008) give an overview of different dimensions of flood damage that are commonplace in the literature. For example, if a residential area floods, tangible damage to homes or infrastructure is hard to avoid. Additionally, floods can lead to intangible psychological damages which are difficult to value. Table 1 shows different dimensions of flood damages, inspired by Jonkman et al. (2008), that I measure using high-frequency bank transaction data. This is because most payments in the Netherlands are done digitally (DNB, 2021), which means they show up in transaction data records. These flood damage dimensions are directly paid for by the impacted individuals. Therefore, transaction data includes the flood damage dimensions from Table 1.

Table 1: Different dimensions of flood damage inspired by Jonkman et al. (2008) that may be measurable with high-frequency bank transaction data.

Flood damage dimension	Example
Residences	Repair costs of flooring, windows, doors, etc.
Durable goods	Replacement costs of washing machines, kitchen appliences, bycecles, books, etc.
Non-durable goods	Replacement costs of food, bevarages, clothing, shoes, etc.
Vehicles	Repair or replacement costs of cars or motorcycles
Clean up costs	Labour costs for professional cleaners or do-it-yourself items such as drying machinery
Adjustment in consumption patterns	Change in money spend at grocery stores because of increase/decrease in home cooking
Temporary housing of evacuees	Hotel costs that are made during repair/dry-up period

2.2 Damage models

There exist many different flood damage estimation methodologies. In this section I provide a brief overview of some commonly used flood damage estimation methods: proprietary reinsurance catastrophe models, vendor catastrophe models, the simple stage-damage curve model called Multi-Coloured Manual and a spacial scale model HIS SSSM which was used to estimate damages in Limburg.

A large part of damages caused by natural disasters is insured. For example, in 2021, global damages caused by natural disasters were $280 billion of which $120 billion were insured (MunichRe, 2022). A large portion of this risk is not retained by local insurance companies but ceded to reinsurance companies such as Munich Reinsurance Company and Swiss Re Ltd, as these risks are difficult to diversify for local insurance companies (Brahin, 2015). To value damages, reinsurance companies use propriety models. Unfortunately this means I do not know how they model damages. However, damage estimates between reinsurance companies vary substantially. For example, in 2012 Munich Reinsurance Company and Swiss Re Ltd. estimated damages of floods in Australia to be valued at \$2.8 and \$6.1 billion, respectively (Sampson et al., 2014).

Not every reinsurance company has the capability to develop propriety catastrophe models. To fill this void, vendor models were created, which are proprietary in nature also. The most commonly used models are created by AIR Worldwide and Risk Management Solutions (Bermuda:Re+ILS, 2021). Sampson et al. (2014) discuss the modules proprietary catastrophe models should contain. First, a stochastic module, that is capable of generating a database of plausible events. Second, a hazard module, used to simulate a series of events. For example, it should be able to simulate a map of water depths in case of a flood. Third, a vulnerability module, which calculates expected damages, similar to damage function, as a function of the water depths in the stochastic module. These models are developed for a global scale, which creates uncertainty in the damage functions because these are influenced by a large number of local factors (Sampson et al., 2014). For example, building type, construction method and precautionary measures. The previously discussed proprietary reinsurance models are likely of similar composition to the vendor models (Sampson et al., 2014).

The Multi-Coloured Manual (MCM) is a step-by-step guide to estimates benefits of flood risk management for the United Kingdom (MCM, 2022). MCM contains methods for estimating damages to property, vehicle damage, evacuation costs and more. It returns relatively simple damage estimates as a function of water depth (damage function) for different objects. On the one hand, the MCM is relatively easy to use and it applies to 75% to 80% of cases (MCM, 2022). However, it relies on damage functions assumptions. One can argue this is less problematic than in vendor models, since MCM is only applicable to the UK.

HIS SSSM a method for estimating expected damage and victims as a result of flooding (Vrisou van Eck & Kok, 2001). It is a more sophisticated method that MCM as it uses spacial data files to infer damages across large regions. Similar to MCM it relies on damage function assumptions focused on the Netherlands. An estimate of the residential damages for the Geul area and a more in-depth explanation of the model is given in the following sections.

2.3 Flood assessment

ENW (2021) discuss the cause and effect of the 2021 Limburg flood. They investigate different impact categories. In particular, hydrological, civil-technical, economic, social, crisis response and health. ENW assess the damage caused by the 2021 Limburg flood using HIS SSSM. They estimate total tangible flood damage to be in the range of 200-250 million EUR for the Geul area. Additionally, roughly 27% of this damage, or 54-68 million EUR is attributed to residents. Moreover, roughly 3,150-3,411 adults were affected in the area (see subsection 4.3). This equates to on-average 15,830-21,596 EUR per affected individual. Additionally, there is a substantial variability between individual residential damages estimates. This estimates contains some of the damages in Table 1. However, HIS SSSM does not include clean up costs, adjustment in consumption patterns and temporary housing of evacuees. In the following section, we will briefly discuss HIS SSSM.

2.4 Standard flood damage method

Slager and Wagenaar (2017) propose an update to the methodology of HIS SSSM to estimate flood damages that is based on a bottom-up approach that uses the flooded surface area, purpose of use and hydraulic conditions such as inundation depth. Below I briefly discuss this method.

HIS SSSM estimates damages by taking the sum of all affected building or units that is assigned a damage value. To obtain this sum several steps are taken. First, the severity of the flood is investigated and the replacement value of each object is determined. Particularly, the water depth at each xy-coordinate in the inundation area should be determined using software such as SOBEK. For example, the water depth is estimated at the location of each object during the flood. Second, a building register like the Basisregistratie Adressen en Gebouwen (BAG) is used to label all buildings in the inundated area (see section 4.3 for more information). Specifically, the BAG records information such as a ZIP code, xy-coordinates, surface area and purpose of use for every building in The Netherlands by law . As such, every object can be assigned a purpose-of-use, surface area and water depth. Additionally, homes can be partitioned into single-family homes and apartments. In summary, the number of units, the water depth at inundation and a maximum damage value per surface area is obtained.

Second, total flood damage (S) is estimated by,

$$S = \sum_{i=1}^{N} \alpha_i n_i s_i, \tag{1}$$

where $\alpha_i \in [0, 1]$ is the damage factor for category i, n_i is the number of units in category i, s_i is the maximum damage per unit (per m^2) for category i and N is the total number of categories. The damage factor α_i maps the hydraulic conditions onto a number that represents the percentage of maximum damage done.

For example, Figure 2 shows a sketch of a damage factor function as a function of water depth that HIS SSSM may use. Imagine an apartment building with two floors. The green line represents the ground floor and the blue line represents the first floor. Water damage will affect the ground floor first, and increasingly so when it rises. The first floor will only be affected by higher water levels. In this example, 40% of maximum damage is reached for a 1 meter water level and a 3 meter water level for the ground floor and first floor, respectively.

To explain this method further, I assume the maximum damage of a two-floor apartment building is 1,000 EUR/m^2 (Slager & Wagenaar, 2017). Additionally, I assume a regional flood with a maximum water depth of 3 m affects 20 apartment buildings with eight 60 m^2 apartments on each floor. Figure 2 shows that the $\alpha_{\text{ground floor}} = 0.9$ and $\alpha_{\text{first floor}} = 0.4$. As such, I can estimate the value of total damage to be 12.5 million EUR by,

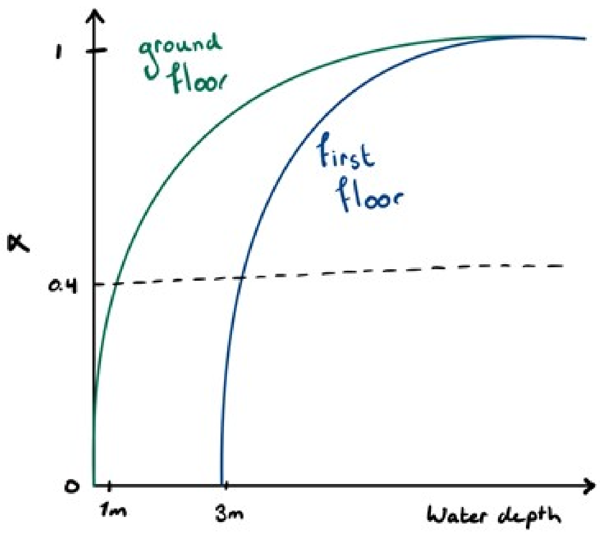

Figure 2: Schematic of water depth versus damage factor inspired by Slager and Wagenaar (2017).

$$\begin{aligned} \mathrm{S} &= \sum_{i=1}^{N} \left(\alpha_{\text{ground floor}} \cdot n_i \cdot s_i + \alpha_{\text{first floor}} \cdot n_i \cdot s_i\right), \\ &= \sum_{i=1}^{160} (0.9 \cdot 60 \cdot 1{,}000 + 0.4 \cdot 60 \cdot 1{,}000) \approx 12.5 \text{ million EUR.} \end{aligned} \tag{2}$$

The final value of the flood damage is heavily influenced by the assumptions on the maximum damage amount of a given object. This can differ substantially between different units of similar categories. In the Netherlands, maximum damage amounts are determined on the basis of macro-figures provided by the CBS, a governmental organisation charged with publishing statistical information (De Bruin et al., 2015). Similar to the example before, the maximum damage amount for a single family home is 1,000 EUR per m^2 (Slager & Wagenaar, 2017). This does not take into account heterogeneity between different units. This is especially relevant for a local case like the 2021 Limburg flood, where damage amounts may be substantially different from national averages. For example, the difference between replacing laminate and hard-wood flooring alone can run up into the 100+ EUR per m^2. Therefore, differences can be substantial if hardwood flooring is more popular in Limburg than it is on average in The Netherlands. Additionally, replacement costs of objects may increase substantially after a catastrophe due to high demand and labour shortages (Grogan & Angelo, 2005). Moreover, these estimates only include some of the damages in Table 1. For example, adjustment in consumption patterns or temporary housing is not modelled. Taken together, the underlying assumptions can significantly influence the total flood damage estimation.

2.5 Sample surveys and transaction data

Sample surveys have been at the core of measuring economic activity ever since the pioneering work by Hansen et al. (1953). This method is used to collect information from a subset of a population, to draw conclusions about the entire population. Although, sample surveys are known to be efficient tools to gather scientifically robust information if properly executed and well designed, they have several limitations. First, surveys have seen declining response rates, leading to lower quality and increased costs (Jarmin, 2019), especially for small areas like Valkenburg. Second, a bias arises in surveys because of the difference between stated preference and revealed preference. This means people act in a different way than they proclaim. For example, individuals say they spend less on goods and services after a flood. However, in reality they spend more. For example, Loomis (2011) give an overview of the difference between individuals stated willingness to pay and what they actually pay for products in an experiment setting and find significant differences. Third, outcomes are usually available with lag and at low frequencies (Chetty et al., 2020).

Banks and other financial intermediaries are collecting large datasets of individual transaction and account records as an increasing portion of economic activity is recorded digitally (Carvalho et al., 2021). This is one of the reasons that national statistical agencies agree that these data will play a more pronounced role in the 21st century (Jarmin, 2019).

To the best of my knowledge, Gelman et al. (2014) are the first to use high-frequency transaction and account balance records from a large heterogeneous sample for measuring economic activity. They use data from Check, an app that collects financial data for its users from different financial service providers, to test consumption smoothing theories. One drawback of this study is overrepresentation of males and younger adults.

Data analogous to the AAB transaction data has been used also. Sobolevsky et al. (2017) use data by Banco Bilbao Vizcaya Argentaria over the year 2011 to find a correlation between official socioeconomic indices and individual consumer spending. They claim their modelling approach performs well in predicting the socioeconomic indices. However, they do not formally test the significance of the output.

Neuteboom et al. (2021) use high-frequency geo-located transactions from AAB clients in the Netherlands to measure the consumer response to the COVID-19 pandemic. They find that consumer spending decreases because of self-isolation, not the imposed lockdowns by the Dutch government. Additionally, they find that consumers adapt their behaviour as the pandemic evolves. Consequently, latter stages of the pandemic are less damaging to the economy than earlier stages.

3 Methodology

I am interested in evaluating the effects of floods on consumption by retail and business clients of AAB using high-frequency transaction and account record data. Therefore, I require a modelling approach that allows me to evaluate differences in outcomes between clients that are exposed to a flood (e.g. residents of the city centre of Valkenburg) and similar clients that are not exposed. In particular, I want to know what would have happened to the exposed individuals if the flood did not occur. In literature, these individuals are referred to as units.

I want to compare the period preceding the flood with the period afterwards. As such, I need to model in two dimensions simultaneously, both in time and across different units. Data with observations about different units across time are also called panel data. If I manage to infer how the affected units would have behaved differently after the flood, I can establish the causal relation between the flood and changes in consumption. In particular, I can estimate how much affected units spend compared to what they otherwise would have.

One issue that arises with this approach is that I only observe one outcome for each unit. For example, if a resident of Valkenburg is affected by a flood and sequentially adjusts consumption, I do not observe the consumption of the same resident that is unaffected by the flood. In statistics this variable is called the unobserved counterfactual.

There are a few additional factors I need to consider when handling panel data. First, the share characteristics of the group that I am investigatin. Such a group of units is also called a cohort. For example, the units in a single cohort having similar characteristic such as age, income and wealth. Second, there are patterns in consumption over time across different cohorts. Specifically, the Valkenburg cohort will likely spend more in the weeks before Christmas, but so would other cohorts in the Netherlands. When I add parameters that can capture these effects, I can automatically remove those effects from the flood-induced differences (e.g. Heij et al. (2004)). In other words, it removes most of the omitted-variable bias from the model.

In practice, there are two panel-data models commonly used to work with unobserved counterfactuals and fixed effects: difference-in-differences and synthetic control (Arkhangelsky et al., 2021). Additionally, Arkhangelsky et al. (2021) propose a new method called synthetic difference-in-differences. These methods are generally used to measure the effect of a policy (Arkhangelsky et al., 2021). For example, the effect of a smoking tax on cigarette consumption. Difference-in-difference is the most parsimonious model with the smallest amount of parameters. It is generally applied when there are many cohorts (or units) that receive a treatment (e.g. the smoking tax), and when a parallel trend assumption can be made (Arkhangelsky et al., 2021). This assumption means that the cohorts are parallel in trend before the treatment. If I want to investigate the effect of the smoking tax in the Netherlands, I can use cigarette consumption of countries, where no cigarette tax is introduced, that have a similar trend in consumption before the smoking tax starts. For example, when cigarette consumption was decreasing every year before the smoking tax, this should also be the case in the other countries. Synthetic control is generally used on a smaller number of units and compensates for the lack of parallel trends by re-weighting the units. The main difficulty with this model is that the consumption before the treatment have to match exactly by applying the weights. Synthetic difference-in-differences combines the attractive features of these models while relying on fewer assumptions (Arkhangelsky et al., 2021). To the best of my knowledge, synthetic difference-in-differences has never been applied to estimate the economic impact of a flood. In the following sub-sections, I will delve deeper into the mechanics of these models.

3.1 Difference-in-differences

One of the most popular methods to measure the aggregate causal effect of a treatment on a cohort is difference-in-differences, first introduced by Snow (1855). The unobserved counterfactual is another cohort that has a parallel trend in the endogenous variable relative to the treatment cohort. The endogenous variable is the variable to be explained by the model (y). In our case the treatment is the inundation or evacuation of residents and the endogenous variable is consumption via transaction and accounts record data. To explain this method, I will use the schematic in Figure 3. Let's say the red line represents aggregate amount of transactions for the treatment cohort (e.g. the Valkenburg group). Furthermore, the blue line is a control cohort made up from a random selection of clients and I assume that the transactions before the treatment are parallel to the treatment cohort. After the beginning of the treatment there appears to be a positive level shift for the treatment cohort relative to the control cohort. Because the consumption for the treatment and control cohort is parallel ex-ante, I can imagine consumption would have followed the black striped line if the treatment did not occur. Thus, I assume that the black striped line is the unobserved counterfactual. The average effect of the treatment τ is the average difference between the consumption of the treatment cohort and the unobserved counterfactual.

Mathematically I can decompose the consumption y of cohort $i \in 1, \ldots, N$ at period $t \in 1, \ldots, T$ as,

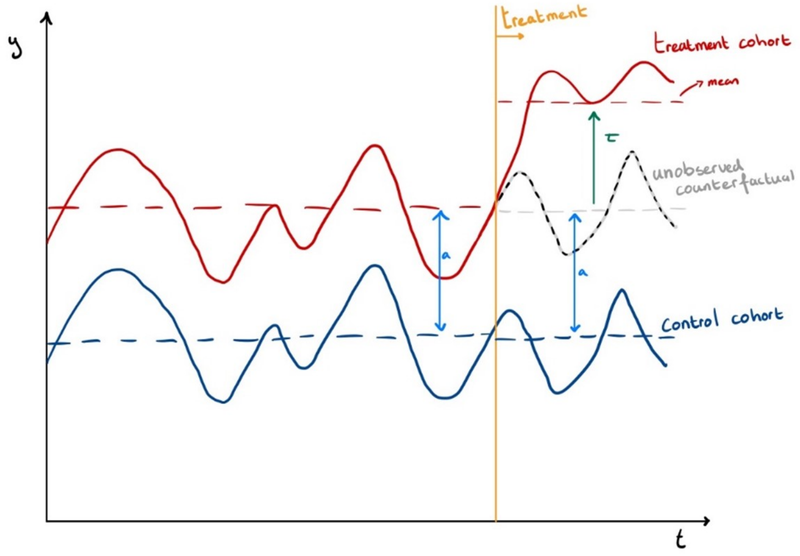

Figure 3: Schematic of difference-in-difference treatment effect.

$$y_{it} = \mu + \alpha_i + \beta_t + W_{it}\tau^{\text{dd}} + \varepsilon_{it}, \tag{3}$$

where $\mu \in \mathbb{R}$ is a constant that is time invariant and equal across cohorts, $\alpha_i \in \mathbb{R}^N$ is a time invariant cohort fixed effect, $\beta_t \in \mathbb{R}^T$ is a time fixed effect that is equal across cohorts, $W_{it} \in \{0, 1\}$ is a treatment indicator where 0 indicates no treatment and 1 indicates treatment, $\tau^{\text{dd}} \in \mathbb{R}$ is the average difference-in-differences treatment effect on the treated cohort, ε_{it} is the residual (captures what I cannot explain with the model).

The control cohort can be an average of different untreated cohorts. For example, when I apply this method to a treatment cohort such as the Valkenburg group, I create a control cohort by randomly sampling N groups with similar size and taking the arithmetic mean over each period to obtain the averaged control cohort time series.

The average treatment effect τ^{dd} can be estimated with the following estimator,

$$\hat{\tau}^{\text{dd}} = \left(\frac{1}{T_{\text{post}}} \sum_{t=T_{\text{pre}}}^{T-1} y_{0,t+1} - \frac{1}{T_{\text{pre}}} \sum_{t=1}^{T_{\text{pre}}} y_{0t} \right) - \left(\frac{1}{T_{\text{post}}} \sum_{t=T_{\text{pre}}}^{T-1} y_{c,t+1} - \frac{1}{T_{\text{pre}}} \sum_{t=1}^{T_{\text{pre}}} y_{ct} \right), \tag{4}$$

where i is equal to 0 and c for the treatment and averaged control cohort respectively, T_{post} is the length of the series after the treatment, T_{pre} is the length of the series before the treatment (Arkhangelsky et al., 2021). As one example, when I consider the consumption of the Valkenburg cohort, I can estimate τ^{dd} by first computing the average consumption ex-post minus the average consumption ex-ante for the Valkenburg cohort. Then I subtract the average consumption ex-post minus the average consumption ex-ante of the averaged control cohort. One benefit is that τ^{dd} does not depend on the estimates of the fixed effects. This is because these terms disappear from the estimator. Specifically, the fixed effect for the Valkenburg cohort α_i will disappear because I assume it is equal ex-post and ex-ante. Additionally, the time fixed effects β_t disappears because I assume it is equal for the Valkenburg cohort and control cohorts.

3.2 Synthetic control

The main downside of difference-in-differences is the parallel trends assumption. Abadie and Gardeazabal (2003) first introduced synthetic control to alleviate this issue. Synthetic control methods re-weights the untreated cohorts to compensate for the lack of parallel trends. Particularly, this re-weighted control cohort is set to match the treatment cohort ex-ante. Figure 4 shows a schematic of the synthetic control method. Again, let's assume that the untreated cohorts are made up from a random selection of clients with a similar size to the treated cohort. Now I do not make the parallel trend assumption as before. Instead I add a degree of freedom for every untreated cohort and allow the model to find a weight for each untreated control cohort such that the weighted sum of their consumption y matches that of the treatment cohort. When I am able to perfectly match the weighted control cohort to the treated cohort I obtain something similar to Figure 4.

Because I re-weight the untreated cohorts to match the treated cohort I no longer take into account unit fixed effects (α_i). Mathematically, I decompose the consumption y of cohort i at period t as,

$$y_{it} = \mu + \beta_t + W_{it}\tau^{\text{sc}} + \varepsilon_{it} \tag{5}$$

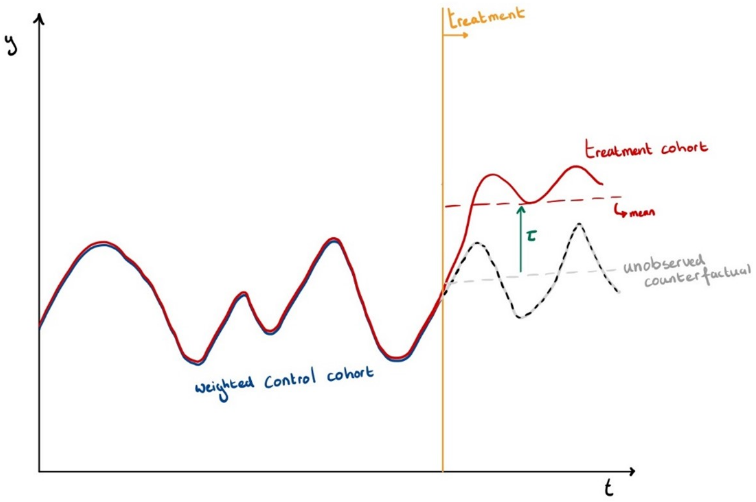

Figure 4: Schematic synthetic control treatment effect.

where $\tau^{\text{sc}} \in \mathbb{R}$ is the average synthetic control treatment effect on the treated cohort which can be estimated with,

$$\hat{\tau}^{\text{sc}} = \left(\frac{1}{T_{\text{post}}} \sum_{t=T_{\text{pre}}}^{T-1} y_{0,t+1} \right) - \left(\frac{1}{T_{\text{post}}} \sum_{t=T_{\text{pre}}}^{T-1} y_{c,t+1} \right) \text{ where } y_{ct} = \sum_{n=1}^{N} \widehat{\omega}_n^{\text{sc}} y_{nt}, \tag{6}$$

where $\omega_n^{\text{sc}} \in \mathbb{R}$ are weights (Arkhangelsky et al., 2021). In words, when I match the consumption of the weighted control cohort with the treated cohort ex-ante, and I assume that this relation remains stable ex-post, I can simple subtract the average consumption of the weighted control cohort from the average consumption of the treated cohort ex-post in order to estimate the treatment effect. Thus, I assume that the weighted control cohorts consumption ex-post is equal to the unobserved counterfactual of the treatment cohort.

I use slightly modified version of Arkhangelsky et al. (2021) formulation of the estimator of the control weights ω^{sc} which is given by,

$$\begin{aligned} \widehat{\omega}^{\mathrm{sc}} &= \arg\min_{\omega\in\Omega} l(\omega) \text{ where } l(\omega) = \sum_{t=1}^{T_{\mathrm{pre}}} \left(\sum_{n=1}^{N} \omega_n y_{nt} - y_{0t} \right)^2 , \\ \Omega &= \left\{ \omega \in \mathbb{R}_+^N : \sum_{n=1}^{N} \omega_i = 1 \right\} , \end{aligned} \tag{7}$$

where $\mathbb{R}_+$ denotes the positive real line and the elements of vector ω are the weights for the untreated control cohorts. The vector ω^{sc} is the argument of the function $l(\omega)$ at the location where the output is minimal. In particular, I search ω^{sc} for which the summed consumption of the weighted control cohorts ($i = n$ for $n \in 1, 2, .., N$) is closest to the treatment cohort ($i = 0$) ex-ante. By bounding the elements of ω between zero and one and constraining the sum of the elements of ω to one, I obtain a weighted average consumption when I sum the product of the elements of ω with the consumption of cohorts for each period ($\sum_{n=1}^{N} \omega_n y_{nt}$ for $t = 1, 2, \ldots, T$). I can solve Equation 7 numerically with the Sequential Least Squares Programming algorithm (Dieter, K., 1988).

3.3 Synthetic difference-in-differences

While synthetic control removes the parallel trend assumption, it add other issues. First, I assume that I can match the consumption of the weighted control cohorts to the treatment cohort. In practice, this is quite difficult. This can result in a bias, where I will find a treatment effect in periods where the treatment did not occur. Second, if the data is noisy, Equation 7 can cause over-fitting. Let's say I can decompose all data into a signal and noise component. The signal represents the information in the data I am interested in and the noise is the meaningless information in the data. Figure 5 shows two schematics. On the left, I happen know what the true signal is and can directly measure the noise in the data as the distance between the data and the signal. On the right, I allow the model to exactly fit the data. This may reduce the distance between each data point and the black model line but loses most of the signal. Preferably, I only capture the signal when modelling. When dealing with many parameters such as in synthetic difference I should be watchful. In particularly, the ω vector adds N parameters to the model. I can also interpret this as adding N degrees of freedom to the model to find a 'better' fit. If I

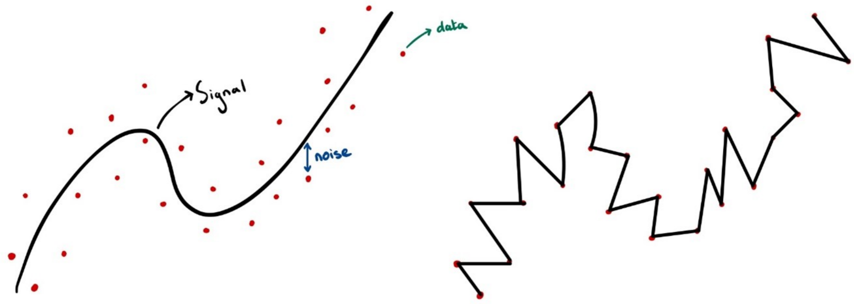

Figure 5: Left 'true' signal, data and noise. Right, example of an overfitted model that loses the information contained in the signal.

predict the unobserved counterfactual (Figure 4) with an over-fitted model I introduce a bias into the treatment effect.

Arkhangelsky et al. (2021) propose a new method called synthetic difference-in-differences that combines the attractive features of difference-in-differences and synthetic control while alleviating the aforementioned issues. I use the same decomposition of consumption as with difference-in-differences which is given by,

$$y_{it} = \mu + \alpha_i + \beta_t + W_{it}\tau^{\text{sdid}} + \varepsilon_{it}, \tag{8}$$

where $\tau^{sdid} \in \mathbb{R}$ is the average synthetic difference-in-differences treatment effect on the treated cohort which can be estimated by,

$$\hat{\tau}^{\text{sdid}} = \left(\frac{1}{T_{\text{post}}} \sum_{t=T_{\text{pre}}}^{T-1} y_{0,t+1} - \frac{1}{T_{\text{pre}}} \sum_{t=1}^{T_{\text{pre}}} y_{0t} \right) - \left(\frac{1}{T_{\text{post}}} \sum_{t=T_{\text{pre}}}^{T-1} y_{c,t+1} - \sum_{t=1}^{T_{\text{pre}}} \hat{\lambda}_t y_{ct} \right) \text{ where,}$$
$$y_{ct} = \sum_{n=1}^{N} \widehat{\omega}_n^{\text{sdid}} y_{nt}, \tag{9}$$

where $\hat{\lambda}_t \in (0,1)$ for $t = 1, 2, \ldots, T_{\text{pre}}$ is an estimated period weight and $\hat{\omega}_n^{\text{sdid}} \in (0,1)$ is an estimated control weight (Arkhangelsky et al., 2021). The intuition behind the

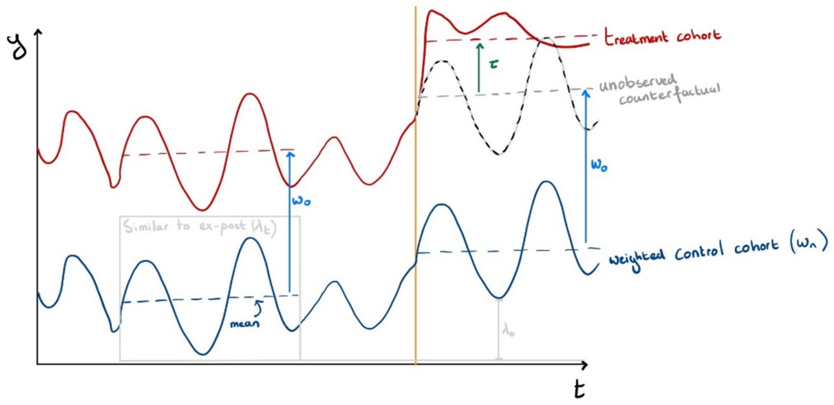

Figure 6: Schematic of the synthetic difference-in-differences treatment effect (τ).

control weights is similar as before. However, they add a constant and alleviate the overfitting problem with regularisation. Additionally, the time weights λ_t allow the model to focus on consumption in periods ex-ante that are similar in consumption ex-post for the control cohorts. Intuitively this makes sense, if the flood occurs in the summer, I prefer to measure the effect against previous summers that did not experience a flood.

Figure 6 shows a schematic of the synthetic difference-in-differences (SDID) model. The observations of the treatment cohort are shown in red and the observations of the weighted control cohort in blue. The model finds weights such that the weighted control cohorts consumption matches the treatment cohorts consumption. However, different from synthetic control, SDID adds an extra parameter ω_0 to take into account a level difference between the two series. Additionally, period weights λ_t and constant λ_0 are added which allows the model to find periods that are similar on average to the consumption of the weighted control cohort ex-post.

I use the formulation of Arkhangelsky et al. (2021) for the estimator for ω^{sdid} and ω_0, which is given by,

$$
\begin{aligned}
\left(\widehat{\omega}_0, \widehat{\omega}^{\text{sdid}}\right) &= \arg \min_{\omega_0 \in \mathbb{R}, \omega \in \Omega} l\left(\omega_0, \omega\right) \text{ where },\\
l\left(\omega_0, \omega\right) &= \sum_{t=1}^{T_{\text{pre}}} \left(\omega_0 + \sum_{n=1}^{N} \omega_n y_{nt} - y_{0t}\right)^2 + \zeta^2 T_{\text{pre}} \left|\left|\omega\right|\right|_2^2,\\
\Omega &= \left\{\omega \in \mathbb{R}_+^N : \sum_{n=1}^{N} \omega_{\text{i}} = 1\right\},
\end{aligned}
\tag{10}
$$

where ζ is set to roughly match the size of a typical one-period change in consumption for the untreated control cohorts ex-ante, $||\omega||_2^2$ is the squared second norm of the unit weights vector which ensures the weights do not deviate too much from $1/N$ to prevent over-fitting while ensuring a unique solution. For example, the global minimum of $||\omega||_2^2$ with 4 elements is obtained at (1/4, 1/4, 1/4, 1/4). The value of $||\omega||_2^2$ increases when the difference between the elements in ω increases. In other words, regularisation adds a cost to the optimisation function when more weight is put on a specific untreated cohort(s). Additionally, regularisation ensures the uniqueness of the weights (Doudchenko & Imbens, 2016). The regularisation parameter ζ is estimated by,

$$
\begin{aligned}
\zeta &= T_{\text{post}}^{1/4}\, \hat{\sigma} \quad \text{ with},\\
\hat{\sigma}^2 &= \sum_{n=1}^{N} \sum_{t=1}^{T_{\text{pre}}-1} \left(\Delta_{nt} - \bar{\Delta}\right)^2, \text{ where } \Delta_{nt} = y_{n(t+1)} - y_{nt},\\
\text{and } \bar{\Delta} &= \frac{1}{T_{\text{pre}}-1} \sum_{n=1}^{N} \sum_{t=1}^{T_{\text{pre}}-1} \Delta_{nt},
\end{aligned}
\tag{11}
$$

which measures the intensity of one-period changes in the endogenous variable scaled by the number of periods in the ex-post window. This means that the regularisation parameter increases when the expected intensity over the period T_{post} increases. In other words, we trust the weights of ω^{sdid} less when consumption y for the untreated cohorts is expected to fluctuate more ex-post. Consequently, we penalize the optimisation function $l(\omega_0, \omega)$ in Equation 10 such that it limit the captured noise. Additionally, I follow Arkhangelsky et al. (2021) to estimate λ and λ_0 by,

$$
\begin{aligned}
\left(\hat{\lambda}_0, \hat{\lambda}\right) &= \arg \min_{\lambda_0 \in \mathbb{E}, \lambda \in \Lambda} l\left(\lambda_0, \lambda\right) \text{ where,} \\
l\left(\lambda_0, \lambda\right) &= \sum_{n=1}^{N} \left(\lambda_0 + \sum_{t=1}^{T_{\text{pre}}} \lambda_t y_{nt} - \frac{1}{T_{\text{pre}}} \sum_{t=T_{\text{pre}}}^{T-1} y_{n,t+1} \right)^2, \\
\Lambda &= \left\{ \lambda \in \mathbb{R}_{+}^{N} : \sum_{t=1}^{T_{\text{pre}}} \lambda_t = 1 \right\}.
\end{aligned}
\tag{12}
$$

In words, I try to find $\hat{\lambda}_0$ and $\hat{\lambda}$ such that the difference between the weighted sum of the untreated control cohorts ex-ante plus a level $\left(\lambda_0 + \sum_{t=1}^{T_{\text{pre}}} \lambda_t y_{nt}\right)$ is as close as possible to the arithmetic mean of the untreated control cohorts ex-post $\left(\frac{1}{T_{\text{pre}}} \sum_{t=T_{\text{pre}}+1}^{T} y_{nt}\right)$. It is possible to regularise Equation 12 as well. However, I want to allow the model to load up on periods ex-ante that are like the observations after the flood. For estimating τ^{sdid} I follow algorithm 1 from Arkhangelsky et al. (2021).

Algorithm 1: SDID estimator steps

Data: W_{it}, y_{it} for $i \in 0, 1, \ldots, N$, and $t \in 1, 2, \ldots, T_{\text{pre}}, T_{\text{pre}} + 1, \ldots, T$

Result: Point estimate $\hat{\tau}^{\text{sdid}}$

Compute ζ using Equation 11;

Estimate $\hat{\omega}^{\text{sdid}}$ using Equation 10;

Estimate $\hat{\lambda}$ using Equation 12;

Estimate $\hat{\tau}^{\text{sdid}}$ using Equation 9;

3.4 Large-sample inference

The estimated treatment effect using the methods in the previous section will most probably yield a result different from zero in any scenario. This raises the question: how do I know with certainty that the effect of the treatment can be attributed to the treatment. To answer this question I need to have a method which I can use to formally test our result. Specifically, I want to apply a formal test to the null hypothesis $\tau = 0$ against the alternative $\tau \neq 0$. For this, it is useful to know the distributional properties of τ.

Arkhangelsky et al. (2021) formally show that,

$$\hat{\tau}^{\text{sdid}} \rightarrow N\left(\tau, V_\tau\right) \tag{13}$$

where τ is the real but unknown treatment effect, V_τ is the real but unknown variance of the real treatment effect and (N($\cdot$)) is the normal distribution. I follow Arkhangelsky et al. (2021) placebo variance estimation method to estimate V_τ in algorithm 2. For example, let's say I have 20 untreated control cohorts. First, I remove the treatment cohort from y and give the placebo to the first untreated cohort and estimate the effect with the remaining untreated cohorts. Second, I repeat this for the remaining 19 untreated cohort. Finally, I estimate the variance of the treatment effect by estimating the variance of the placebo treatment effects.

Algorithm 2: Placebo variance estimation

Data: y_{it} for $i \in 1, \ldots, N$, and $t \in 1, 2, \ldots, T_{\text{pre}}, T_{\text{pre}} + 1, \ldots, T$

Result: Variance estimate $\widehat{V}_\tau^{\text{placebo}}$

$j \leftarrow 0$;

while $j < N$ **do**

  Control cohort j receives placebo and is set as the treatment cohort;

  Estimate $\hat{\tau}_j$ against remaining untreated control cohorts;

  $j \leftarrow j + 1$;

end

$\widehat{\mathrm{V}}_\tau^{\text{placebo}} \leftarrow \frac{1}{N} \sum_{j=1}^{N} \left(\hat{\tau}_j - \frac{1}{N} \sum_{j=1}^{N} \hat{\tau}_j\right)^2$

I now have all the information to formally test the null hypothesis against the alternative hypothesis by,

$$t_{\hat{\tau}} \overset{\text{def}}{=} \frac{\hat{\tau}}{\sqrt{\widehat{V}_\tau^{\text{placebo}}}} \sim t(n-k), \tag{14}$$

where $t_{\hat{\tau}}$ is the test statistic, t is the student-t distribution with $n - k$ degrees of freedom, n is the number of observations and k is the number of estimated parameters. I use the student-t distribution because I approximate the relation in Equation 13 with an estimate of V_τ. However, in practice $n-k$ becomes very larges which causes the student-t distribution to converge back the normal distribution. Using Equation 13 I can build a $1-\alpha$ confidence interval by,

$$\tau \in \hat{\tau} \pm Z_{\alpha/2}\sqrt{\hat{V}_\tau^{\text{placebo}}} \tag{15}$$

where $z_{\alpha/2}$ follows the standard normal distribution (for large samples). When this confidence interval does not include zero, I can reject the null hypothesis with significance level of α. For example, the 99% confidence interval is given by,

$$\tau \in \hat{\tau} \pm 2.58\sqrt{\hat{V}_\tau^{\text{placebo}}} \tag{16}$$

In other words, there is a 99% probability that the interval $\tau \in \hat{\tau} \pm 2.58\sqrt{\widehat{V}_\tau^{\text{placebo}}}$ contains the true parameter τ.

The validity of the placebo variance estimation relies on the homogeneity of variances assumption. This means that the control cohorts have similar variance relative to the treatment cohort. We can test this assumption by performing a test of the difference of two variances for each treatment/control pair (Heij et al., 2004). The hypothesis of the tests are: $H_0 : \sigma_1^2 = \sigma_2^2$ against $H_1 : \sigma_1^2 \neq \sigma_2^2$. I can perform the test by computing the test statistic by,

$$F = \frac{s_1^2}{s_1^2}, \tag{17}$$

where s_i for $i \in \{1, 2\}$ are the estimated variances of the transaction data before the flood. I reject the null when the p-value is smaller than α (commonly 0.05 or 0.01). The p-value is computed by,

$$p\text{−value} = \text{Fcdf}(F, df_1, df_2), \tag{18}$$

where Fcdf is the CDF of the F distribution, df_1 are the degrees of freedom of the first sample and df_2 are the degrees of freedom of the second sample. The degrees of freedom can be computed with $n-1$, where n is the length of the sample. For example, take two samples of $n = 100$. The first sample has a variance of 1 and the second sample has a variance of 3. First, I compute the test statistic which is $F = 0.33$. Second, I compute the p-value which is $\text{Fcdf}(0.33, 99, 99) = 0.00$. Thus, I can formally reject H_0 at an $\alpha = 0.01$ significance level. Consequently, the assumption of homogeneity does not hold and we cannot use this sample in the placebo variance estimator.

4 Data

4.1 Flood and emergency ordinances

To answer the research questions, I need to know which areas where impacted by the 2021 Limburg flood. Therefore, I obtain flood and evacuation area coordinates from Slager et al. (2021). These data contain areas of inundation, evacuations, blackouts and sludge blockages in the Amersfoort coordinate system. Figure 7 shows a subset of the data of flood and evacuation information of the area surrounding Valkenburg (ENW, 2021). The line shows the river Geul, red and yellow are evacuated areas, lightblue inundated areas, light green indicates power blackouts and darkblue ares are blockages caused by sludge. Some inundated areas were not evacuated while others were. This was caused by an underestimation of the severeness of the flood which surprised many as the crisis moved from south to north (ENW, 2021). Additionally, because of the relatively small scale of the flood, I do not differentiate between inundated and evacuated areas and inundated and unevacuated areas. The two areas of interest I use in this thesis are inundated and uninundated and evacuated.

4.2 Clients

To infer the economic impact of a flood I need to find ABN AMRO Bank (AAB) clients that live in the affected area. Consequently, I obtain client information data from ABB ($9.197 \cdot 10^6$ clients). These data contain information such as account numbers, types of clients (individuals/business), names, birthdates, and residential zip-codes. I transform the data by filtering for individuals from the Netherlands. Furthermore, I remove all deceased clients that were in the data prior to the flood.

4.3 Building information

I require residential coordinates of ABBs clients to create groups containing the inundated and evacuated. For this purpose, I obtain the Basisregistratie Adressen en Gebouwen (BAG) (Kadaster, 2022). The BAG ($9.096 \cdot 10^6$ observations) is maintained by munici-

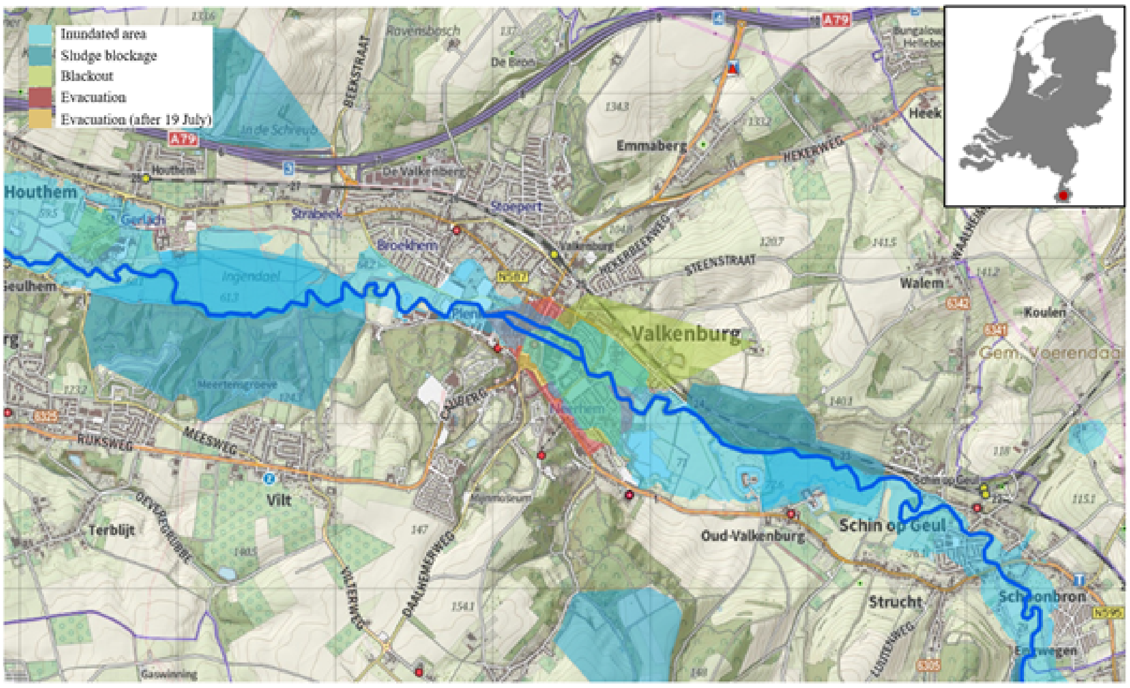

Figure 7: Map of flood and evacuation information from ENW (2021) for the area surrounding Valkenburg. The darblue line shows the river Kleine Geul, red and yellow are evacuated areas, lightblue inundated areas, light green indicates power blackout and darkblue ares are blockages caused by sludge.

palities and is updated daily. It contains information such as status, construction year, use-purpose, zip-code and most importantly, coordinates (Amersfoort coordinate system) on every building in the Netherlands. Different buildings can have the same zip-code. As such, different clients can live on the same zip-code with slightly different locations (e.g. on the same street at a different number). Therefore, I do not know exactly where a client lives, but do know all coordinates of the buildings on the street where the clients lives. As such, I transform the data by computing the centre point of all coordinates of a single zip-code and use that as the estimate for the location of all the addresses in a zip-code. Next, I join the client data with these location estimates by matching zip-codes of the client data with the transformed zip-codes. Thus, I obtain an estimate for the residential location in Amersfoort coordinate system for each ABB client. Moreover, I retrieve

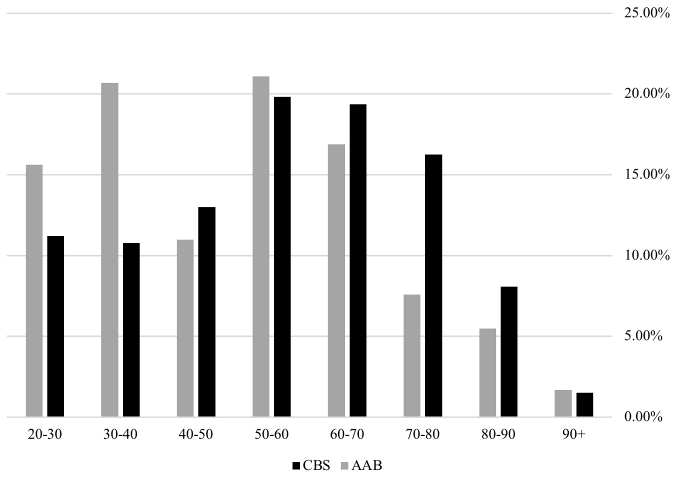

Figure 8: Age distribution of Valkenburg aan de Geul population and AAB inundated cohort (CBS, 2022).

all clients that live in the inundated Valkenburg area and evacuated-but-not-inundated areas. Additionally, I sample 30 random cohorts for both the evacuated and inundated cohort with a similar sizes from the entire Dutch AAB population. Finally, I anonymize the data by removing all names and postal codes from the final cohorts. In summary, I store 249 inundated Valkenburg clients and 30x249 control clients. Additionally, I obtain 1,996 evacuated-but-not-inundated clients and 30x1,996 control clients.

Figure 8 shows the age distribution of the population of the municipality of Valkenburg aan de Geul and the inundated cohort (CBS, 2022). The total populaton size of Valkenburg aan de Geul is equal to 16,353. Furthermore, ENW (2021) find that roughly 3,840 to 4,160 individuals were impacted. I am only interested in the adult population. Therefore, I assume the distribution of the entire population is similar to the inundated population. Therefore, I remove the CBS ratio of the 0 to 20 year from the ENW estimate, which is approximately 18%. This results in a adult inundated population of 3,150 to 3,411 people. Additionally, Figure 8 shows that the inundated cohort is younger on average than the population.

4.4 Transactions

To infer economic impact of the 2021 Limburg flood, I obtain high-frequency SEPA and POS transaction data from AAB from 1 Jan 2018 to 1 Feb 2022 ($6.480 \cdot 10^9$ and $4.285 \cdot 10^9$ observations, respectively). These data cannot be shared publicly because they are proprietary and privacy sensitive. SEPA is a system that facilitates cashless transactions in Europa. Table 2 shows summary statistics for the raw SEPA and POS transactions. SEPA clearly shows outlier behaviour with a minimum and maximum in the order of tens of billions. Additionally, mean and median of transaction amounts are quite different. Consequently, skewness is likely present in the data. Furthermore, the distance between minimum and 25% quantile and the 75% quantile and maximum show signs of high kurtosis. The PIN data seems more realistic. However, since outliers are present in the SEPA data may also be the case in the PIN data. Additionally, the data shows signs of skewness and high kurtosis as well.

To get a better understanding of the data I split the SEPA transactions into two sets. One bounded between -100,000 and 0 EUR, the other bounded between 0 and 100,000 EUR. Additionally, I bound the PIN transactions from -5,000 to 0 EUR (as these are mostly negative). By bounding I keep more than 99% of the data while removing most outliers. Additionally, the data is easier to interpret. SEPA$^+$, which is bounded between 0 and 100,000, includes transactions such as salaries, pensions and allowances. Mean transactions size is relatively high at 740 EUR while the median is approximately 64 EUR, which indicates that large transactions positively skew the distribution. Moreover,

Table 2: Summary statistics of raw SEPA and PIN transaction amounts from 1 Jan 2018 to 1 Feb 2022 for the Netherlands.

Type	Size ($\cdot 10^9$)	Mean [EUR]	Stdev [EUR]	Min [EUR]	25% [EUR]	50% [EUR]	75% [EUR]	Max [EUR]
SEPA	6.48	-79.80	$1.34 \cdot 10^7$	$-4.90 \cdot 10^{10}$	-53.11	-7.99	48.50	$4.90 \cdot 10^{10}$
POS	4.30	-25.49	123.70	$-50,000$	-25.42	-11.40	-5.00	$50,000$

Table 3: Summary statistics of bounded SEPA and POS transactions from 1 Jan 2018 to 1 Feb 2022 for the Netherlands. SEPA low and high are bounded between -100,000 and 0 EUR and 0 and 100,000 EUR, respectively. Additionally, POS is bounded between -5,000 and 0 EUR.

Type	Size ($\cdot 10^9$)	Mean [EUR]	Stdev [EUR]	Min [EUR]	25% [EUR]	50% [EUR]	75% [EUR]	Max [EUR]
$SEPA^+$	2.28	739.96	3780.24	0.00	20.00	63.80	300.00	$100,000$
$SEPA^-$	3.63	-454.27	3072.61	$-100,000$	-144.94	-45.00	-15.24	0.00
POS	4.27	-25.05	68.57	$-5,000$	-25.50	-11.48	-5.00	0.00

standard deviation is high at approximately 3800 EUR, indicating some transactions run into the tens of thousands. $SEPA^-$, which is bounded between -100,000 and 0 EUR, includes transactions such as online purchases and rent payments. Interestingly, the absolute value of the mean and median are lower than in $SEPA^+$. However, skewness is still present in the data. Summary statistics do not change substantially by bounding PIN transactions but for the standard deviation which is much lower as expected.

I transform the data by matching the account numbers from the cohorts from the previous section to the SEPA and POS data. As such, I obtain high-frequency anonymized SEPA and PIN data from 1 Jan 2018 to 1 Feb 2022 for each cohort. Furthermore, I combine $SEPA^-$ and POS data for each cohort with the bounds from the previous paragraph. In the following section, these transaction amounts will be referred to as transactions. All transformations have been programmed in Python for this thesis.

Each transaction contains textual information. It may contain a description of why the transaction was made or the name of the counterparty that was on the other side of the transactions. Unfortunately, because of data quality issues, these labels are sparse and often wrong. This means many transactions are not labelled or incorrectly labelled. This is especially true for SEPA transactions, where I expect most of the damages will be recorded.

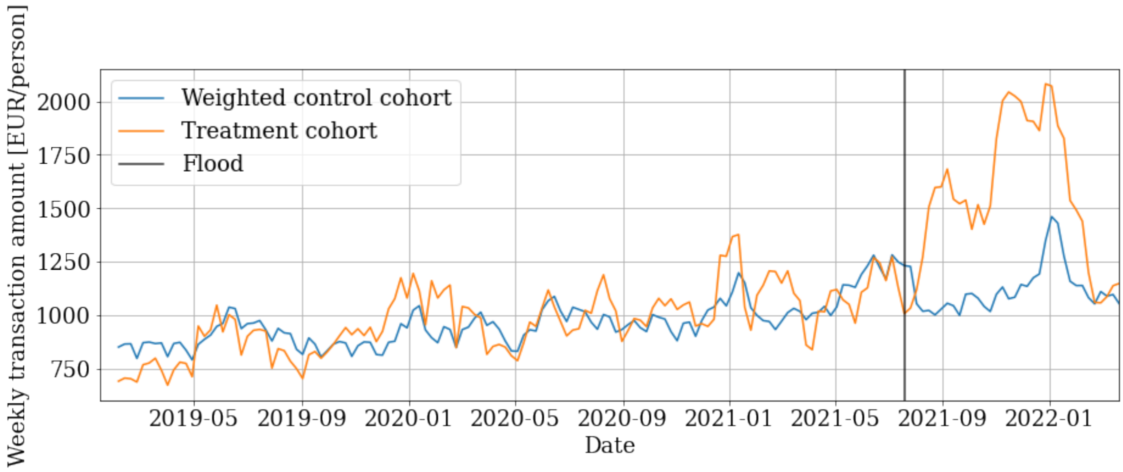

Figure 9: Four-week rolling-window of weekly transaction amounts for the control cohort with estimated SDID unit weights ($\hat{\omega}^{\text{sdid}}$) and the treatment cohort from Jan 1, 2019 to Apr 4, 2022.

5 Results

The following section shows the results from the model explained in section 3. First, the duration of the economic impact caused by the 2021 Limburg flood is investigated. Second, the value of economic impact caused by the flood is determined. Finally, I look into the impact to different transaction sizes.

5.1 Duration economic flood impact

Figure 9 shows a four-week rolling-window of total weekly transactions for the inundated treatments cohort and weighted control cohort. The latter is created by taking the weighted sum of the control cohorts total weekly transaction using $\hat{\omega}^{\text{sdid}}$ from Equation 10. Total transactions are increasing slowly over time for both cohorts. Additionally, around Christmas the transactions peak. Moreover, the transactions of the treatment cohort follow the modelled control relatively well before the flood, albeit with more fluctuations. This is expected since the modelled control cohort is a weekly weighted sum of the original control cohorts, which can lower the standard deviation because of subadditivity ($\sigma_{1+2} \leq \sigma_1 + \sigma_2$). This does not influence the estimated effect, because it

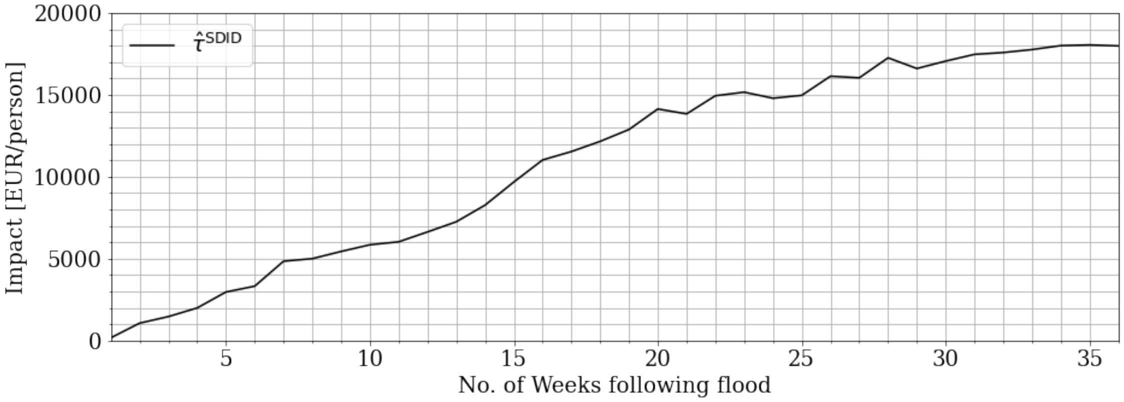

Figure 10: Economic impact point estimates ($\hat{\tau}^{\mathrm{sdid}}$) for an expanding window.

only takes into account level differences. Furthermore, I formally test for differences in the weekly standard deviations between the treatment cohort and all the control cohorts before the flood, separately. I cannot reject the null hypothesis of equal variances for any of the pairs at a 1% significance level. Consequently, the assumption of homogeneity in the variances holds. In other words, the variances of the treatment and control cohorts are statistically the same. Moreover, weekly transactions seem to increase substantially after the flood for a period lasting several months before reverting back to that of the weighted control cohort.

Figure 10 shows model estimates $\hat{\tau}^{\mathrm{sdid}}$ from algorithm 1 in an expanding window for the inundated treatment cohort. The expanding window starts with the transaction amounts of the fist week after the flood to compute $\hat{\tau}_1^{\mathrm{sdid}}$. Then, it computes $\hat{\tau}_2^{\mathrm{sdid}}$ using the transaction amounts from first and second week. Next, $\hat{\tau}_3^{\mathrm{sdid}}$ is computed using transaction amounts from the first, second and third week. This is done until all weeks after the flood are included in the final estimate $\hat{\tau}_N^{\mathrm{sdid}}$ where N is the total number of weeks after the flood. The result of Figure 10 can be interpreted as the cumulative economic impact caused by the flood up to week t. Interestingly, the economic impact is increasing and concave down in time before normalising. This means, the economic impact levels off to a value around 18,000 EUR per person after 35 weeks, which is in line with the visualisation in Figure 9. Taken together, these findings show that the transaction data

Table 4: SDID model results using aggregated and anonymized total transactions for the inundated and uninundated & evacuated sample. Total transactions is the average of the sum of all transactions for the treatment cohort. Additionally, τ shows the point estimate and standard error of the SDID model for the economic impact of the flood (see Equation 16). Moreover, the duration is the number of weeks which is included in estimating the effect. *** shows a 1% significance level for the estimate. Duration of uninundated and evacuated is not available since there is no significant effect.

Sample	Total trans. amount [EUR/p]	τ (s.e.) [EUR/p]	Duration [weeks]
Inundated	$57,099$	$18,045^{***}$ $(2,161)$	35
Uninundated & evac.	$39,612$	348 $(1,516)$	N/A

allows measurement the duration of economic impact caused by the 2021 Limburg flood and support hypothesis 2 and 3 from section 1.

5.2 Economic flood impact

Table 4 shows the model estimates of the economic impact caused by the 2021 Limburg flood. For the inundated cohort, the economic impact is approximately 18,045 EUR per person on average at a 1% significance level from 19 Jul to 3 Apr. Consequently, from the total transaction amount of 57,099 EUR per person, approximately 32% can be attributed to the flood (18,045/57,099). How do I determine the significance level? Arkhangelsky et al. (2021) show that the SDID estimator is asymptotically normal. Therefore, I can draw a distribution of the economic impact using $\hat{\tau}^{\text{sdid}}$ and $\hat{V}_\tau^{\text{placebo}}$ as estimates for the mean and variance of the asymptotic normally-distributed economic impact. Figure 11 shows the inferred distribution of the total average economic impact per inundated residents after the 2021 Limburg flood. Specifically, the null hypothesis of $\tau = 0$ against the alternative hypothesis $\tau \neq 0$ is rejected when the probability (p-value) that $\tau = 0$ is small. In this case, the value is equal to 0.00%, which is well below 1%. Therefore, I formally reject the null hypothesis of $\tau = 0$ with a 1% significance level. This supports hypothesis that high-frequency transaction data allows me to compare between flood-affected individuals and similar but unaffected individuals to measure the economic impact caused by the flood. Another way to interpret the estimate is that there is a 99% probability that the interval 12,470 to 23,620 EUR per person contains the *true* economic impact caused by the 2021 Limburg flood.

Table 4 shows no significant economic impact for the uninundated and evacuated cohort caused by the 2021 Limburg flood. The estimated effect is equal to approximately 348 EUR per person on average with a standard error of 1,516 EUR. This indicates the model estimates a 99% confidence interval for the economic effect to be -3,563 to 4,260 EUR per person. Thus, using this approach, there is no way of telling whether there is an effect (since zero is contained in the interval). Consequently, I am not able to measure the economic impact for this group using high-frequency bank transaction data. However, the larger size of the cohort reduces the standard error relative to the inundated cohort (1,516 versus 2,161).

Interestingly, the point estimate of the total economic impact per person of the inundated is contained in the range estimated by the ENW (2021) of 15,830-21,596 EUR per person (average of 18,713 EUR). However, the range of the SDID estimate is wider, which could be caused by the following reasons. First, the HIS SSSM, used by the ENW to estimate damages, may underestimate the variability in maximum damage values. Second, the outliers in high-frequency transaction data creates noise which inflates the placebo variance. This is because the placebo variance is computed by estimating $\hat{\tau}^{\text{sdid}}$ for all control cohorts and estimating the standard deviation. If the transaction data of the control groups contains outliers, this will increase the standard deviation. It is important to note that the variance of the economic impact is derived from the variance of control cohorts after the flood. Consequently, it is not a function of the economic impact itself and should not be interpreted as the range of the economic impact. However, it is useful to formally test hypotheses. Furthermore, larger sample sizes should lower the standard error, similar to the uninundated and evacuated cohort.

Table 5 shows SDID model estimates for three transaction bins. These bins contain

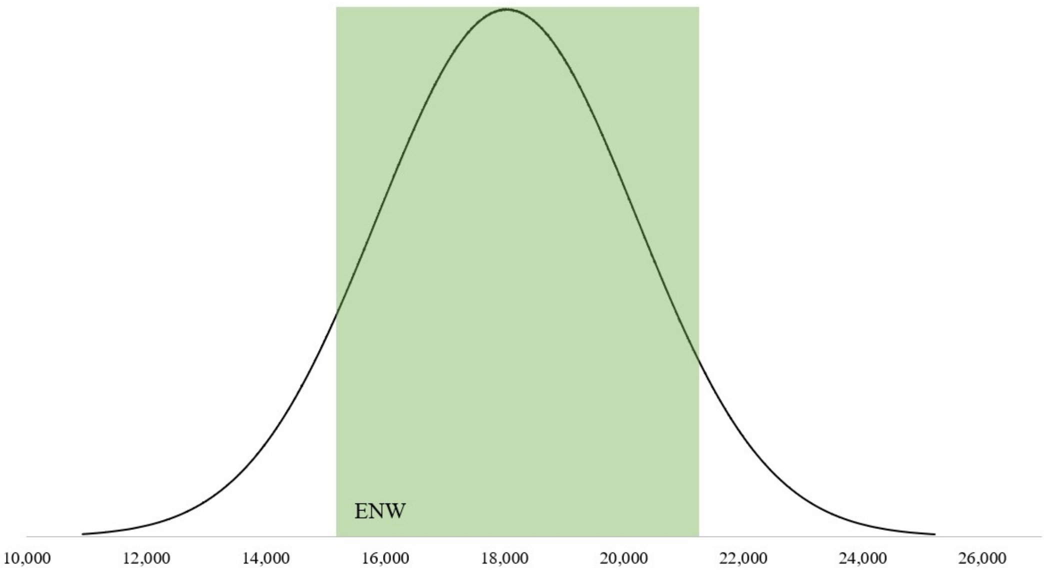

Figure 11: Distribution of estimate economic impact for inundated residents caused by the 2021 Limburg flood in EUR per person. ENW (2021) flood damage estimate is indicated by the green box.

transaction sizes between certain amounts and allow me to get more insights into which categories of transactions are causing the significant economic impact. First, the economic impact between 0 and 1,000 EUR per transaction, is approximately 969 EUR per person on average, at the 5% significance level. Additionally, I estimate that the duration of the impact is roughly 23 weeks in a similar manner as before. Second, the economic impact of the 2021 Limburg flood, for transactions between 1,000 and 5,000 EUR, is approximately 2,268 EUR per person on average, at a 5% significance level. The duration for this bin is roughly 31 weeks. Finally, $5,000 - 100,000$ per transaction, shows a significant difference in transactions between the treatment and control cohort of approximately 13,614 EUR per person on average at the 1% level with a duration of approximately 34 weeks. However, the assumption of homogeneity across variances does not always hold. Consequently, some control cohorts are removed that do not pass the test. As such, the results should not be compared one-on-one with the results in Table 4.

The first bin $(0 - 1,000$ EUR) in Table 5 includes transactions like the purchases of

Table 5: SDID model results using aggregated and anonymized total transactions for the inundated sample segregated in bins. Total transactions is the average of the sum of all transactions for the treatment cohort. Additionally, τ shows the point estimate and standard error of the SDID model for the economic impact of the flood (see algorithm 1 and algorithm 2). Moreover, the duration is the number of weeks which is included in estimating the model. *** and ** shows a 1% and 5% significance level of the estimate, respectively.

	Transaction bins [EUR]		
	$0 - 1,000$	$1,000 - 5,000$	$5,000 - 100,000$
τ [EUR/p]	969^{**} (406)	$2,268^{**}$ $(1,092)$	$13,614^{***}$ $(1,515)$
Total trans. amount [EUR/p]	23,364	9,789	22,923
Duration [weeks]	23	31	35

appliances, home renovation goods, minor construction work and relocation costs. These transactions are probably relatively easy to finance for the affected individuals, since it only takes up approximately 4% of total transactions in that bin. Additionally, it takes a relatively short amount of time for the transactions of the treatment group to return to levels of the control groups relative to the other bins. This indicates the relative ease of replacement of durable goods (e.g. washing machine) in a local flood, since supply at stores is enough to meet the demand of the affected region. Furthermore, it could indicate affected individuals are able to pay for the replacements themselves instead of waiting for the insurance money to arrive to make the transaction. Moreover, it may indicate individuals returning to their homes since relocation costs have diminished.

The second bin ($1,000-5,000$ EUR) in Table 5 includes transactions like home renovations such as flooring repair or major car repairs. They represent 23% of total transactions, are probably more difficult to finance than the first bin, but are manageable still with savings. Duration is higher than the smallest bin. This may be caused by the longer times of completion of home renovation projects. It may be possible that local labour shortages prolong this effect since the region is not able to source enough workers to renovate so

Table 6: SDID model estimates for POS transactions with a 99% confidence interval and groceries category.

Sample	Total trans. amount [EUR/p]	τ (s.e.) [EUR/p]	Duration [weeks]
Total POS	$6,135$	$1,096^{***}$ (171)	35
Groceries	$1,446$	-101^{**} (43)	37

many homes at once.

The third bin ($5,000 - 100,000$ EUR) in Table 5 includes larger home renovations such as kitchen and bathroom replacements. These transactions represent roughly 60% of total transactions made in this bin. It seems plausible that these transactions are harder to finance than the previous two bins. On the one hand, affected individuals are likely to hold off on renovations until they have received insurance payouts. On the other hand, there can be be smaller upfront payments that are done to finance big renovation projects that leak into the first or second bin. For example, a kitchen renovation may require an up-front payment of 5,000 EUR before renovations begin, for a total cost of 20,000 EUR. The duration of the third bin is the longest at 35 weeks. As stated before, local labour shortages can prolong the duration since the region is not be able to source enough workers to renovate so many homes at once.

Unfortunately, I am unable test whether the examples of the previous three paragraphs are true because of data limitations. This is because the labelling of the transactions is very sparse for SEPA transactions. However, I can go deeper into the POS transaction data, because the descriptions of transactions are generally more informative than in SEPA transactions. Table 6 shows that the economic impact to POS transactions of the 2021 Limburg flood is approximately 1,096 EUR per person, at the 1% significance level. All of these transactions are in the first two bins of Table 5. While POS transactions are roughly 19% of total transactions in the first two bins, the economic impact caused by the 2021 Limburg flood is roughly one-third of the total effect of POS transactions. This is expected since these bins also contain regular SEPA transactions such as health insurance and rents, that are unlikely to be affected.

Interestingly, the economic impact to groceries transactions is -101 EUR per person on average for a duration of minimally 37 week, at a 5% significance level. Groceries represent roughly 30% of total POS transactions when the flood effect is excluded. Consequently, it is a very important measure of non-durable goods spending (e.g. food, clothing and toilet paper). Additionally, the effect does not appear to end before Apr 3, indicating longer-term behaviour, especially since these goods are replaced at a high frequency. Groceries transactions are approximately 7% lower for the treatment cohort relative to the control cohort ($\frac{1{,}446}{1{,}446+101} - 1$). In summary, the significant negative effect shows there are (longer term) changes in consumer behaviour caused by the flood.

6 Discussion

In this sections I discuss the findings of this thesis and the limitations of my analysis. The aim of this research was to estimate the economic impact caused by the 2021 Limburg flood using geo-located high-frequency transaction data.

In summary, I have been able to leverage high-frequency bank transaction data to estimate the duration and size of economic impact of the 2021 Limburg flood. High-frequency bank transaction data give a complete view of consumer spending. Additionally, differences in consumer spending between an inundated cohort and its unobserved counterfactual can be used to estimate the economic impact of a flood for inundated individuals. Furthermore, transaction data allows for the comparison between flood affected individuals and similar but unaffected individuals to measure the economic impact and duration caused by the flood.

Modelling high-frequency bank transaction data using synthetic difference-in-difference I found that the economic impact of the 2021 Limburg flood is approximately 18,045 EUR on average for the inundated cohort. The economic impact for the inundated is approximately 32% of transactions during the 35 week period after the flood, or approximately 46% of transactions relative to the transactions of the control cohorts. Furthermore, I found that the duration of the economic impact is approximately 35 weeks starting from the week after the flood. Additionally, a large portion of the impact, or 13,614 EUR per person on average, is in the form of transaction sizes over 5,000 EUR. Furthermore, the impact with transaction sizes ranging between 1,000 and 5,000 is approximately 2,268 EUR per person on average. Moreover, the impact for transactions smaller than 1,000 EUR is approximately 969 EUR per person on average.

The fact that I find an economic impact over a long period for the inundated may be explained by (i) large transactions of durable goods that are normally done over several years are compressed into a smaller period, (ii) renovations are constrained by insurance payouts and labour shortages, (iii) long dry-up times of inundated residences. For example, a kitchen may be replaced once every 20 years by a local construction company.

The loss of many kitchens at once in the area will increase demand for local shops and contractors more than supply can catch up. Additionally, walls and floors may be soaking wet, causing wait times because of the drying process. I conclude that high-frequency bank-transactions data when modelled with a synthetic difference-in-differences model is most helpful when evaluating the economic impact and duration of floods for inundated individuals.

Data quality issues are likely the cause of not being able to partition transactions into the flood damage dimensions from Jonkman et al. (2008). Each transaction carries textual information. For example, an insurance payout from a particular insurance company will have a relevant description and company bank account number. However, for each insurance company, these descriptions and bank account number are different and hard to label accurately. On the one hand, these issues make me unable to test differences in flood damage dimensions. On the other hand, banks are continuously improving labelling function. Therefore, more should be possible in the future and a repeat of this analysis should lead to more detailed results.

ENW (2021) found a damage estimate of 15,830 to 21,596 EUR per affected individual with HIS SSSM, similar to the SDID estimate of 12,470 to 23,620 EUR per affected individual. However, the range of the SDID estimate is larger and includes more flood damage dimension. For example, clean up cost, adjustment in consumption patterns and temporary housing of evacuees is not included in HIS SSSM. It may be that the adjustment in consumption patterns is negative and of similar magnitude as clean up costs and temporary housing of evacuees. The wider range of the SDID can be explained by the relatively small sample I have used (249 clients in the cohorts), which makes the placebo variance more susceptible to outliers.

My inundated sample is different in composition than the average of the inundated population. Nationally, AAB is a large bank with roughly 22% market share, but this is mostly contributed to larger cities (Neuteboom et al., 2021). Additionally, the AAB sample in Valkenburg aan de Geul is younger on average than the population, 3,280 adult individuals were impacted in the Geul area, my sample holds approximately 250 individuals, or 8% of the population reported by ENW (2021). Furthermore, standard errors may reduce when the sample size increases, because the control cohorts size also increases.

This approach may be structurally over or under estimating the economic impact. On the one hand, insurance companies or governments may overestimate damages. The difference may be used by the inundated to pay off mortgages or pay for holidays. Therefore, transaction data cannot yet be used as an independent economic impact estimate. However, when better labelling functions are available, it may be possible to filter out these transactions. On the other hand, the inundated may not be able to finance damages when insurance companies or the government underestimate the claim. If these individuals have inadequate savings or income to replace or repair damaged items, then the damages would not show up in the data. For example, if someone loses a newly installed kitchen worth 20,000 EUR, and replaced it with an inferior kitchen of 15,000 EUR, the SDID model would report the 15,000 EUR as the damage value. The approach is also likely to miss damages to rented houses and apartments because these damages will not be paid for by the renter but by the real estate investor. One possible solution in future research is to filter out individuals that pay a monthly rent. This should be possible when labelling functions improve.

Another reason that the approach may be structurally over estimating the economic impact is because I cannot separate between households and individuals. A household can have several bank accounts. For example, take a household with one wage-earner and one homemaker. Both have a private bank account and one shared account. In a perfect scenario, all accounts are held at AAB. However, the couple can have bank account at different banks. If only the account of the wage-earning is held ABB, and all damage dimensions are paid for using that account, it will overestimate the economic impact of that individual, because in reality it is shared by two individuals.

I was unable to find an economic impact or duration for the uninundated and evacuated using high-frequency bank transaction data and the SDID model. This may be explained by the evacuations leading to changes in consumption that are smaller than the average variability of the control cohorts or nonexistent. Therefore, I cannot differentiate between the impact of the evacuation and other unrelated impacts. This issue may be addressed by looking into specific categories when labelling functions are sufficient. For example, it may be possible to find a change in consumption at hardware stores, while total consumption are unaffected.

Looking forward, high frequency bank transaction data should be the basis for estimating the economic impacts of future natural disasters. Additionally, when labelling functions improve, and flood damage dimensions can be estimated separately, it can lay the foundation for calibrating excising damage models. For example, I recommend calibrating existing damage functions, similar to Wing et al. (2020) study using claims from the National Flood Insurance Program in the United States. This can improve future damage estimate from governments, insurance companies and reinsurance companies to better price disasters, possibly making infrastructure project and insurance policies more efficient. Additionally, I recommend repeating an experiment similar to Sampson et al. (2014) to evaluate the output of proprietary reinsurance and vendor catastrophe models. Furthermore, investigating the income side of the bank transaction data will give more insight into how affected individuals financially manage the impact or natural disasters. In particular, it could give more insight into the differences between social layers. For example, it can show detailed differences between home owners and renters, or wealthy and poor individuals. I also recommend pooling transaction data of major banks to limit over estimating the economic impact when a household holds accounts at several banks.

7 Conclusion

In the following section I revisit the research questions from section 1 followed by answering the questions one-by-one.

Research question 1: How can high-frequency transaction data be used to estimate the duration of the 2021 Limburg flood impact to individuals?

Answer: Using synthetic difference-in-differences, the economic impact can be estimated by computing the difference in transactions between the affected individuals and the unobserved counterfactual. Furthermore, it is assumed that the unobserved counterfactual observes the behaviour of the affected individual as if the flood did not occur. To do this exercises, it is required to have both Point-of-Sale and direct cashless transaction date between individuals and companies (SEPA in Europe). Furthermore, it is required to have access to not only the transaction data of the affected individuals but also unaffected individuals. To construct the unobserved counterfactual, several cohorts of unaffected individuals must be created of similar size as the affected cohort. The data is then aggregated per cohort and used as input in the model. The model constructs the unobserved counterfactual and computes the effect for each week after the flood up to now. This results in a time series of economic effects, that levels off some time after the flood, indicating the duration of the effect.

Research question 2: How can high-frequency transaction data serve as a proxy for the economic impact to individuals caused by the 2021 Limburg flood?

Answer: The answer to this question is mostly similar to the previous question. However, the result of the previous question is used to answer this question. When the economic duration in known, the final estimate of the economic effect can be determined by using the final value of economic impact from the time series of economic effects. Additionally, the standard deviation can be estimated using the placebo variance estimator. The economic impact of the disaster exists and is measurable when the estimate is statistically significant.

Research question 3: What is the economic impact and duration of the impact caused by the 2021 Limburg flood to individuals using high-frequency transaction data?

Answer: The economic impact caused by the 2021 Limburg flood to inundated individuals was equal to 18,045 EUR on average per person and lasted 35 weeks. Furthermore, there is no measurable economic impact to uninundated and evacuated individuals.

References

Abadie, A., & Gardeazabal, J. (2003). The economic costs of conflict: A case study of the basque country. *American Economic Review*, *93*(1), 113–132.

Arkhangelsky, D., Athey, S., Hirshberg, D., Imbens, G., & Wager, S. (2021). Synthetic difference-in-differences. *American Economic Review*, *111*(2), 4088–4118.

Bermuda:Re+ILS. (2021). Air and rms dominate modelling among bermuda reinsurers [Accessed: 2022-08-25].

Brahin, P. (2015). *The essential guide to reinsurance*. Swiss Re.

Carvalho, V. M., Garcia, J. R., Hansen, S., Ortiz, Á., Rodrigo, T., Rodríguez Mora, J., & Ruiz, P. (2021). Tracking the covid-19 crisis with high-resolution transaction data. *Royal Society Open Science*, *8*, 210218.

CBS. (2022). Inwoners per gemeente [Accessed: 2022-05-01].

Chetty, R., Friedman, J., Hendren, N., Stepner, M., & Team, O. I. (2020). The economic impacts of covid-19: Evidence from a new public database built using private sector data.

De Bruin, K., Wagenaar, D., Slager, D., De Bel, M., & Burzel, A. (2015). The updated and improved method fro flood damage assessment ssm2015: Explanation, motivation and comparison to his-ssm (ssm2015).

Dieter, K. (1988). *A software package for sequential quadratic programming.*

DNB. (2021). Corona stuwt gebruik pin in 2020: Voor het eerst betalen consumenten in alle branches vaker met pin dan contant [Accessed: 2022-08-22].

Doudchenko, N., & Imbens, G. (2016). *Balancing, regression, difference-in-differences and synthetic conrtol methods: A synthesis.*

ENW. (2021). *Hoogwater 2021 feiten en duiding.*

Gelman, M., Karvi, M., S. Shapiro, Silverman, D., & Tadelin, S. (2014). Harnessing naturally occurring data to measure the response of spending to income. *Science*, *345*(6193), 212–215.

Grogan, T., & Angelo, W. (2005). Katrina keeps inflation roaring. *255*, 66–68.

Heij, C., de Boer, P., Franses, P., Kloek, T., & van Dijk, H. (2004). *Econometric methods with applications in business and economics*. Oxford University Press.

Jarmin, R. (2019). Evolving measurement for an evolving economy: Thoughts on 21st century u.s. economic statistics. *Journal of Economic Perspectives*, *33*(1), 165–184.

Jonkman, S., Bočkarjova, M., & Bernardini, P. (2008). Integrated hydrodynamic and economic modelling of flood damage in the netherlands. *Ecological Economics*, *66*, 77–99.

Kadaster. (2022). Basisregistratie adressen en gebouwen [Accessed: 2022-03-01].

Kreienkamp, F., Philip, S., Tradowsky, J., Kew, S., Lorenz, P., & Arrighi, J. (2021). Rapid attribution of heavy rainfall events leading to the severe flooding in western europe during july 2021. *World Weather Attribution*.

Loomis, J. (2011). What's to know about hypothetical bias in stated preference valuation studies? *Journal of Economic Surveys*, *25*(2), 363–370.

MCM. (2022).

MunichRE. (2021). Retrieved March 8, 2022, from https://www.munichre.com/en/company/media-relations/media-information-and-corporate-news/media-information/2022/natural-disaster-losses-2021.html

MunichRe. (2022).

Neuteboom, N., Kapetanios, G., Ritsema, F., Golec, P., & Ventouri, A. (2021). *How the consumer response to the covid-19 pandemic changes over time*.

Sampson, T., C.C.and Fewtrell, O'Loughlin, F., Pappenberger, F., Bater, P., Freer, J., & Cloke, H. (2014). The impact of uncertain precipitation data on insurance loss estimates using a flood catastrophe model. *Hydrol. Earth Syst. Sci.*, *18*, 2305–2324.

Slager, K., & Wagenaar, D. (2017). *Standaardmethode 2017 schade en slachtoffers als gevolg van overstromingen*.

Snow, J. (1855). On the mode of communication of cholera. *1*(7).

Sobolevsky, S., Massaro1, E., Bojic, I., Arias, J., & Ratti, C. (2017). Predicting regional economic indices using big data of individual bank card transactions.

Vrisou van Eck, N., & Kok, M. (2001). *Standaardmethode schade en slachtoffers ten gevolge van overstromingen.*

Wing, O., Pinter, N., Bates, P., & Kousky, C. (2020). New insights into us flood vulnerability revealed from flood insurance big data. *11*(1444).