KING'S BUSINESS SCHOOL

Data Analytics for Finance & Macro Research Centre

KING'S College LONDON

Consumption during the covid-19 pandemic: Lockdown or fear? Evidence from transaction data for the Netherlands

P. Golec, G. Kapetanios, N. Neuteboom, F. Ritsema, A. Ventouri

Working paper No. 2020/4 | September 2020 | ISSN 2516-5933

Consumption during the covid-19 pandemic: Lockdown or fear? Evidence from transaction data for the Netherlands

Pascal Golec* George Kapetanios† Nora Neuteboom‡
Feiko Ritsema§ Alexia Ventouri¶‖

September 2020

*ABN AMRO Bank; Gustav Mahlerlaan 10, 1082 PP Amsterdam; pascal.golec@nl.abnamro.com

†Kings Business School, Kings College London; george.kapetanios@kcl.ac.uk

‡Kings Business School, Kings College London and ABN AMRO bank; Gustav Mahlerlaan 10, 1082 PP Amsterdam; nora.neuteboom@nl.abnamro.com

§ABN AMRO Bank;Gustav Mahlerlaan 10, 1082 PP Amsterdam; feiko.ritsema@nl.abnamro.com

¶Kings Business School, Kings College London; alexia.ventouri@kcl.ac.uk

‖We are grateful to ABN AMRO bank for providing us with their data. All individual data used in this analysis has been anonymized. Special thanks to Chief Economist Sandra Phlippen, Sinem Hocioglu and Rolf van der Velden for their support and comments.

Abstract

Using transaction data from 2 million customers of ABN AMRO bank, this paper distinguishes the economic effects of voluntary responses to Covid-19 from those attributable to government lockdown measures. We compare municipalities that experienced large Covid-19 outbreaks with municipalities that had few or no cases, and find that the scale of the outbreak in a municipality has a strong negative effect on physical transactions by consumers, including for sectors that were forced to close during the lockdown. We show that these responses are largely rational, correlated with the intensity of the local outbreak rather than provoked by general perceptions of the outbreak. Our findings imply that fear of the virus stimulates self-isolation, which has a negative economic impact at the local level. Therefore, one potential path for long-term economic recovery is to diminish the effect of fear and restore consumer confidence by addressing the spread of the virus itself.

Keywords — Covid-19, coronavirus, transaction data, consumption, consumer spending, lockdown, policy implications, economic impact

1 Introduction

The outbreak of the coronavirus named Covid-19 has disrupted the Netherlands economy and is spreading globally. Countries around the globe have increasingly implemented strict public health measures to respond to the outbreak. These measures range from social distancing to complete lockdown, invariably constraining economic activities with serious ramifications. However, it is not clear to what extent the drop in economic activity during the lockdown was attributable to these restrictive policies or to sickness and public fear. When people are either infected or afraid of getting infected by the virus, they will choose to voluntarily stay at home. Understanding the different impacts that the two factors have on economic activity is key for policy making. Quantifying the extent to which the economic fallout was driven by fear or by lockdown measures should help policy makers to make better decisions in terms of lifting lockdown measures and dealing with a second Covid-19 crisis. GDP in the Netherlands fell by 1.5% from the fourth quarter of 2019 to the first quarter of 2020, primarily because of a reduction in consumer spending, which fell by 2.7% quarter on quarter, the largest decrease since 1987. Previous literature shows that transaction data closely resembles both domestic consumption data (Chetty et al., 2020 [17]) and retail spending (Carvalho et al., 2020 [15]). Using transaction data with geolocation tags for 299 municipalities in the Netherlands, this paper examines the impact of the fear factor on economic activity. Our empirical strategy separates the effects of voluntary distancing from those of lockdown measures by comparing municipalities that have seen large Covid-19 outbreaks with municipalities that had few to no Covid-19 cases.

The first case of Covid-19 in the Netherlands was identified on 27 February 2020 and the first Covid-related death occurred on 6 March. On 12 March, the Dutch government announced an "intelligent lockdown". Effectively, the next day, all events (concerts, sports) and all meetings with more than 100 people were forbidden. Bars, restaurants and other public places or venues where people gather had to close. The government encouraged people to work from home. All Dutch universities suspended physical teaching. In the days after, some measures where tightened. Schools and childcare facilities were closed, all those in contact-based roles had to stop performing their jobs unless it was possible to maintain a distance of 1.5 metres from clients at all times. The timing and severity of the measures were generally comparable to most of northern Europe (such as Germany, the Netherlands and Norway), but less restrictive than in southern Europe, where the virus spread more rapidly (such as Italy, France and Spain). Figure 1 shows the overall government response index for the Netherlands, measured by stringency, over time.[1]. It also shows the change in consumption year on year and the amount of new hospitalized Covid-19 cases. Part of the lockdown measures were relaxed on 11 May. The measures were eased further on 1 June, as the economy "reopened" (see the list of relaxed measures in Figure 1).

The Dutch government's strategy of nationwide lockdown resulted in two additional patterns that are relevant to our empirical estimation strategy: (i) While the lockdown policy was implemented nationwide on 12 March, both the incidence of the illness and its timing have varied substantially across the provinces. Therefore, while the lockdown measures induced homogeneous expenditure dynamics across the Netherlands, spatial heterogeneity remains because the Covid-19 virus may have had different effects in different municipalities across the country. In this study we single-out the effect of the local magnitude of the Covid-19 pandemic on private consumption. We find that the amount of new hospitalized Covid-19 cases in a municipality has a strong negative and statistically significant effect on volume of physical transactions by consumers. In other words, municipalities in the

[1] This government response index proxies the strictness of the lockdown policy over time by country. See: https://www.bsg.ox.ac.uk/research/research-projects/coronavirus-government-response-tracker

Netherlands that have seen a large Covid-19 outbreak have struggled more in economic terms than municipalities that seen few or no Covid-19 cases. This effect cannot be explained by lockdown measures, because in the Netherlands the lockdown was imposed for all municipalities on the same date and with the same stringency. Moreover, we show that the magnitude of the Covid-19 outbreak negatively impacts sectors that were forced to close during the lockdown. This suggests that people and businesses located in areas that were little affected by the virus were inclined to follow the lockdown rules more loosely, or that local government were less strict in enforcing the rules. Our findings imply that if people get the disease or know people who got infected, they adapt their spending behavior and spend less in physical transactions. We test this hypothesis by measuring the impact of magnitude of local Covid-19 cases on the amount of supermarket visits, by comparison to the baseline before the Covid-19 outbreak. We find that Covid-19 is strongly negatively related to the frequency of supermarket visits, indicating that people are more afraid to go outside when there is larger virus outbreak in their local area. We further investigate this effect of "fear" by looking at online grocery orders. Indeed, consumers living in badly affected areas order more groceries online, in order to prevent having to go outside the house to buy groceries in person. Digging further into this notion of fear, we investigate whether it is "rational" fear that people feel because of the intensity of the local outbreak, or whether it is "irrational" fear more related to the "perception" of how big the outbreak is, rather than the real local numbers. We look at how consumers respond to outbreaks in their province, compared to their local outbreaks, and test the relationship with supermarket visits. We find that supermarket visits are not related to the virus outbreak on a province level, but only on a local level, form which we draw the conclusion that it is mainly "rational" fear driving consumer behavior.

On the basis of these results, we find that despite the nationwide lockdown, people living in badly affected areas changed their economic behavior differently compared to people living in areas less affected by the virus. People in badly affected areas spent less offline, and therefore local economic activity has been hurt more in municipalities that have seen larger Covid-19 outbreaks. Accordingly, we believe that the only path to full economic recovery in the long run may be to diminish the effect of fear and restore consumer confidence by addressing the spread of the virus itself.

The remainder of this paper is organized as follows. In section 2, we discuss relevant theory and earlier empirical evidence from other countries. Section 3 describes our transaction data for the Dutch municipalities and section 4 deals with the model and methods that we use in this paper. In section 5 we discuss the results. Section 6 concludes the study.

2 Literature review

Our work builds on and contributes to a rapidly evolving literature on measuring the economic impacts of Covid-19 (see e.g. Brodeur et al. (2020) [12]). Our paper combines two strands of research: (i) literature that uses high-frequency transaction data to measure the economic impact of the Covid-19 outbreak, and (ii) literature that focuses directly on the impact of lockdown policies.

Recent papers use high-frequency transaction data, analogous to the data we assemble here, to analyze aggregate consumer spending (e.g., Sobolevsky et al. 2017 [31], Carvalho et al. 2020 [14], Baker et al. 2020 [10], Baker et al. 2020 [9], Chen, Qian, and Wen 2020 [16], Andersen et al. 2020 [5], Hacioglu, Känzig and Surico, 2020 [24]). These papers identify a number of important findings, such as: i) there are concentrated impacts on spending in certain industries (e.g. food and accommodation); ii) some social distancing is a result of voluntary choices rather than legislation and iii) there has been a severe drop in consumption caused the pandemic. In this paper we find an average drop in consumption of 20% during the lockdown period for the Netherlands. This is similar to the figure found by Andersen et al. 2020 [5] (27%), but less than those found by Hacioglu, Känzig and Surico, 2020 [24]) for the UK economy (a decline of 40% to 50%), Baker et al. 2020 [10] for the US economy (50%) and Carvalho et al. 2020 [14] for the Spanish economy (around 50%). Chen, Qian and Wen 2020 [16] find a 42% decline for the Chinese economy. However, these different figures are are not fully comparable because there are major differences in collection methods, time periods and whether the series are adjusted for inflation, holidays and other fluctuations.

The first paper that uses transaction data to addresses economic activity on a local level is that of Carvalho et al. (2020) [15]. They find that by the end of March 2020, the effect of the lockdown on expenditure growth was very similar across regions, irrespective of the number of confirmed cases. Using panel data, they confirm that neither GDP per capita nor the daily evolution of the regional incidence of the illness correlate robustly with the daily regional expenditure growth rate. They do, however, find a heterogeneity at a more granular level of disaggregation. When looking at the Madrid region in detail, Carvalho et al. (2020) [15] find that the fall of expenditure induced by the lockdown is larger in local areas within the region where the pandemic has caused more distress. Chetty et al. (2020) [17] examine the effect of executive lockdown orders on changing consumer spending in the US. Many states enacted stay-at-home orders and shutdowns of businesses in an effort to limit the spread of Covid-19 infections and later reopened their economies by removing these restrictions. They find that spending fell sharply in most states before formal state closures. Moreover, states' reopenings had little immediate impact on economic activity. They conclude that health concerns are the core driver of reductions in spending, rather than government-imposed restrictions. Chetty et al. (2020) [17] use US state-level data, which contain large heterogeneity across and within states in the precise form that the lockdown and the reopening took. They therefore argue that their estimates should be viewed as a broad assessment of the average impact of typical reopening efforts on aggregate economic activity and the authors defer from a more detailed analysis of how different types of reopenings affect different sectors.

Goolsbee and Syverson (2020) [22] examine to what extent the reduction in economic activity was due to government restrictions or to people's voluntary decision to stay at home. They perform a detailed analysis of customer visits on a county level with the use of cellular phone record data. By comparing consumer behavior within the same commuting zones but across counties with different government restrictions, they find that lockdown orders account for only a modest share of the decline in economic activity. They argue that although overall consumer traffic fell by 60%, legal restrictions explain a decline of only 7%. Our study differs from that of Goolsbee and Syverson (2020) [22] in that we use actual spending data instead of personal mobility data. What is more, Goolsbee and Syverson (2020) [22] implicitly assume that the number of visits corresponds to the amount of economic activity. Our data show that although people may visit stores less often during the Covid-19 crisis, they spend more time during every visit (see Figure 2). This implies that travel to stores may decline rapidly, but spending does not decline at the same rate. The transaction data on a municipality level allow us to investigate actual spending at the local level. Similar to Goolsbee and Syverson (2020) [22], Gupta et al. (2020) [23] study smart-device data that proxy mobility patterns. They find large declines in mobility in all states since the start of the pandemic, even those without major mitigation mandates. This indicates a substantial share of the fall in mobility was not induced by strong lockdown mandates such as "stay-at-home" orders. Their findings show that state level emergency declarations account for about 55% of the decline between the first week of March and the second week in April, with the remaining 45% of the decline attributable to secular trends that they interpret as the private (residual to policy) response to the pandemic.

In another study, Coibion, Gorodnichenko and Weber (2020) [19] find that the costs of lockdowns across counties in the US are very large. Specifically, they find that lockdowns are associated with a drop in overall spending equal to 31 log percent. They measure spending (and other economic indicators) using survey data from the Kilts Nielsen Consumer Panel. They also account for whether a county is in a lockdown. They argue that local public health authorities are likely to impose a lockdown as soon as a single case of a Covid-19 infection in a location is confirmed. Coibion, Gorodnichenko and Weber (2020) [19] measure the effect of lockdowns by comparing late and early adopters of lockdown policies, but this method may miss general equilibrium effects. They do not measure the impact of a health crisis during a lockdown and presume that all variation in economic activity during the crisis is caused by the lockdown.

Chen, Qian and Wen 2020 [16] find a similar result to that of Goolsbee and Syverson (2020) [22] for China. They study the drop in card and QR scanner transactions through UnionPay. They also find that the effect on consumption is stronger in cities that have had more Covid-19 cases. More specifically, they argue that in the 20 cities that received the highest inflow of Wuhan residents (the epicentre of the Covid-19 outbreak), consumption decreased by 12% more than in other cities in their sample. For cities reporting zero cases (as of late March), the decrease in offline consumption was 13% less than for cities with positive Covid-19 cases in the same time period. They find that a doubling of the number of infected cases at the city-level was associated with a 2.8% greater decrease in offline consumption. In addition, they find that fears regarding local hospital capacity and the first local Covid-19 related death drive consumption down further. They find that when the health situation worsened, consumption declined further, even for daily necessities. They conclude that management of the public health crisis is crucial for reinvigorating the economy. As fear of the virus spreads, (physical) consumption decreases. This hypothesis is supported by a recent study by Altig et al. [2], who constructed a newspaper- and Twitter-based sentiment measure that is closely aligned with the perceptions of households. The sentiment measure peaked during the Covid-19 outbreak as news and social media were becoming more negative in tone.

Other papers study the effects of Covid-19 and the imposition of lockdown on different aspects of the economy. Aum, Lee and Shin [7] perform a regional analysis of the effect of the Covid-19 outbreak on unemployment in Korea. Korea did not implement a lockdown, like may Western countries, but instead relied on testing and contact tracing. They find that a one per thousand increase in infections caused a 2 to 3% drop in local employment. Non-causal estimates of this coefficient from the US and UK, which implemented large-scale lockdowns, range from 5 to 6% , suggesting that at most half of the job losses in the US and UK can be attributed to lockdowns. The authors conclude that the primary culprit in the Covid-19 recession is Covid-19 itself, rather than lockdowns, and so the lifting of lockdowns around the world may lead to only modest recoveries unless infection rates also fall. Hiroyasu and Todo (2020) [25] quantifies the economic effect of a possible lockdown in Tokyo, Japan. Applying an agent-based model they find that when Tokyo is locked down for a month, the indirect effect on other regions would be twice as large as the direct effect on Tokyo, leading to a total production loss of 5.3% of Japan's annual GDP. Moser and Yared (2020) [29] study the optimal lockdown policy and argue that governments are faced with the choice to limit the extent of future lockdowns in order to support more optimistic investor expectations and credibility. Others, such as Alvarez, Argente and Lippi (2020) [3], study the optimal lockdown policy using the SIR (Susceptible, Infected and Recovered) epidemiology model and a linear economy to formalize the planner's dynamic control problem. They find that the optimal policy prescribes a severe lockdown beginning two weeks after the outbreak, covering 60% of the population after a month, and gradually withdrawn to cover 20% of the population after 3 months. They also argue that absence of testing increases the economic cost of the lockdown.

Our research is most related to the study by Chen, Qian and Wen (2020) [16]. We do, however, examine different geographical levels, as Chen, Qian and Wen 2020 [16] focus specifically on cities, whereas we consider municipalities. Also, China has had a more heterogeneous lockdown policy across different cities, because some where forced into full quarantine (such as Wuhan) whereas other cities faced milder lockdowns. Also the imposition dates of the lockdowns differ between cities in China. The lockdown policy in the Netherlands, which was consistent among regions in terms of stringency and imposition date, offers an empirical advantage for distinguishing the effect of Covid-19 from that of the lockdown policy.

3 Data

We use unique transaction data from ABN AMRO cardholders, the Stringency Index (which measures the strictness of lockdown measures) provided by the Blavatnik School of Government and the amount of hospitalized Covid-19 cases provided the Dutch National Institute for Public Health and the Environment (RIVM). In the following section we describe each dataset in more detail. We show how we transform the raw transaction data and explain how we acquire geolocation data on a municipality level. We end this section by exploring the descriptive stats and plots of the variables that we use for our estimation.

Transaction data We use transaction data from ABN AMRO cardholders. ABN AMRO is the third biggest bank in the Netherlands and has approximately 18% of the total market share in the country [34]). As a consumer bank, ABN AMRO has around 3.1 million unique account holders. This covers around 22% of the total adult (18 +) population. ABN AMRO is a broad retail bank present in all parts of the country and catering to all types of customers. Our data are therefore largely representative of the adult population of The Netherlands in terms of gender, age and income. Collectively, ABN AMRO account holders spend over 65 million euros on a daily basis, with an average transaction size of 23 euro. On average, over the sample period, our dataset comprises 2,745,651 physical pin transactions a day and around 344,753 online transactions a day. We acquire Point of Sale (PoS) data from pin terminals documenting every transaction. The data for each transaction include a timestamp, the amount in euros, the corresponding account number, the counter party description and the counterparty account number. For every ABN AMRO account holder we have some additional meta-data, such as gender, address and date of birth. We have used a labeling function based on keywords in the transaction description to identify the category the transaction belongs to. This is mostly done manually, whereby transaction descriptions where deciphered by analysts in order to determine their correct category. So, for example, if the description contained the word IKEA, it was added to the category "house fitment furniture stores". In this paper, we use (1) the aggregate physical debit card spending of account holders within municipalities, (2) the aggregate physical debit card spending of account holders on groceries within municipalities and (3) the aggregate online spending on groceries within municipalities by means of iDEAL (sepa) transactions.[2]

[2]Note that our dataset does not include transactions made by foreign (non-Dutch) tourist in the Netherlands. The total economic contribution from foreign tourists in the Netherlands is relatively small; tourism represents around 4.4 percent of GDP according the Statistical Bureau. Around 61.5 percent of this figure is contributed by domestic tourism and 38.5 percent by foreign tourism, giving a total contribution of foreign tourists to Dutch GDP of around 1.7 percent.

Our transaction data only incorporate accounts held by individuals and households. We exclude corporate accounts (SMEs) by excluding transactions backed by a debit card issued to a corporation as a “company card”. The purpose of this is to ensure that we make a correct analysis of domestic consumption that is not biased by corporate expenditure. [3] We don’t have access to credit card transaction data. In the Netherlands credit cards are not widely used as a payment option. This is confirmed by Jonker et al. (2017) [27], who argue that debit cards are the dominant payment method in the Netherlands. Moreover, the number of local Covid-19 cases should not meaningfully affect the ratio between debit and credit card usage compared to other parts of the country, and therefore the omission of credit card data should not bias our analysis. Our dataset offers several advantages for studying consumption. First, relative to survey-based consumption datasets, such as the consumer sentiment survey of the Central Bureau of Statistics (CBS), our sample has a much larger coverage of the Dutch economy. Moreover, survey data measure expectations, but do not measure actual private consumption in the way transaction data do.

Stringency index Throughout the Netherlands, reopening from the lockdown was a a gradual process in which certain industries and types of businesses opened before others. We make use of the stringency index for the Netherlands based on the Oxford Covid-19 government response tracker. This tracker collects information on several different common policy responses that governments have taken to respond to the pandemic, using 17 indicators such as school closures and travel restrictions. The data from the 17 indicators are aggregated into a set of four common indices, reporting a number between 1 and 100

[3] Note that this is a drawback of the Carvalho et al. (2020) [15] data. The study cannot distinguish the identity of the buyer in each transaction and therefore the data represent a mixture of final consumption expenditure by households and corporate firms’ intermediate input purchases.

to reflect the level of government action. Figure 1 depicts the stringency index for the Netherlands[4].

Covid-19 Data We report the number of new hospitalized Covid-19 cases each day, for each municipality, using publicly available data from the National Institute for Public Health and the Environment, [5]. National Institute for Public Health and the Environment provides the cumulative number of individuals who have tested positive for Covid-19, cumulative amount of deaths and cumulative amount of hospitalized cases on a daily basis by municipality. We use the latter in our research because testing capacity was very limited in the beginning of the crisis and biased towards healthcare workers (and hence big cities that have hospitals) and the amount of deaths is a lagging variable (people are often infected for many days/weeks before they pass away). For the amount of hospitalized cases we report two daily series by municipality: a seven-day moving average of new daily totals and a seven-day moving average of the cumulative total to the given date. Figure 3 depicts the aggregate cumulative and new amounts of Covid-19 hospitalized cases in the Netherlands. The figure shows that the peak in new Covid-19 cases was on 27 March 2020. Figure 1 shows the year on year change in transactions, the stringency index and the new Covid-19 hospitalized cases. From the plot it is seen that the initial spike in new Covid-19 cases corresponds with a severe drop in transactions. Moreover, as new Covid-19 cases fall, the average spending increases again.

[4] https://www.bsg.ox.ac.uk/research/research-projects/coronavirus-government-response-tracker

[5] https://www.rivm.nl/

3.1 Data transformations

We analyze year on year growth rates of the aggregate amount of spending by municipality on a particular date. In order to control for seasonality trends in expenditure we proceed as follows: We pair every day following 2 January, 2020 with its equivalent day in the equivalent week of the previous year:

$$\Delta Expenditure_i = \frac{Expenditure_{it} - Expenditure_{i\,t-365}}{Expenditure_{i\,t-365}} \tag{1}$$

where $Expenditure_{it}$ is *total transactions* $\times$ *average amount in euros* on day t for municipality i. In other words, in order to filter out daily seasonality, we compare Fridays in 2020 with Fridays in 2019, i.e. the same day of the week and almost exactly the same place in the month. In addition to adjusting for seasonal patterns, we also manually adjust for calendar effects. These are events and holidays related to the calendar that usually have an irregularly recurring pattern. We compared specific holidays with the same holiday in the previous year, despite the fact that it may have fallen on a different week- or weekend day. Moreover, we deleted data for 29 February 2020 because this date only occurs in a leap year, as a additional day. By adjusting for these patterns, data can be more effectively compared over time. [6] The data depict large periodic fluctuations across days. We address such fluctuations through aggregation, e.g. reporting 7-day moving averages to smooth daily fluctuations.

Figure 4 shows the 7-day moving average aggregate expenditure for the Netherlands in the period January until June. This figure includes all transactions by consumers; offline,

[6] This method does not fully take into account the fact that spikes in spending due to pay days may fall on different weekdays in different years. However, similar to what Andersen et al. [6] find for Denmark, paydays in the Netherlands are not all fixed on one particular day and vary across occupations and employers.

online and cash withdrawals. In line with Carvalho et al. (2020) [15], we find a pronounced increase in expenditure prior to the national lockdown on 13 March, but immediately after the announcement a large and sustained reduction in expenditure. To calculate the impact of the imposed lockdown, we take the period 13 March until 1 June, when the lockdown was relieved (see table 1 for an overview of when certain lockdown measures were relieved). Our data show that consumer spending has declined by an average of about 19 percent year on year since the government unveiled its measures on 12 March.[7] This calculation omits the 'panic-buying' days of 12 and 13 March. The largest drop in expenditure is seen on Sunday 22 March, when total spending decreased by 35 percent year on year. Figure 4 also shows that the transaction data already start to see improvements before the lockdown was relieved in the latter half of May and the beginning of June.

[7] Transaction data are very volatile, therefore we use the following methodology to calculate our summary measure of the effect of the lockdown on consumer expenditure. We look at the average spending in the period 10 January to 27 February. We have chosen this specific period because it ends before the first coronavirus incident in the Netherlands, which was reported on 27 February, and commences after the holiday effect of December should have ceased to affect the data. Google Mobility data (see section https://www.google.com/covid19/mobility/data_documentation.html?hl=en) confirms that out of the ordinary patterns in persons' movements to retail, workplace and grocery/pharmacy stores only started to appear from 9 March onwards. Therefore we can safely conclude that the period 10 January until 27 February can be considered as "normal times". We correct the spending during the lockdown for the difference in average growth compared to the period 10 January until 27 February by using the following equation:

$$Exp_{lockdown} = (avg\ exp \mid t \in lockdown) - (avg\ exp \mid t \in pre - lockdown) \tag{2}$$

where $Exp_{lockdown}$ is the average decline in expenditure year on year during the lockdown period measured for period t = 12 March until 1 June, $avg\ exp \mid t \in lockdown$ is the measured decline in expenditure during the lockdown period measured for period t = 12 March until 1 June and $avg\ exp \mid t \in pre-lockdown$ is the measured increase in expenditure for the period t = 10 January until 27 February.

3.2 Geolocation

Geolocation data for pin terminals are not readily available. However, for a more specific analyses of consumer spending behavior in different regions, it is important to know where point-of-sales (POSs) are located. Therefore we use transaction-weighted density-based clustering to predict POS geolocations (see Appendix A).

ABN Amro has bought a training set of physical locations of pin terminals in 2015. This contains around 718 million transactions made by 2 million customers at 2.1 million points of sale. This data serves as the ground-truth for geolocations. For every account-holder we have the zipcode of their registered home address. These zipcodes are merged with an external dataset featuring all Dutch zipcodes and their latitudes and longitudes[8] in order to obtain geolocations (expressed as latitude and longitude). Customers' home locations often have a clear relationship to payment point locations, with most payment points being situated in a dense cluster around consumer home location points. The purpose is thus to identify one cluster in these points, and to classify other transactions as outliers. Then using only the main cluster points, we proceed by determining the payment point location. The final implementation consists of several steps, of which clustering is the first, followed by determining the location and confidence using weighted transactions and finally correcting low-confidence predictions. The full methodology is described in Appendix A. This methodology offers the advantage of predicting the correct province with 95 percent accuracy, the correct municipality with 86 percent accuracy and the correct city with 67 percent accuracy.

For every municipality we calculate the average expenditure by

[8] http://geonames.org

$$\frac{\sum_i^T expenditure}{T}$$

where T is the amount of days in our dataset. In order the ensure that we have adequate coverage of transactions within every municipality (i), we delete the lowest 5 percentiles of municipalities based on the ABN AMRO clients to total inhabitants ratio. We also delete municipalities for which we have incomplete data, which leaves us with a dataset of 299 municipalities out of 355 in the country. For the online iDEAL (sepa) transactions, we locate transactions at the home address of the ABN AMRO client, which is part of the meta data that is provided for every ABN AMRO account-holder.

For the geographical distribution, figure 5 shows that ABN AMRO clientbase is over-represented in the province Noord-Holland, and under represented in provinces such as Gelderland, Overijssel, Limburg and Drenthe. Large cities in the Netherlands are mainly concentrated in Noord-Holland, Zuid-Holland and Noord-Brabant. This reflects the history of ABN AMRO, which was formed by a merger between two banks which were based in Noord-Holland and Zuid-Holland and were focused on industry and services, not on agriculture. Therefore, we have an over-representation of clients from urban areas, while having relatively fewer clients from rural areas.

3.3 Descriptive analysis

We find a clear break in the trend in total pin transactions as the Covid-19 virus starts to emerge and the Netherlands imposes a lockdown. Figure 4 shows that the year on year change in total spending shows a clear trend break at the start of the Covid-19 crisis. Similar to findings by Carvalho et al. (2020) [15], we find evidence of hoarding just before the imposition of the lockdown on March 12. After that period, we see a sharp contraction in total spending that starts to recover from May onwards. Figure 6 shows the development of different consumer categories over time. Notice that after the end of the lockdown on June 1, sectors such as entertainment and travel still show substantial contraction, as many businesses in these sectors were not allowed to fully reopen. Also sectors that largely reopened at the start of June subsequently struggled to lure back customers. Spending on clothing, footwear and wellness have not rebounded to pre-coronavirus levels, and the pattern is similar for hospitality venues. Only retail remained buoyant, especially consumption in supermarkets and online groceries shopping. Figure 7 and figure 8 show average spending on offline and online groceries, respectively. From the graphs it is seen that both categories show increased spending from the middle of March onwards. In particular, online grocery shopping saw a large jump, at the end of April, at the peak, people spend around four times as much on online groceries as a year before. However, as figure 8 shows, as the Covid-19 crisis starts the wane during May and June, the total increase in online grocery shopping also starts to fall. Figure 9 shows the year on year changes in total consumption, offline supermarket consumption and online supermarket consumption in one plot. This figure clearly shows that whereas overall consumption dropped significantly after mid-March, supermarket and online supermarket expenditure show an inverse trend and performed remarkably well.

The transaction data also show the sharp drop in expenditure in sectors that were forced to close during the nationwide lockdown. Figure 10 shows the year on year change in transactions in sectors that could remain open and sectors that were forced to close, aggregated over all municipalities in the Netherlands (see Table 2 for a list of spending categories in closed sectors). Figure 10 also shows the stringency index for the Netherlands. It is striking to see that with every step the government took to lift the lockdown measures (proxied by the change in the stringency index) the total transactions in closed sectors increases. In our research, we disentangle the effect of the lockdown from the effect that is attributable to other trends, i.e. the private response to the epidemic (residual to the policy response). Figure 10 shows that the total consumption is very closely correlated with consumption in the sectors that were still open during the lockdown measures (correlation of 98%). Indeed, as the lockdown in the Netherlands was relatively mild, the sum of the sectors that had to close represented only a small proportion of total consumer expenditure. This strengthens our view that solely looking at the economic impact of the lockdown measures is not sufficient in order to asses the economic impact of the epidemic.

Figure 12 and figure 13 plot the year on year transactions, the inverse of new Covid-19 hospitalized cases and the stringency index proxy of the lockdown for two municipalities of relatively small population size, Veendam and Bernheze. Veendam had almost no Covid-19 cases, and although the transaction data show an initial spike just before the imposition of the lockdown and a small drop afterwards, it recovers almost immediately to normal levels of expenditure. This is contrary to the municipality of Bernheze, where the drop in transactions is very pronounced (50% at is lowest level), and moves alongside the number of new confirmed Covid-19 cases. This same conclusions hold for two big cities, Amsterdam and Groningen, with the first seeing a large Covid-19 outbreak and the latter experiencing a much smaller outbreak. Figure 15 and Figure 14 show that in Amsterdam the total transactions dropped to 50 percent, whereas in Groningen transactions bottomed out at around 35% below the prior year level. Notice that for Amsterdam the transactions had not returned to positive levels at the end of the lockdown.

We see a similar pattern in the amounts spent in supermarkets. We hypothesize that, when municipalities are harder hit by the spread of the virus, people are more afraid to go to physical stores. Figure 16 and Figure 17 show the total spending in physical stores on groceries for Veendam and Bernheze respectively. Again, these figures show that whereas Veendam saw an initial spike in supermarket consumption (hoarding), the amount spent in supermarkets quickly returned to near-regular levels, albeit slightly higher than before the Covid-19 outbreak. This is contrary to Bernheze, were the initial spike in supermarket consumption was followed by a steep drop, even becoming negative for a while. For Amsterdam, figure 19 shows that after the period of hoarding, supermarket consumption dropped by over 20%. Moreover, it has remained below average levels throughout the Covid-19 crisis. Groningen also experienced a steep drop during the start of the virus outbreak, but saw a slightly better recovery in supermarket consumption in the weeks thereafter.

Figure 20 shows a comparative bubble chart. The vertical axis shows the average year on year drop in transactions between March 14 and April 16 and the horizontal axis shows the logarithm of cumulative hospitalized Covid-19 cases up to April 16. The size of the bubble indicates the population size of the municipality. The trend-line shows a negative correlation between cumulative Covid-19 cases and growth in transactions. It also shows that in general larger municipalities (in terms of population size) have faced larger Covid-19 outbreaks and hence experienced a larger drop in transactions.

Figure 21 shows the year on year change in transactions between March 13 and June 1 (the end of the nationwide lockdown) by region. Blue indicates an increase in transactions and red indicates a decrease in transactions. Figure 22 shows a geographical map of the cumulative amount of hospitalized Covid-19 cases on June 1. Dark red indicates a large number of Covid-19 cases and light red indicates few Covid-19 cases. Municipalities that are coloured white have seen zero Covid-19 cases. Comparing the two maps reveals that municipalities that have seen larger Covid-19 outbreaks, mostly located in the centre and south of the Netherlands, also show a larger drop in transactions.

4 Model and research design

Throughout this paper, we use the following terms:

- *The dependent variable* is denoted Y_{it} for municipality i at time t.
- *The target variable of interest* is denoted X_{it} for municipality i at time t.
- *Control variables* are denoted X_i and X_t for the sets of control variables $(X_{i1}, x_{i2}, \ldots, X_{in})$ and $(X_{t1}, X_{t2}, \ldots, x_{tn})$, which are time-invariant and entity-invariant, respectively [9].

.

To measure whether the severity of the Covid-19 outbreak at the local (municipality) level has a significant impact on transactions, we use a fixed effects (FE) panel regression.

FE estimation is performed by *time demeaning* the data. Demeaning deals with unobservable factors because it takes out any component that is constant over time and entity. By assumption, this would be the entire effect of the unobservable variables (see Appendix B for an investigation into what sort of entity fixed effects we are capturing in the regression). We use the following equation:

$$Y_{it} = \alpha_i + \gamma_t + \beta_1 X_{it} + \varepsilon_{it}, \tag{3}$$

where Y_{it} is the observation for the ith cross-section unit at time t for $i = (1, 2, \ldots, N)$ and $t = (1, 2, \ldots, T)$. α_i denotes unobserved characteristics for each cross-sectional unit that don't vary over time; a $m \times 1$ vector of unobserved common effects. γ_t are unobserved characteristics for each time unit t that don't vary over entity; hence a $k \times 1$ vector of unobserved common effects. X_{it} is a $1 \times n$ (include constants) of observed independent

[9] These variables control for factors that could potentially impact Y_{it} and/or X_{it}.

variables, including our variable of interest, new hospitalized Covid-19 cases. ε_{it} are the individual-specific (idiosyncratic) errors assumed to be distributed independently of X_{it} and α_i. By including fixed effects α_i and γ_t, we are controlling for the average differences across municipalities in any unobservable predictors, which allows us to eliminate the omitted variable bias.[10]

In general, however, the unobserved factors α_i could be correlated with X_{it}. Pesaran (2006) [30] shows that by including the cross-sectional averages in the regression the differential effects of unobserved common factors are eliminated. This approach is favoured above the principal components approach brought forward by Coakley, Fuertes and Smith (2002) [18] because we want to avoid inconsistent results in the situation that the unobserved factors and the included regressors are correlated. Moreover, the approach by Pesaran [30] allows us to use ordinary least squares (OLS) when we specify an auxiliary regression where the observed regressors are augmented by cross-sectional (weighted) averages of the dependent variable (Y_{it}) and observable variables (X_{it} and possibly other control variables

[10] The random effects (RE) model is more appropriate when the entities in the sample can be thought of as having been randomly selected from the population, but a fixed effects model is more plausible when the entities in the sample effectively constitute the entire population. A RE model allows all unobserved effects to be relegated to the error term by specifying the model as

$$Y_{it} = \beta_1 X_{it} + v_{it}$$

where $v_i t = \omega_{it} + \varepsilon_{it}$. The RE technique doesn't estimate the fixed effects separately for each cross-sectional unit, and so gives fewer estimated parameters, increased degrees of freedom, and smaller standard errors. A critical assumption of the RE model is that the unobserved individual effect (ω_i) isn't correlated with the independent variable(s). For the research question at hand, we cannot eliminate the possibility that there are unobserved factors that are correlated with X_{it}. Moreover, although ε_{it} satisfies the classical linear regression model (CLRM) assumptions, the inclusion of ω_i in the composite error $v_{it} = \omega_i + \varepsilon_{it}$ results in a CLRM assumption violation. If you relegate the individual effects (ω_i) to the error term, you create a positive serial correlation in the composite error. As a result, a RE estimation requires feasible generalized least squares rather than ordinary least squares (OLS) to appropriately eliminate serial correlation in the error term and to produce the correct standard errors. Therefore, the FE specification is the more appropriate method for us to use.

X_i and X_t) (see also Kapetanios and Pesaran, 2005 [28]). We therefore run the regression in two stages:

First stage regressions:

$$Y_{it} = \beta_1 \mu_t^Y + \beta_2 \mu_t^X + e_{it}^y X_{it} = \beta_1 \mu_t^Y + \beta_2 \mu_t^X + e_{it}^x \tag{4}$$

where μ_{it}^Y and μ_{it}^X are the cross-sectional averages of Y_{it} and X_{it} respectively over time t. e_{it}^y and e_{it}^x capture the residuals of equation 4.

Second stage regression:

$$e_{it}^y = e_{it}^x + w_{it}, \tag{5}$$

Under the assumptions explained in Pesaran, 2006 [30], for any fixed m in α_i these residuals provide consistent estimates of ε_{it} in the multifactor model (3) and could be used as "observed data" to obtain estimates of the factors α_i. The factor estimates can then be used directly as (generated) regressors in regression equation 5. Effectively, in the second stage we try to explain the variance of Y_{it} with the variance of X_{it}, thereby eliminating all the other fixed effects.

To further investigate the effects of our selection of control variables, we also run a Panel OLS model including only time fixed effects and control variables that vary among municipalities (but not over time):

$$Y_{it} = \gamma_t + \beta_1 X_{it} + \varepsilon_t, \tag{6}$$

where $X_{it} = X_{it} + X_i$ and X_i is a set of control variables $(x_1, x_2, \ldots, x_n)$, e.g. $n \times 1$ vector of observed common effects. Notice that this is a less strict specification compared to equation (3) because we cannot account for entity fixed effects.

Moreover, to the test the importance of the lockdown measures on physical spending, we run a FE model that includes entity fixed effects and control variables that vary over time (but not over municipality):

$$Y_{it} = \alpha_i + \beta_1 X_{it} + \varepsilon_t, \tag{7}$$

where $X_{it} = X_{it} + X_t$ and X_t is a set of control variables $(x_1, x_2, \ldots, x_v)$, e.g. $v \times 1$ vector of observed common effects. As this specification does not include time fixed effects, this is also a less strict specification than equation (3).

5 Results

Table 3, column (1) shows the regression results for the fixed effects panel regression (3). The dependent variable (Y_{it}) is the year on year change in total volume of transactions by municipality i for time t. The key explanatory variable is the new amount of hospitalized Covid-19 cases (X_{it}) in municipality i for time t. The regression also includes both municipality and time fixed effects. We cluster standard errors at the municipality level (see Appendix B for an investigation of what the entity fixed effects in this regression are potentially capturing). We find a strong statistically significant negative coefficient on the explanatory variable (X_{it}). The results show that on average, each additional hospitalized Covid-19 case reduces year on year transactions by -2.79%. Given that in the Netherlands the stringency and timing of the lockdown were identical for all municipalities, the lockdown effect in the regression is captured by time fixed effects.[11]

Table 3, column (2) shows the results including the stringency index as a control variable (X_{1t}), which is a proxy of the strictness of the lockdown measure over time (see equation 7). This index is constant over municipality because the Netherlands faced a nationwide lockdown. The panel regression also includes municipality fixed effects, but no time fixed effects because the stringency variable is constant over municipality but differs over time. Both variables are significant; however, the coefficient on the number of new Covid-19 cases is substantially larger than the coefficient on the stringency index. This suggest that if the lockdown policy becomes one unit more strict on a scale of 0–100, transactions will drop by 0.05%. The beta coefficient of X_{it} suggests that one additional Covid-19 case will cause transactions to drop by 2.79%. These regression results show that economic activity is reduced by both the threat of Covid-19 itself as well as the government lockdown restrictions.

Table 3, column (3) shows the results including the stringency index, and two other high-frequency macro variables that are constant over municipality but differ over time (X_t in equation 7). The AEX Index is the Dutch stock market index, which proxies investors sentiment. EURIBOR stands for the Euro Interbank Offered Rate, the average rate that European banks issue loans to other banks. These variables proxy the growth of the economy and the level of inflation. The inclusion of both high-frequency macro variables into the fixed effects panel regression has little effect on the results. Table 3, column (3) shows that

[11]For all our regression, the value of R^2 is relatively low. This is contrary to findings of Chen, Qian and Wen (2020) [16] and Goolsbee and Syverson (2020) [22], who report a relative large R^2. This is mainly because of differences in specifications. Both of those studies use dummy variables to single out specific subgroups of panel data, and therefore they measure the effects between two different groups over time. In contrast, we look at the fit of all municipalities (299) over time. Moreover, they transform the data into logged values, whereas we prefer to use a linear model and to keep the amount of transformations to the data as minimal as possible.

the coefficients on target variable, the number of new hospitalized Covid-19 cases (X_{it}), the stringency index and the EURIBOR are statistically significant. Moreover, the inclusion of the two macro variables reduces some of the explanatory power of the stringency index, which reduces by more than half from -0.05 to -0.021. The coefficient on the AEX Index is not statistically significant. This is not surprising given that financial markets have performed remarkable well during the Covid-19 crisis. After a short dip in the beginning of March, the stock index recovered rapidly.

Table 4 shows the results of the two-stage regression (see equation (5)). This specification allows us to eliminate the differential effects of unobserved common factors (see section 4). The result is significant at the 1% level. The interpretation of the beta coefficient is not straightforward because it is the beta coefficient of the residual of the original regression. The sign of the coefficient is as anticipated; the number of hospitalized Covid-19 (X_{it}) cases has a negative effect on year on year average transactions. If we include time-variant control variables (X_t) in the first stage regression (equation 4) and estimate (X_{it}) in the second stage, the variable is still statistically significant and the beta coefficient does not change. This suggests that the set of control variables X_t has no correlation with the variable of interest (X_{it}). Moreover, if we add these control variables to the final regression (equation (5)) we find that all control variables are not statistically significant, whereas X_{it} remains highly significant, which implies that their variance does not have any explanatory power for the variance of Y_{it}.

5.1 Time lags

As we make use of high-frequency daily data, we also investigate the sensitivity of time lags (inter-temporal shifting). Figure 23 shows the beta coefficient on the explanatory variable (X_{it}), the new amount of hospitalized Covid-19 cases, shifted over time. We use the second-stage regression (the regression of the residuals) to show the evolution of the beta coefficient. The horizontal axis t indicates the time in days, the vertical axis shows the beta coefficient on (X_{it}). The figure shows that the beta coefficient has the highest negative explanatory power when $t = 0$. The beta coefficient becomes smaller as t shifts over time, and this is unequally distributed between t_+ and t_-. This makes intuitive sense, as for instance the use of t_{+15} would try to explain the drop in transaction data of t_{-15} days previously based on the hospitalized Covid-19 cases at t, where a causative relationship would not seem possible. Figure 24 shows the p-value of (X_{it}) in the second-stage residual regression. The horizontal axis t indicates the time in days. This figure shows that the explanatory variable (X_{it}) becomes non-significant when the time is shifted towards t_{-45}. This makes sense as the amount of Covid-19 cases was zero for all municipalities before the first Covid-19 case was confirmed on 27 February 2020. In general, these results show that our variable of interest is robust to shifts in time.

5.2 The effect of the lockdown

One may argue that the difference between municipalities may not be the result of the number of local Covid-19 cases, but rather the result of "disobedience" or the lack of "enforcement" of the lockdown rules. In our research we take the assumption that the lockdown measures—as they were implemented nationwide—are similar across municipalities. While this should be true theoretically, in practice some municipalities may be stricter in following the national government rules than others. For example, Amsterdam might have been much more strict in terms of checking whether people actually obey the rules compared to, for example, a small municipality in the North. Figure 11 shows that the variance in closed sectors between big cities is very small and largely constant over time. We test this further by regressing the stringency index on the total transactions in closed sectors. Under this category we take all sectors that were subject to restrictions and were either forced to fully close or continue with very limited capacity, for instance, restaurants permitted only to provide takeaway goods (see Table 2 for a list of consumer categories that we include in the subset of closed sectors). As expected (see Table 6, column (1)) the stringency level is an extremely good predictor of the drop in transactions in the closed sector, with an R^2 of 0.45[12]. Table 6, column (2) shows the same regression but including the Covid-19 variable (X_{it}). For completeness, we include the residual regressions in Table 7, which include stringency in the first-stage regression of equation (4), and also show a statistically significant result for the coefficient on hospitalized Covid-19 cases in equation (5). It is interesting to see that the variable is still significant at the 1% confidence level, which indicates that the small differences between municipalities in supermarket visits are the result of the severity of the local Covid-19 outbreak. Table 8 validates this conclusion. Here, we regress the Covid-19 variable on the variance between municipalities in transactions in the closed sectors. Again the coefficient on the Covid-19 variable is statistically significant, albeit at the 5% level. Intuitively, these results make sense. People are less likely to follow lockdown rules if they don't see the benefit or purpose of these rules. Similarly, local municipalities may be less strict in enforcing the national rules if they have few or no Covid-19 cases in their community. This also tells us that even for the sectors that may feel their loss in revenues is fully caused by the lockdown measures, it does also matter how severely their local community has been hit by the virus. The effect is actually quite strong: With every additional hospitalised Covid-19 case in the municipality, closed sectors see their pin trans-

[12]Note that this panel regression only includes entity effects because stringency is constant for each entity i over time.

actions drop (year on year) by 3.2%. If the lockdown measures had been made tougher, for example during April to the level that Italy faced at the time (a stringency level of 93.5 out of 100), this would have resulted in an additional 10% decrease in pin transactions (year on year) in the closed sectors.

5.3 The effect of fear

Our results show a strong negative sensitivity of consumption to within-municipality changes in the outbreak's severity. Part of this finding may be attributed to the customers' fear and anxiety about going outside, based on the intensity of the local Covid-19 outbreak. We test this hypothesis further by looking at the frequency of visits in supermarkets. Given that supermarkets are generally crowded places and a potential source of Covid-19 contamination, we hypothesize that consumers would want to avoid these places more if they live in a municipality that experiences a larger Covid-19 outbreak. We choose supermarket visits specifically because they are not impacted by the lockdown measures, because supermarkets, convenience stores and other "vital" food retailers were allowed to stay open. Moreover, supermarkets are likely to be situated close to costumers, generally within the same municipality, which minimizes any spatial effects that may occur. CBS data show that, on average, there is a supermarket around 1 kilometre away from every Dutch household.

Aggregate data show that spending in supermarkets was not negatively impacted by the Covid-19 crisis. Figure 7 shows that despite the initial spike in expenditure on groceries, spending in this category remained high throughout the lockdown period. From March 13 until the beginning of July, the expenditure in this category increased, on average, by 25% compared to the previous year. This is partly explained by the substitution effect, as consumers spend less on restaurants, bars and catering at work (see also figure 6). However, the frequency with which consumers visited supermarkets did decline and hence, on average, consumers spent more during each visit. Figure 25 shows that the average amount spent in supermarkets (in euros) was substantially higher throughout the Covid-19 crisis.[13] Before the Covid-19 crisis, consumers spent around 18.5 euros at every supermarket visit, and this number increased to around 22.5 euros during the Covid-19 crisis.

In Table 9 we investigate whether Covid-19 had a statistically significant impact on the number of visits to supermarkets. Similar to Table 3, we perform a panel regression with fixed effects (see equations (3), (6) and (7) and a two-stage panel regression with the residuals (see equation (5)). The fixed effects regressions in Table 9 show that the amount of new hospitalized Covid-19 cases has a strong and statistically significant negative effect on the number of supermarket visits. This result shows that every additional hospitalized Covid-19 case reduces year on year supermarket visits by -1.09%, on average. This beta coefficient becomes slightly larger if we add other control variables and eliminate time effects. Table 9 shows that adding control variables does not change the significance nor the sign of the coefficient on the Covid-19 variable X_{it}. Also, if we regress the two-stage residuals regression (see Table 10) the results do not change. This finding indicates that whereas total spending in supermarkets was not negatively effected by the Covid-19 outbreak, people living in municipalities that have seen a bigger spread of the virus did visit supermarkets less often. These results suggest that the day-to-day change in the within-municipality Covid-19 intensity affects daily supermarket visits, beyond the effect of physical constrains imposed by the government's lockdown measures. On the basis of these results, we can reasonably argue that people living in badly affected areas were more afraid of going outside than people that live in areas that were little affected by the virus.

[13]This also holds for online grocery shopping. Figure 26 shows that the average amount people spent on online groceries increased from around 50 euros to an average of 67 euros in the period form mid-March until July.

We further examine the role of "fear" by looking at the online ordering of groceries. Aside from people going to supermarkets less often, buying more at each visit (as Figure 25 shows), they could also choose to fully self-isolate by ordering groceries online. We hypothesise that if people are afraid to go outside because they live in an area that has seen a large virus outbreak, they will be more inclined to order groceries online and get them delivered to their home address. Whereas online ordering of groceries was a relatively small market in the Netherlands[14] before the Covid-19 outbreak, the country has seen a enormous increase during the epidemic crisis. Figures 6 and 8 show the increase in online grocery shopping during the crisis, which on average more than doubled compared to the period before the crisis.

In Table 11, column (1) we regress the measure of Covid-19 (X_{it}) on the year on year change of the volume[15] in online grocery shopping (Y_{it}) for every municipality (see equation (3)). We find a strong statistically significant effect of the Covid-19 variable (X_{it}). The coefficient is positive, which supports our hypothesis that people living in areas that see larger Covid-19 outbreaks order more online in order to avoid going outside. Table 11, columns (2) and (3), show the regression results from equation (7), including the control variables X_t (stringency, AEX and EURIBOR). Again, these variables seem to capture some of the fixed effects that change over time but stay constant over municipality Table 12 shows the regression results with the residuals (equation (5)). Again, the coefficient on Covid-19 is statistically significant and positively correlated with the dependent variable online grocery expenditure. These results further strengthen our view that fear of getting

[14]Exact figures are not available for the Netherlands, but on the basis of survey data it is estimated that online grocery orders amount to around 5% of the total expenditure on groceries, see https://insights.abnamro.nl/2020/04/waar-halen-we-ons-eten-in-de-anderhalve-meter-samenleving/

[15]volume is measured as the amount of times that people ordered online

the virus, measured by the Covid-19 variable, is important in people's decision to change their behavior and switch from offline to online grocery shopping.

One can wonder how rational this fear is. Research has shown that people are very much influenced by the media (Akerlof and Shiller, 2010 [20]; Baker et al., 2014 [8]; Beaudry and Portier, 2004 [11]; Soroka, 2006 [32]; Tetlock, 2007 [33]). So the fear of getting the virus may not necessarily be caused by the absolute size of the local outbreak, but rather by the "perception" of the magnitude of the local outbreak. People living in the province of Brabant, which was known as the "epicentre" of the Netherlands' outbreak, may simply change their behavior because they feel the virus to be present, whereas the municipality of Brabant in which they live may have seen few or no Covid-19 cases.

In order to disentangle these two effects, we use an identification strategy that considers the amount of Covid-19 cases in the province (denoted by p) where municipality i is located. The Netherlands consists of 12 provinces. We include a variable that distinguishes the Covid-19 outbreak in municipality i from the cross-sectional average over the province ($\overline{X}_{pt}$), weighed by population size (expressed as a figure for hospitalized cases per 1000 people).We do this specifically for the regression on supermarket visits, because these should not be impacted by the amount of Covid-19 cases on a province level. Other transactions, such as spending on clothing, books or hardware could potentially be affected by the scale of the virus outbreak at the province level because some of these shops are not available within a particular municipality, causing people to travel to different municipalities (that may have seen different levels of Covid-19 cases). In the latter case, the fear of getting the virus may actually be rational, despite a low number of municipality cases. As previously noted, the average travel distance for supermarkets is small, which eliminates potential spatial issues. Table 13 shows the results of the fixed effects regression. Whereas the Covid-19 variable is still significant at the 5% level, the variable that measures the difference to the amount of cases on the province level is not statistically significant. We also perform the two-stage residual regression. In the first-stage (equation (4)) we perform three regressions, on Y_{it}, X_{it}^1 and X_{it}^2, with Covid-19 cases by municipality denoted by X_{it}^1 and the Covid-19 outbreak in municipality i minus the cross-sectional average for the province ($\overline{X}_{pt}$) denoted by X_{it}^2. In equation (5) we perform the regression $e_{it}^y = e_{it}^{x1} + e_{it}^{x2} + w_{it}$. Table 14 shows the results and again, the Covid-19 cases by municipality is the only statistically significant variable.

From the regression in Table 13 we can conclude that it is not the average level of the Covid-19 outbreak in the province that make people change their physical consumer behavior, but rather this change is related to the amount of Covid-19 cases in close proximity to them. Although further research is needed to look deeper into how people experience anxiety due to virus outbreaks and how this may result in changing spending behaviors, our results show that the fear of the virus and consequent changes in behavior may actually be largely rational.

6 Conclusion and discussion

From our regression results described in section , we find that the amount of new hospitalized Covid-19 cases in a municipality has a strong and statistically significant negative effect on the change in physical transactions by consumers. In other words, municipalities that have seen a large Covid-19 outbreak have struggled more in economic terms than municipalities that seen few or no Covid-19 cases. This finding holds even for sectors that were subject to restrictions due to the lockdown measures. We find that spending on necessities, proxied by supermarket expenditure, were barely hit by the Covid-19 crisis. However, the number of supermarket visits was reduced, especially in those municipalities that saw a large Covid-19 spread. This suggests that people living in badly affected areas altered their economic behavior differently to people living in little affected areas. This finding was further confirmed by the evidence that people were ordering more groceries online if their area was badly affected. Moreover, we find that it is not the perception of the outbreak that results into people changing their behavior, but rather the real magnitude of the local outbreak. Therefore, we conclude that the only path to full economic recovery in the long run is to diminish the effect of fear and restore consumer confidence by addressing the virus spread itself.

Our research has several policy implications. Our findings show that to in order to minimize the economic impact of the Covid-19 pandemic, governments should respond quickly to contain the spread of the disease. This can be done by stay-at-home policies, but also by implementing policies such as social distancing and contact tracing, if they prove to be more effective. [16] Many other authors have focused on researching the effectiveness of different policies on containing the virus (see, for example, Fang, Yiting and Marshare (2020) [30]; Hou et al. (2020) [26] and Alvarez, Argente and Lippi (2020) [4]). Our research shows that consumption in sectors that faced restrictions during the nationwide lockdown differs in accordance with the magnitude of the local virus outbreak.

Therefore, local lockdowns may prove to be more effective than nationwide lockdowns, because people tend to be more strict in following the rules and/or local enforcement is tougher in areas where there is a severe local outbreak. Our research could potentially also have implications for the distribution of national funds to local governments. Our research indicates that some municipalities have been hit substantially harder by the drop

[16] Aum, Lee and Shin (2020) [7] provide a good overview of the economic effect of intensive testing and contact tracing in South Korea.

in transactions than others, which could be used as an argument to reallocate budgets according to the extent to which a municipality has been hit by the Covid-19 crisis.

Measuring physical transactions is not a perfect substitute for economic activity. The drop in physical transactions can be partly offset by an increase in online transactions, and figure 6 shows that online grocery shopping skyrocketed during the Covid-19 crisis. Also, our regression results (section 6) show that the incidence of Covid-19 has a large positive correlation with online grocery spending. Further research could investigate to what extent online consumption substitutes the reduction in offline consumption. Also relevant is this respect is the reallocation from spending at local businesses to online retail businesses. The local community does not necessarily profit from an increase in online consumption at the expense of offline consumption, because many of the online retailers may not be located within the same municipality.

Moreover, it would be interesting to further research the impact of loss of income on consumer expenditure. Do municipalities that have seen larger virus outbreaks also experience a higher increase in unemployment? And if so, does this affect local economic activity? Initially, spending fell in the early stages of the pandemic, primarily because of health concerns rather than a loss of current or expected income. This is confirmed by Cajner et al. (2020) [13] and Ganong, Noel and Vavra (2020) [21], who found that income losses were relatively modest because relatively few high-income individuals lost their jobs and lower-income households who experienced job losses had their loss in income offset by unemployment benefits. This also holds for the Netherlands, where Covid-19 crisis measures provided financial support to around 2.9 million workers (out of a labour force of about 9 million) [17] However, when companies lose costumers and revenue, they may be forced to

[17] https://www.tweedekamer.nl/debat_en_vergadering/plenaire_vergaderingen/details?date=24-06-2020

cut costs, including labour costs. Such a negative income shock could further deteriorate the economic situation at the local level. Further research into the dynamics of the labour market and consumer spending at the local level in relation to the Covid-19 outbreak is needed the understand its full economic impact.

A final further avenue of research could be to investigate how "fear" works in the setting of a virus outbreak. We have seen that at the start of the lockdown, some spending categories dropped to almost 0%, whereas before the lockdown was eased, consumer spending in these categories had already started to recover. Does the imposition of the lockdown have a positive or negative signalling function? Which spending categories are most impacted by fear? And more importantly, does fear wane over time? I.e. do consumers get used to the virus threat and return to business as usual? If this is indeed the case, we may see a less detrimental economic effect of the virus if and when a second wave arrives.

Table 1: Phased relief of the lockdown restrictions in the Netherlands

11 May 2020
Partial relaxation of restrictions
Outdoor sports for people aged 19 or older allowed
Most close-contact roles, such as hairdressers and physiotherapists, were allowed to reopen
Primary schools reopened

01 June 2020
Reopening of economy
Secondary education started up
Hospitality venues, cinemas and theaters open with 1.5 metre social distancing and no more than 100 people
People were allowed to meet again
Public transport returns to normal timetable (only essential travel)

01 July 2020
Smart restart
All indoor and outdoor sports permitted
No limit for number of visitors to events; reservations and pre-entry health checks are mandatory
Non-essential use of public transport allowed
All close-contact roles back to work

Table 2: Categories of consumer spending that were forced to close during the nationwide lockdown in the Netherlands

Closed sectors during the Covid-19 crisis
Restaurants
Wellness
Casinos
Cinemas
Concert halls
Museums
Theme parks
Bars
Dining
Fastfood
Hairdressers
Solariums
Spas
Other entertainment

Table 3: Dependent Variable: Percentage change in transactions year on year

Regressor	**Column 1**	**Column 2**	**Column 3**
Covid-19 cases	-1.2094^{***}	-2.796^{***}	-2.8521^{***}
	(0.0413)	(0.0470)	(0.0470)
Stringency		-0.0473^{***}	-0.0217^{***}
		(0.0010)	(0.0019)
AEX			-0.0011
			(0.0011)
EURIBOR			-20.402^{***}
			(0.7977)
N	299	299	299
T	386	386	386
FE	Entity and Time	Entity	Entity
R^2	0.0074	0.0598	0.0654

(Standard errors in parentheses)
*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$

Table 4: Dependent Variable: Residuals for percentage change in transactions year on year

Regressor	**Column 1**
Residual Covid-19 cases	-1.6004^{***}
	(0.0499)
N	299
T	386
FE	NO
R^2	0.0089

(Standard errors in parentheses)
*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$

Table 5: Dependent Variable: Percentage change in transactions year on year

Regressor	**Column 1**
Covid-19 cases	-1.3986^{***}
	(0.0463)
Proximity to GP	-0.1127
	(0.1343)
Proximity to hospital	0.5934^{***}
	(0.0058)
Proximity to supermarket	-1.3413^{***}
	(0.1467)
Proximity to restaurant	-2.8558^{***}
	(0.1270)
Proximity to culture	-0.0215^{***}
	(0.0077)
Proximity to public transport hub	-0.1582^{***}
	(0.0057)
N	299
T	386
FE	Time
R^2	0.1521

(Standard errors in parentheses)
*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$

Table 6: Dependent Variable: Percentage change in closed sectors year on year

Regressor	**Column 1**	**Column 2**
Stringency	-0.7836^{***}	-0.7662^{***}
	(0.0025)	(0.0026)
Covid-19 cases		-3.2599^{***}
		(0.1160)
N	299	299
T	386	386
FE	Entity	Entity
R^2	0.4585	0.4622

(Standard errors in parentheses)
*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$

Table 7: Dependent Variable: Residuals for percentage change in closed sectors year on year

Regressor	**Column 1**
Residual Covid-19 cases	-1.2147^{***}
	(0.1317)
N	299
T	386
FE	NO
R^2	0.0007

(Standard errors in parentheses)
*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$

Table 8: Dependent Variable: Variance between municipalities in percentage change in transactions in closed sectors year on year

Regressor	**Column 1**
Covid-19 cases	-0.2525^{**}
	(0.1038)
N	299
T	386
FE	Entity and Time
R^2	0.001

(Standard errors in parentheses)
*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$

Table 9: Dependent Variable: Percentage change in supermarket visits year on year

Regressor	**Column 1**	**Column 2**	**Column 3**
Covid-19 cases	-1.087^{***}	-1.4899^{***}	-1.6376^{***}
	(0.1179)	(0.1133)	(0.1135)
Stringency		-0.0912^{***}	-0.1237^{***}
		(0.0025)	(0.0047)
AEX			-0.0150^{***}
			(0.0026)
EURIBOR			-36.220^{***}
			(1.9265)
N	299	299	299
T	386	386	386
FE	Entity and Time	Entity	Entity
R^2	0.0007	0.0115	0.0153

(Standard errors in parentheses)
*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$

Table 10: Dependent Variable: Residuals for percentage change in supermarket visits year on year

Regressor	**Column 1**
Residual Covid-19 cases	-1.3375^{***}
	(0.1553)
N	299
T	386
FE	NO
R^2	0.0006

(Standard errors in parentheses)
*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$

Table 11: Dependent Variable: Percentage change in online groceries year on year

Regressor	**Column 1**	**Column 2**	**Column 3**
Covid-19 cases	5.0367***	−5.5437***	−3.1878***
	(0.8726)	(0.8345)	(0.8340)
Stringency		2.0882***	2.2304***
		(0.0185)	(0.0346)
AEX			0.4589***
			(0.0192)
EURIBOR			240.54***
			(14.152)
N	299	299	299
T	386	386	386
FE	Entity and Time	Entity	Entity
R^2	0.0003	0.1024	0.1107

(Standard errors in parentheses)
*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$

Table 12: Dependent Variable: Residuals for percentage change in online grocery volume year on year

Regressor	**Column 1**
Residual Covid-19 cases	16.175***
	(1.1305)
N	299
T	386
FE	NO
R^2	0.0018

(Standard errors in parentheses)
*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$

Table 13: Dependent Variable: Percentage change in supermarket visits year on year

Regressor	**Column 1**
Covid-19 cases	-39.272^{**}
	(16.108)
Province Covid-19 cases minus municipality Covid-19 cases	-21.218
	(17.433)
N	299
T	386
FE	Entity and time
R^2	0.0001

(Standard errors in parentheses)
*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$

Table 14: Dependent Variable: Residuals for percentage change in supermarket visits year on year

Regressor	**Column 1**
Covid-19 cases	-74.345^{***}
	(21.345)
Province Covid-19 cases minus municipality Covid-19 cases	-37.192
	(23.146)
N	299
T	386
FE	NO
R^2	0.0002

(Standard errors in parentheses)
*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$

Figures

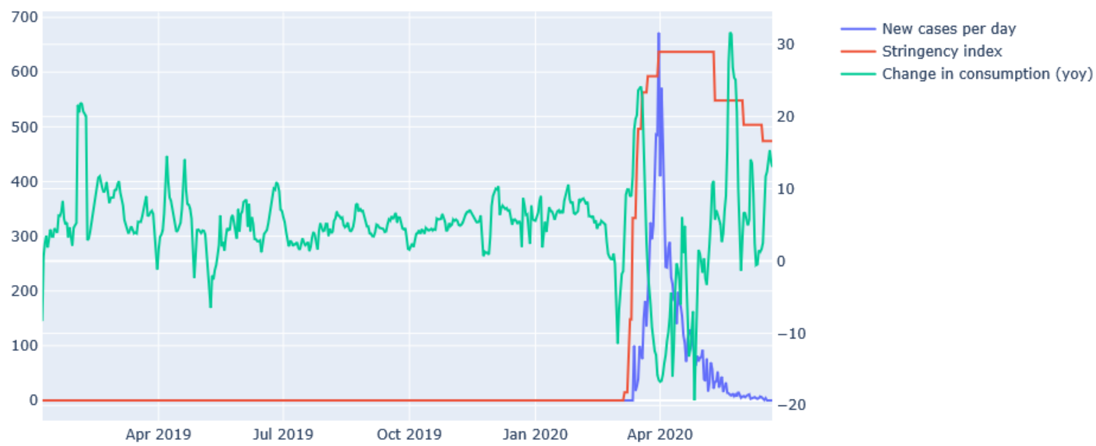

Figure 1: Year on year transactions, new Covid-19 hospitalized cases and the proxy of the lockdown measured by the stringency index in the Netherlands (stringency scaled * 8)

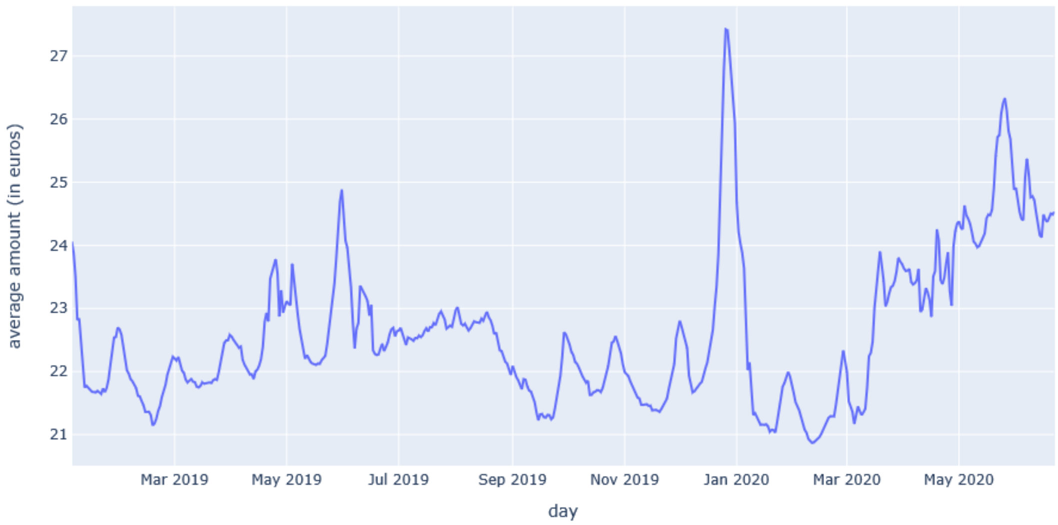

Figure 2: Average amount per transaction in euros spent with debit cards in the Netherlands, 7-days rolling average

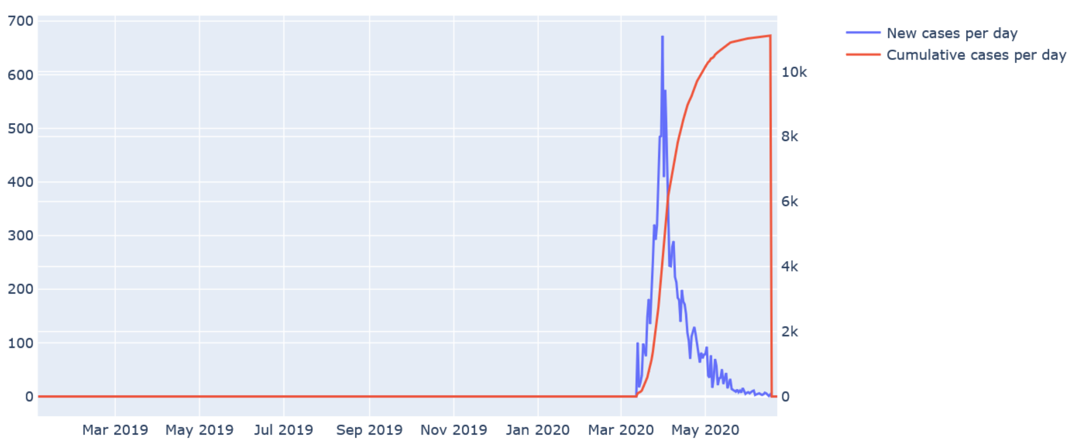

Figure 3: Cumulative and new amounts of hospitalized Covid-19 cases in the Netherlands

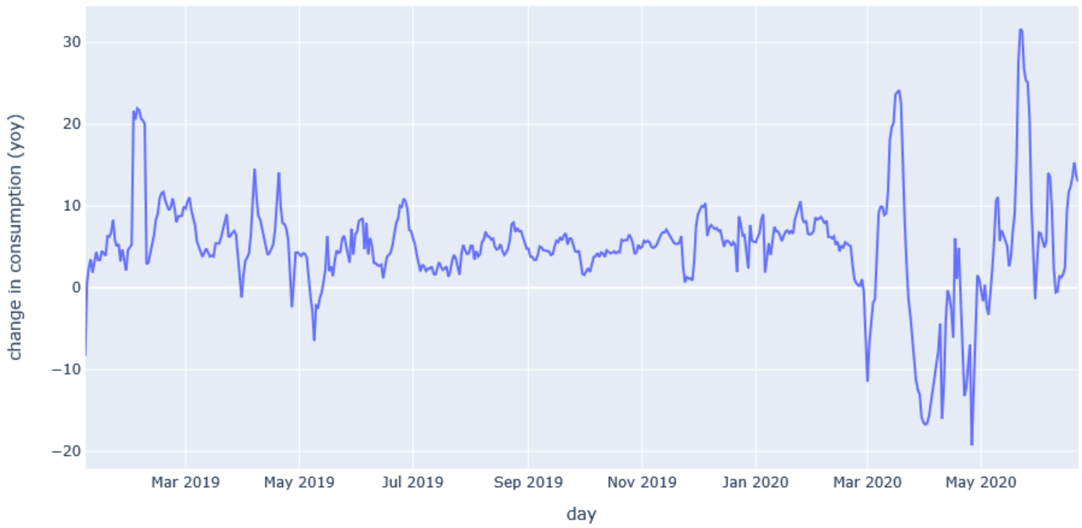

Figure 4: Average year on year change in total expenditure, 7-days rolling average

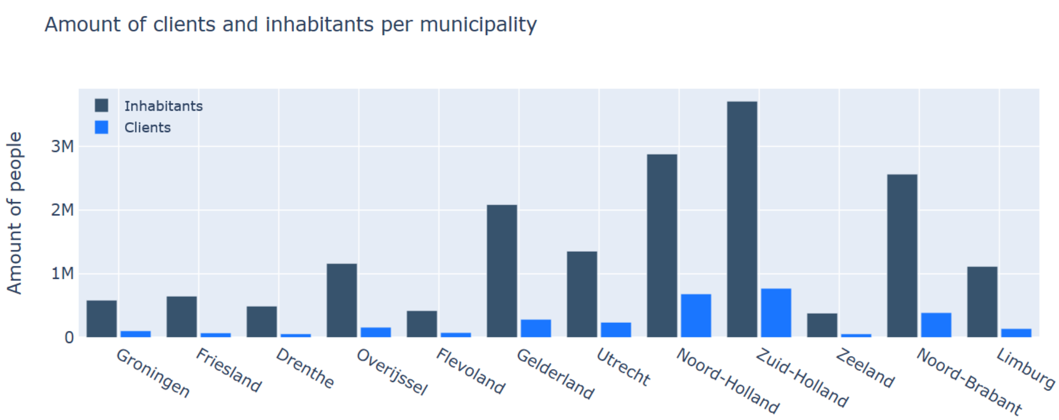

Figure 5: Amount of ABN AMRO clients in different provinces of the Netherlands

Lockdown imposed (16 march) | First relief (18-may) | Restart (1-jun)

Consumer expenditure by category, % change yoy, weekly basis	2 march	9 march	16 march	23 march	30 march	6-apr	13-apr	20-apr	27-apr	4-may	11-may	18-may	25-may	1-jun	8-jun	15-jun	22-jun
Groceries (online)	60,7	135,6	93,0	156,3	181,3	287,3	311,1	212,2	351,4	380,8	254,8	226,5	243,1	226,1	236,4	197,5	221,9
Electronics	3,4	12,7	40,5	29,3	29,2	38,5	32,0	36,2	35,3	23,7	40,2	27,6	29,3	43,9	17,4	20,1	23,2
Groceries (offline)	9,4	45,2	38,3	25,4	24,9	27,8	17,8	26,4	22,9	23,4	27,2	28,1	31,3	21,7	22,8	17,7	18,5
House and garden maintenance	7,9	14,2	5,2	-8,6	-4,7	19,2	6,8	19,6	26,3	31,2	37,3	45,3	14,2	48,1	23,4	24,7	30,8
Others*	24,4	24,8	26,2	28,9	36,9	41,0	27,4	37,0	38,0	40,4	49,5	29,5	49,0	48,1	38,9	23,5	35,2
Wellness	15,5	25,5	0,6	-21,9	-28,0	-30,8	-31,2	-29,3	-26,1	-19,1	3,8	-14,8	6,3	0,5	0,2	1,6	-2,3
Transport	3,5	-9,0	-34,7	-42,9	-43,3	-41,2	-41,0	-33,6	-32,2	-26,4	-23,1	-27,7	-8,7	-19,0	-10,1	-10,7	-5,3
Appearance***	-6,4	-22,4	-61,7	-55,9	-58,4	-53,7	-54,2	-49,4	-38,3	-24,5	-5,1	-17,6	-6,6	-6,8	-9,6	-17,2	-9,0
Restaurants & bars ***	7,3	-12,1	-61,0	-53,6	-49,5	-49,0	-48,1	-40,9	-47,6	-43,9	-39,2	-37,6	-39,3	-1,3	-11,9	-4,3	2,3
Travel	-7,1	-20,8	-66,6	-75,3	-78,6	-81,4	-81,7	-82,6	-83,9	-77,0	-73,5	-71,3	-71,5	-65,7	-63,6	-50,8	-41,5
Entertainment	7,0	-55,1	-91,7	-93,3	-92,3	-89,7	-91,4	-94,1	-92,6	-92,1	-87,7	-86,5	-83,5	-75,5	-80,2	-72,1	-58,0

Total amount of pintransactions, online payments and cash withdrawals, ex. Creditcard payments; nominal data not correct for seasonality, holidays and or other trends

* amongst others; books, toys, flowers; *** online and offline

Figure 6: Transactions in different consumer categories over time in the Netherlands before, during and after the lockdown. Transactions are measured by the total online (iDEAL) and offline (debit card) transactions

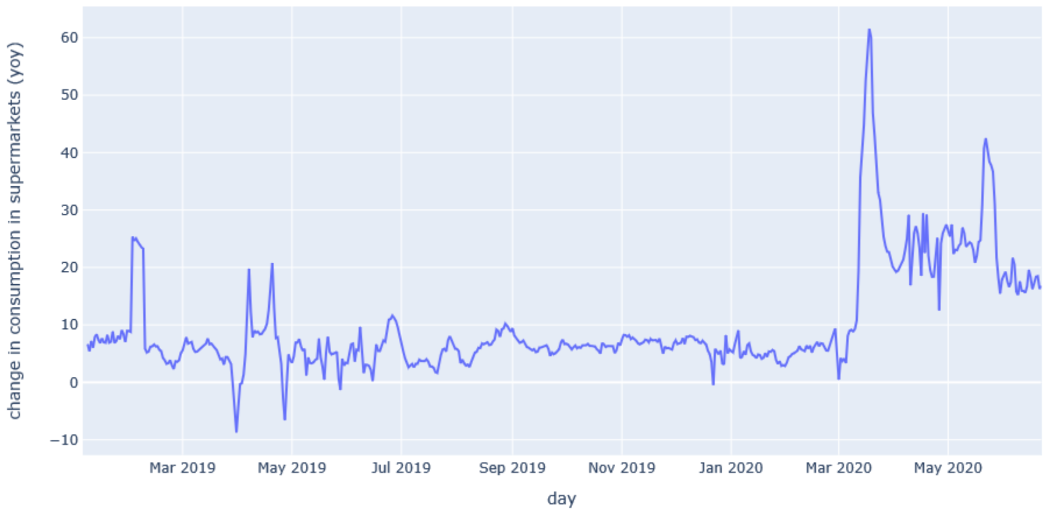

Figure 7: Average year on year change in supermarket expenditure, 7-days rolling average

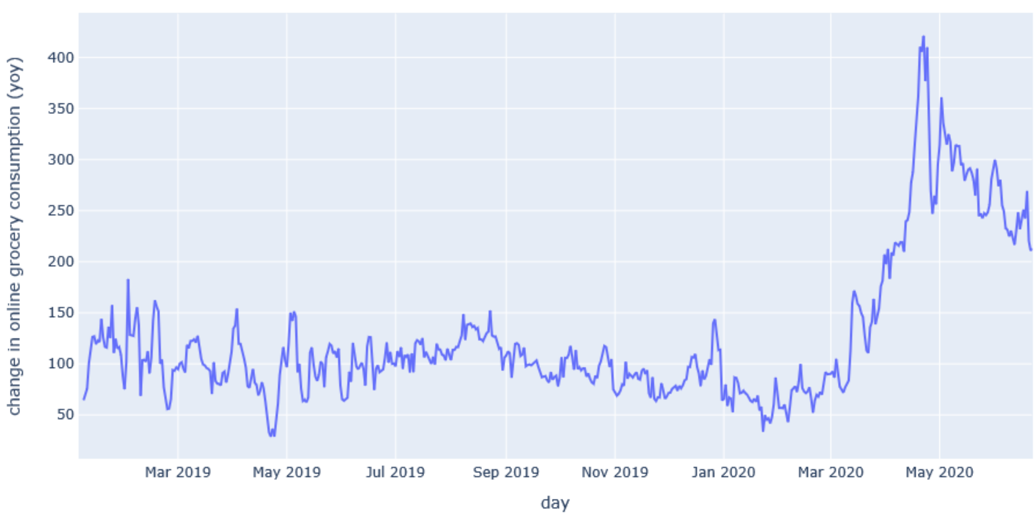

Figure 8: Average year on year change in online grocery expenditure, 7-days rolling average

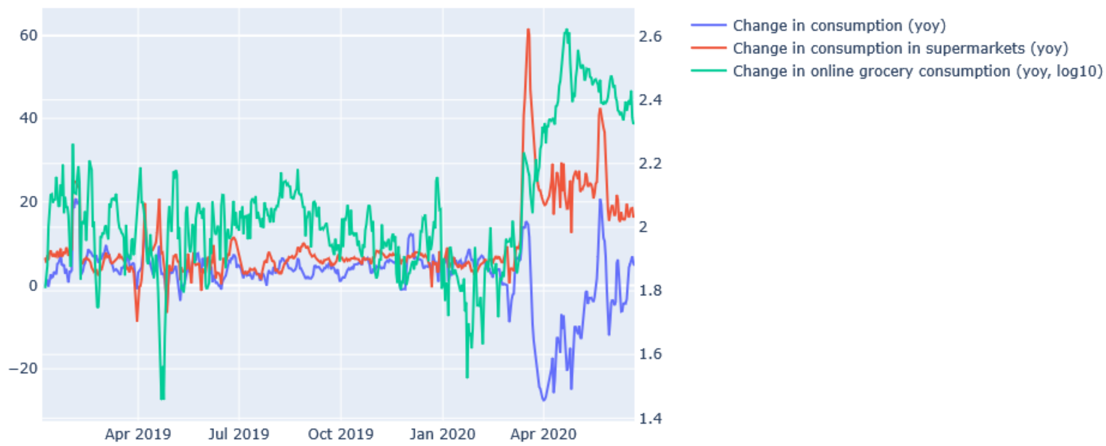

Figure 9: Year on year change in different expenditure types, 7-days rolling average

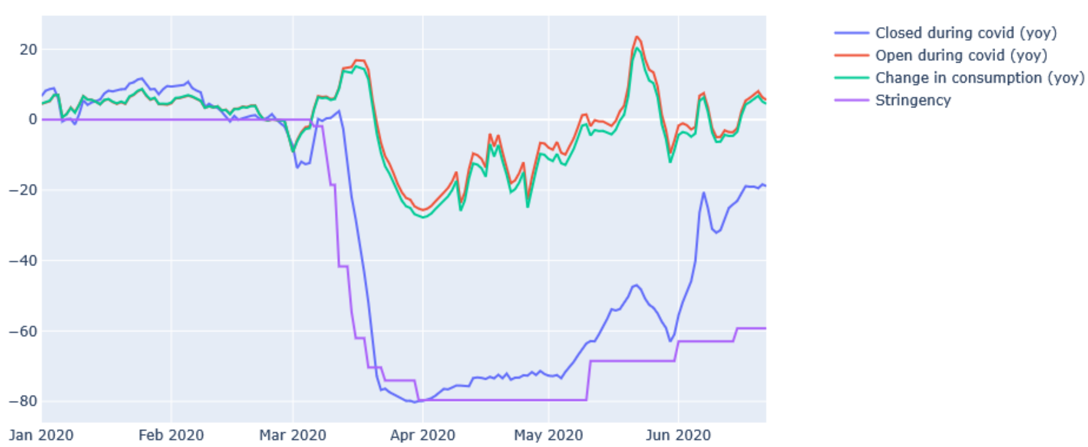

Figure 10: Year on year change in transactions for total consumption, consumption in sectors that were forced to close during the lockdown and sectors that could remain open. Stringency (inverse) is a proxy of the lockdown.

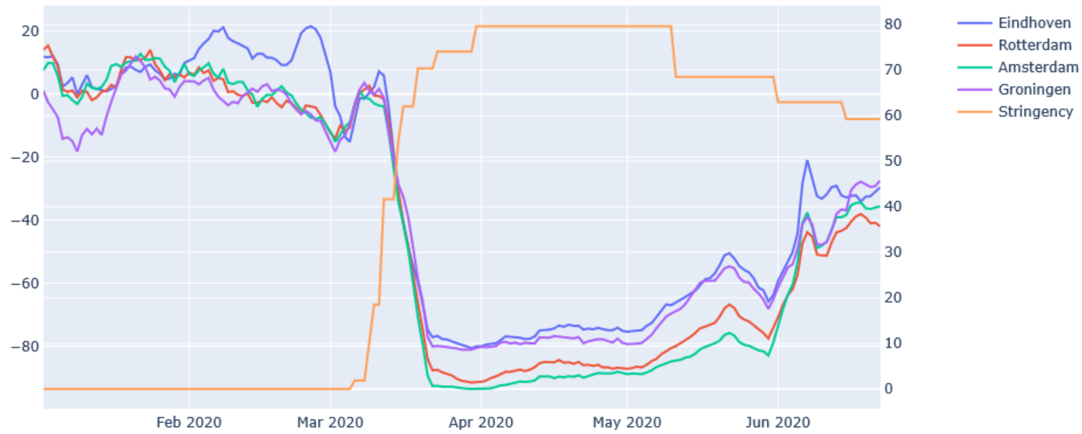

Figure 11: Year on year change in transactions in sectors that were forced to close during the lockdown in four large cities. Stringency is a proxy of the lockdown.

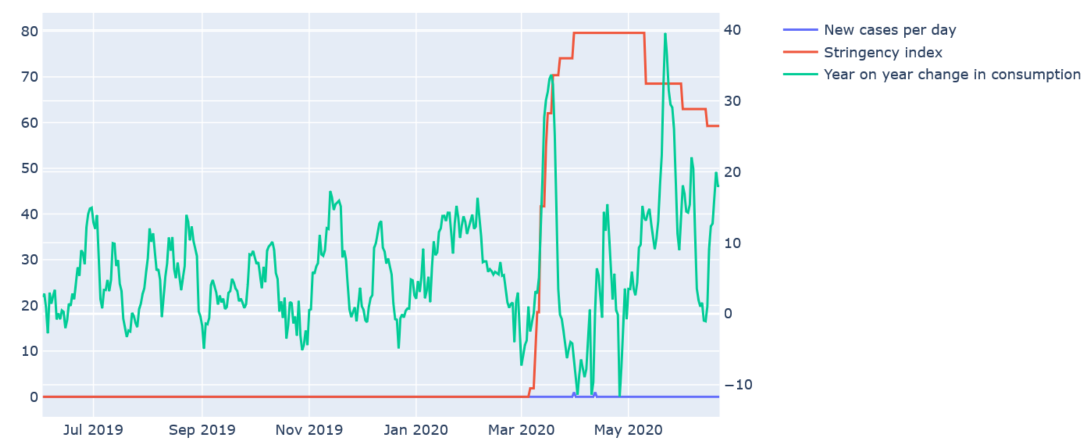

Figure 12: Year on year transactions, new hospitalized Covid-19 cases and the proxy of the lockdown measured by the stringency index for Veendam

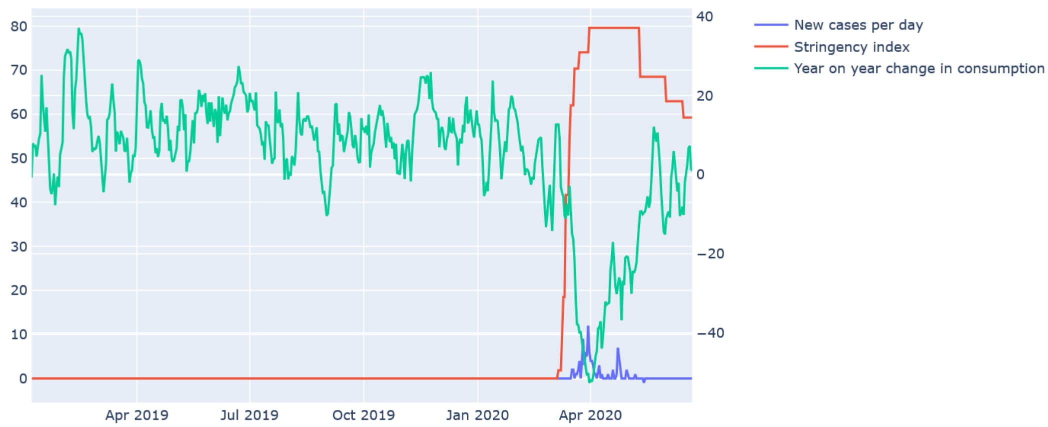

Figure 13: Year on year change in transactions, new hospitalized Covid-19 cases and the proxy of the lockdown measured by the stringency index for Bernheze

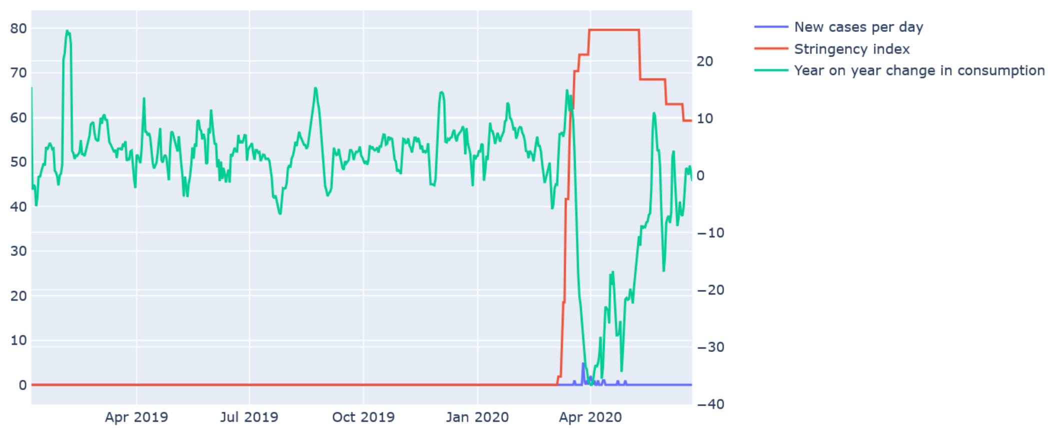

Figure 14: Year on year transactions, new hospitalized Covid-19 cases and the proxy of the lockdown measured by the stringency index for Groningen

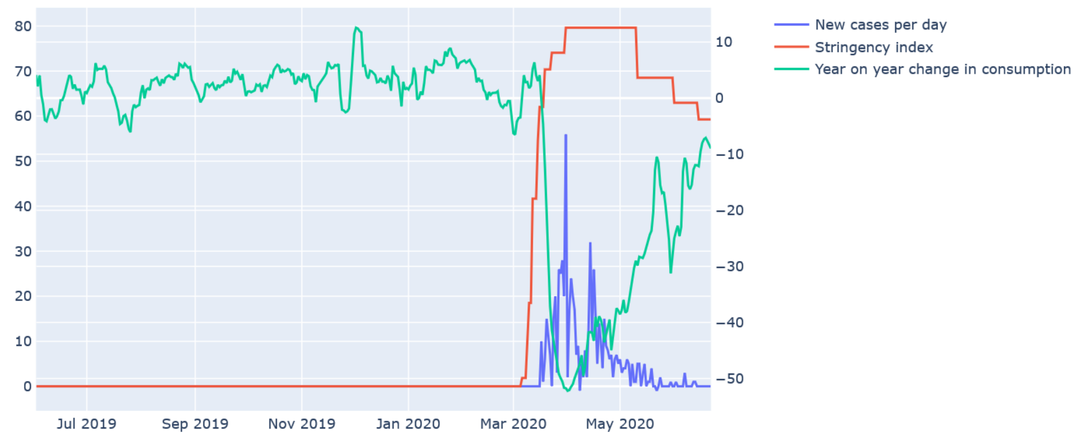

Figure 15: Year on year transactions, new hospitalized Covid-19 cases and the proxy of the lockdown measured by the stringency index for Amsterdam

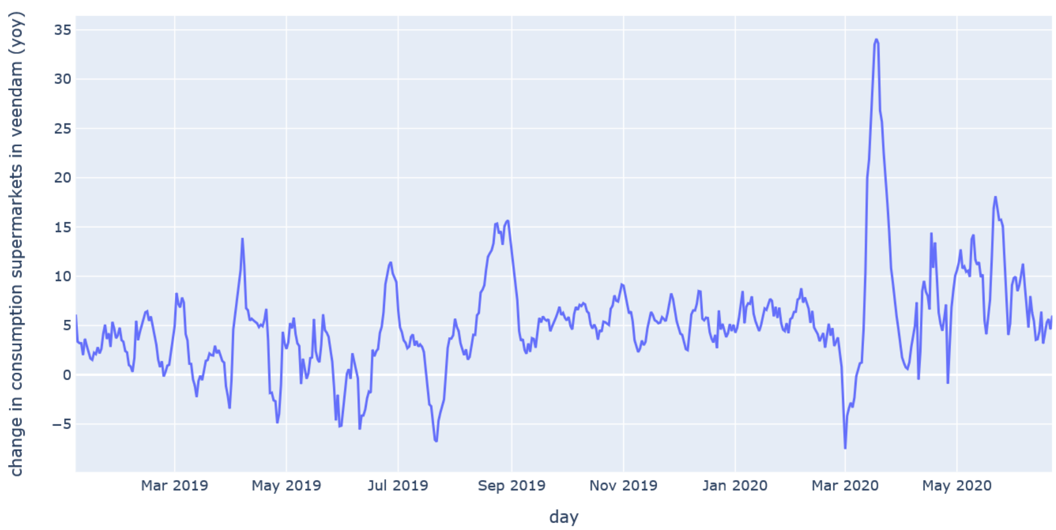

Figure 16: Veendam year on year change in supermarket expenditure, 7-days rolling average

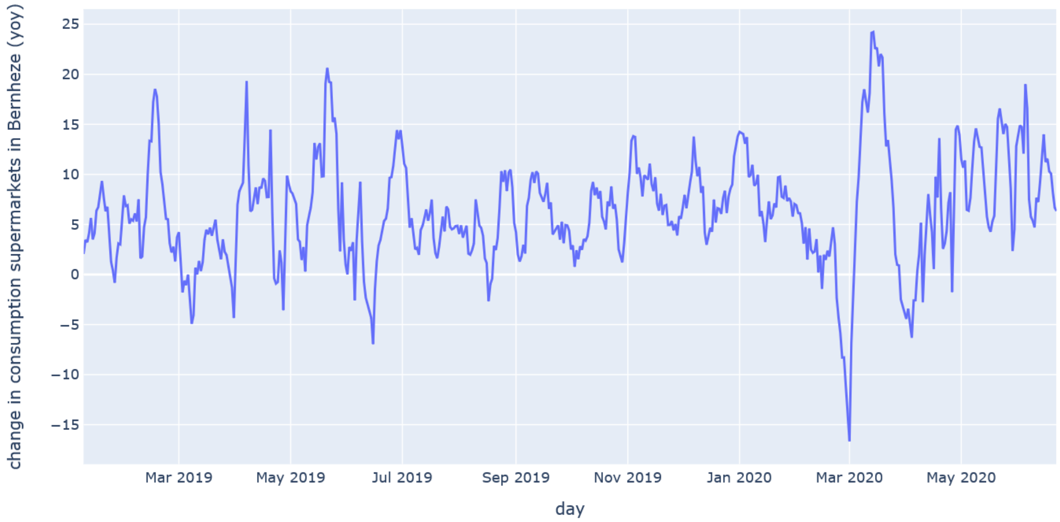

Figure 17: Bernheze year on year change in supermarket expenditure, 7-days rolling average

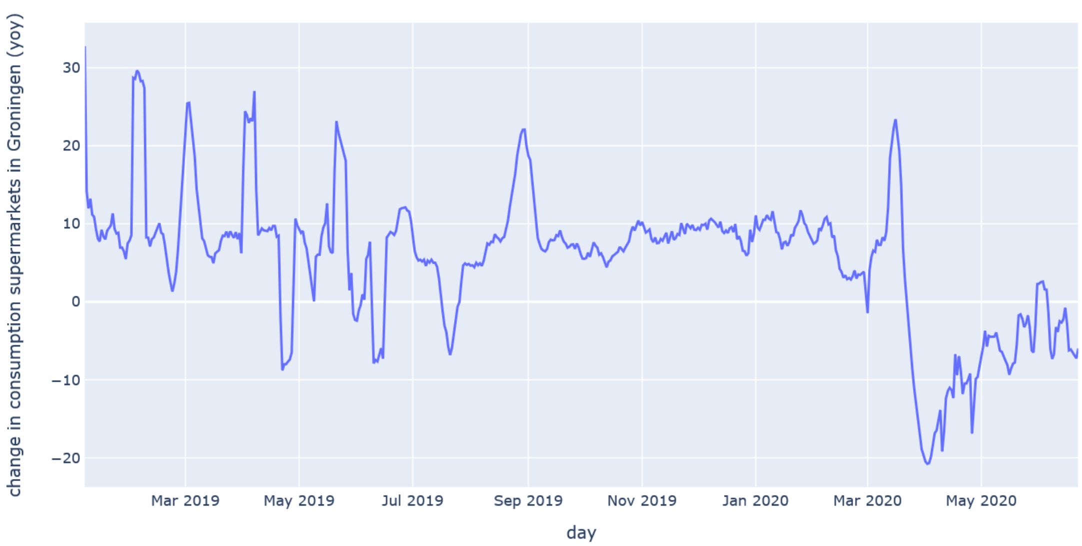

Figure 18: Groningen year on year change in supermarket expenditure, 7-days rolling average

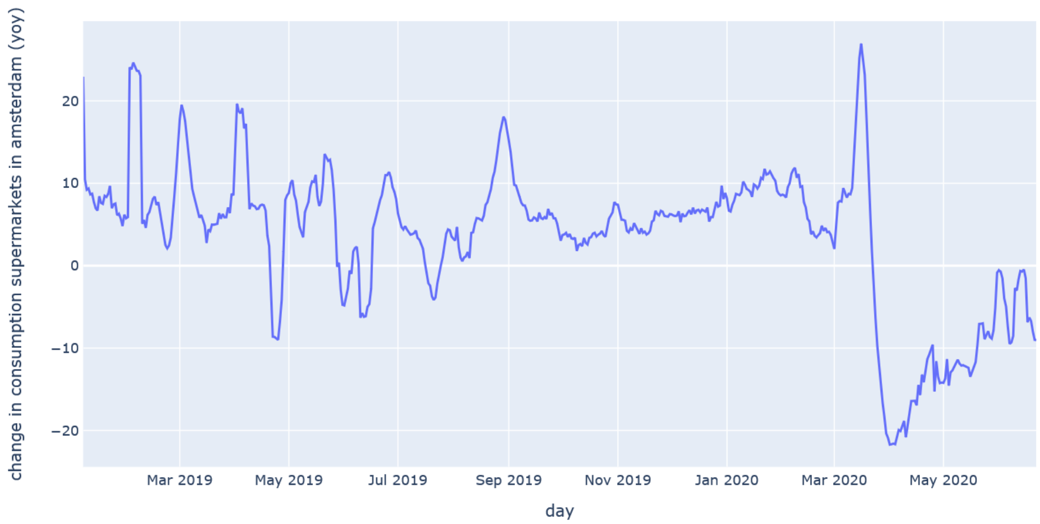

Figure 19: Amsterdam year on year change in supermarket expenditure, 7-days rolling average

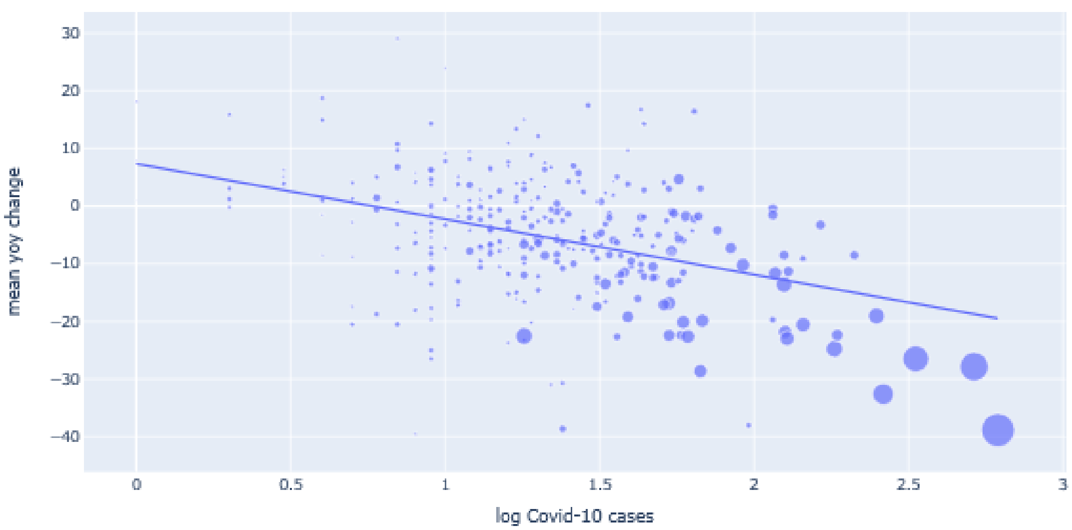

Figure 20: Comparative bubble chart. On the vertical axis, the average year on year drop in transactions between March 14 and April 16 and on the horizontal axis, the logarithm of cumulative Covid-19 cases up to April 16. The size of the bubble indicates the population size of the municipality.

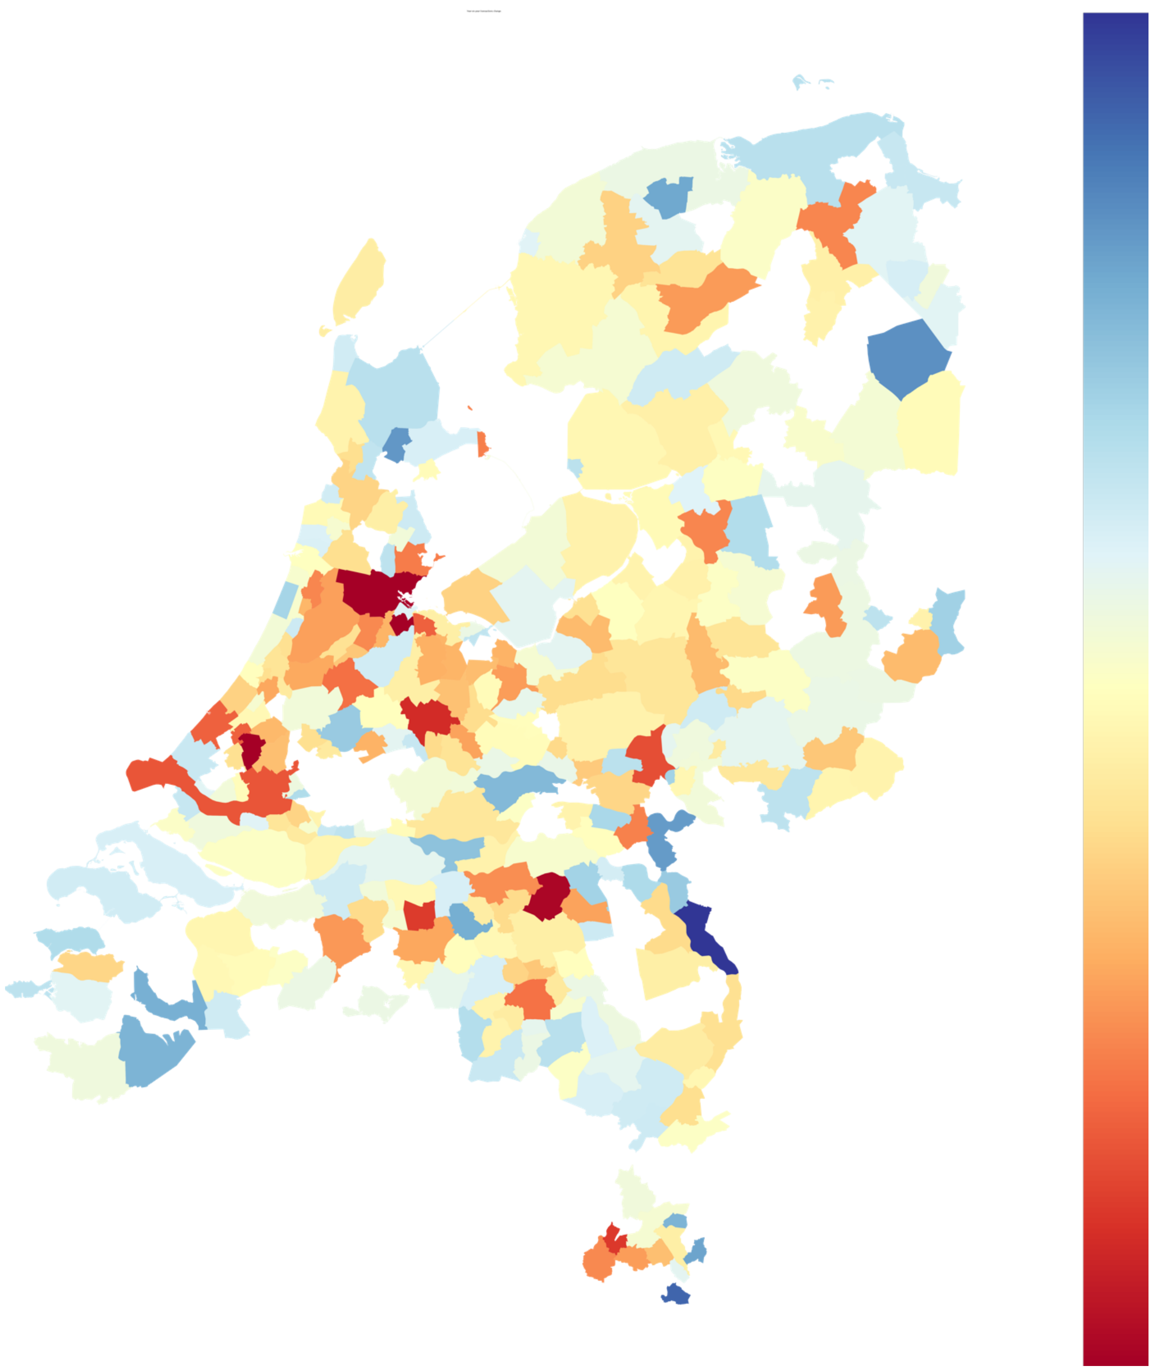

Figure 21: Map of the Netherlands showing the change in transactions (year on year) per municipality between March 13 and June 1 (the end of the nationwide lockdown). Blue indicates an increase in transactions and red indicates a decrease in transactions

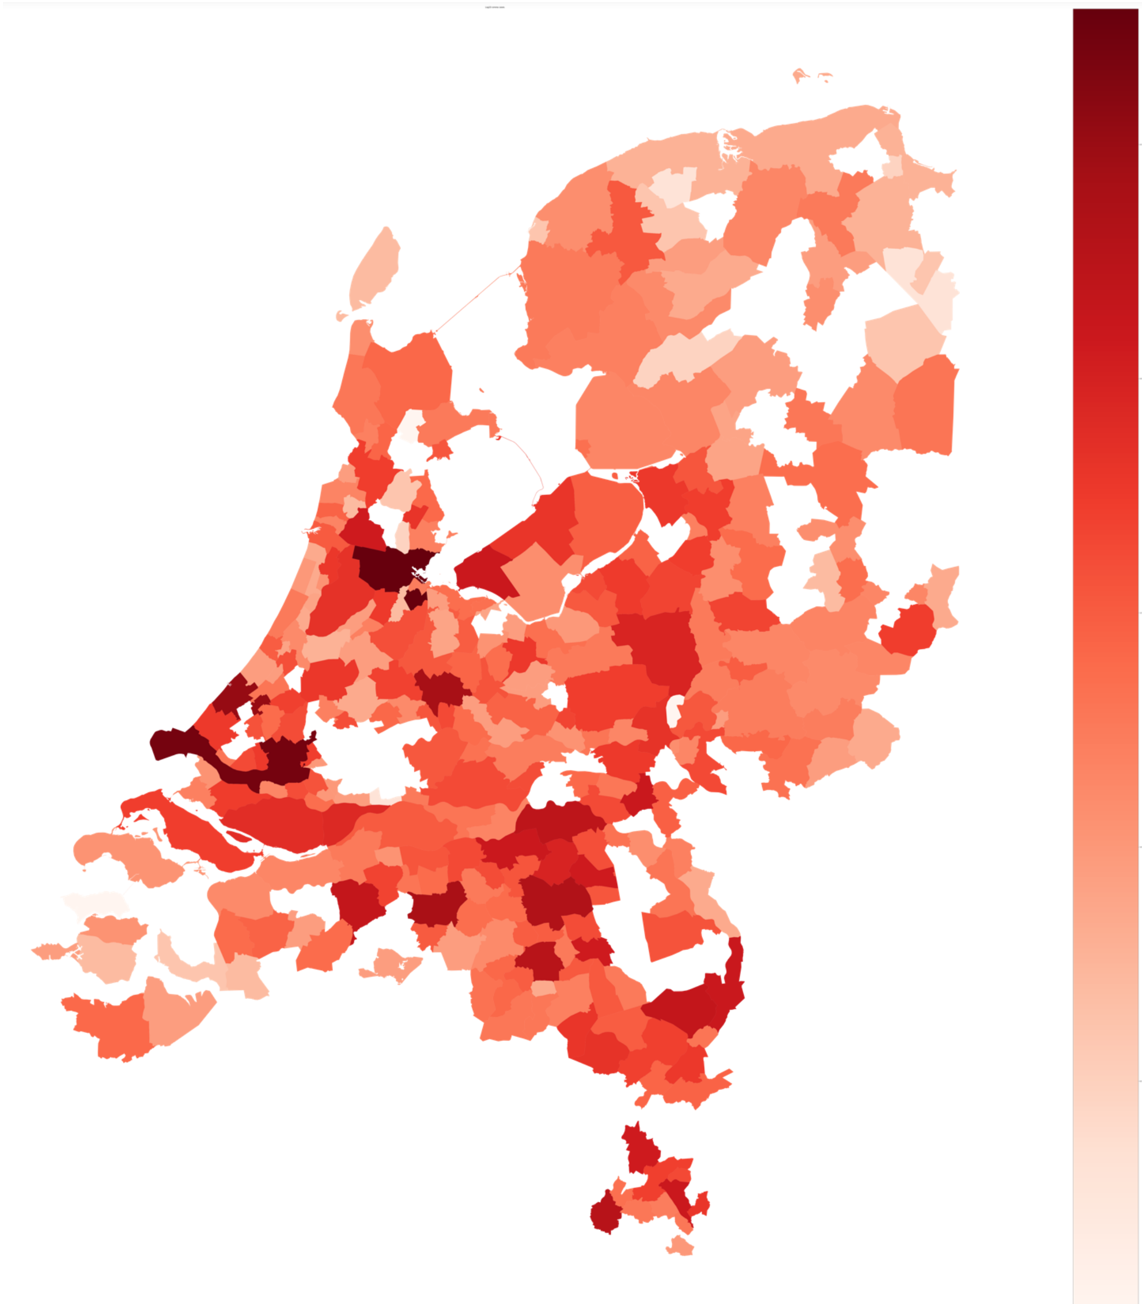

Figure 22: Map of the Netherlands showing the cumulative Covid-19 cases per municipality at June 1. Dark red indicates a large number of Covid-19 cases and light red indicates few Covid-19 cases. Municipalities that are coloured white have seen zero Covid-19 cases.

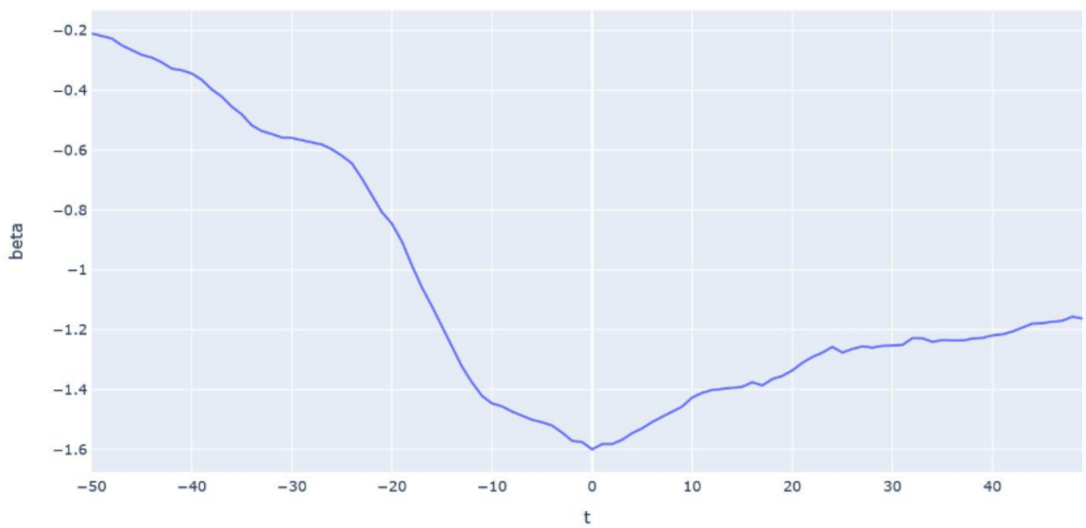

Figure 23: The development of the beta coefficient of the explanatory variable (X_{it}), the new amount of hospitalized Covid-19 cases, shifted over time. We use the second-stage regression (the regression of the residuals) to show the evolution of the beta coefficient. The horizontal axis t indicates the time in days and the vertical axis shows the beta coefficient of (X_{it}).

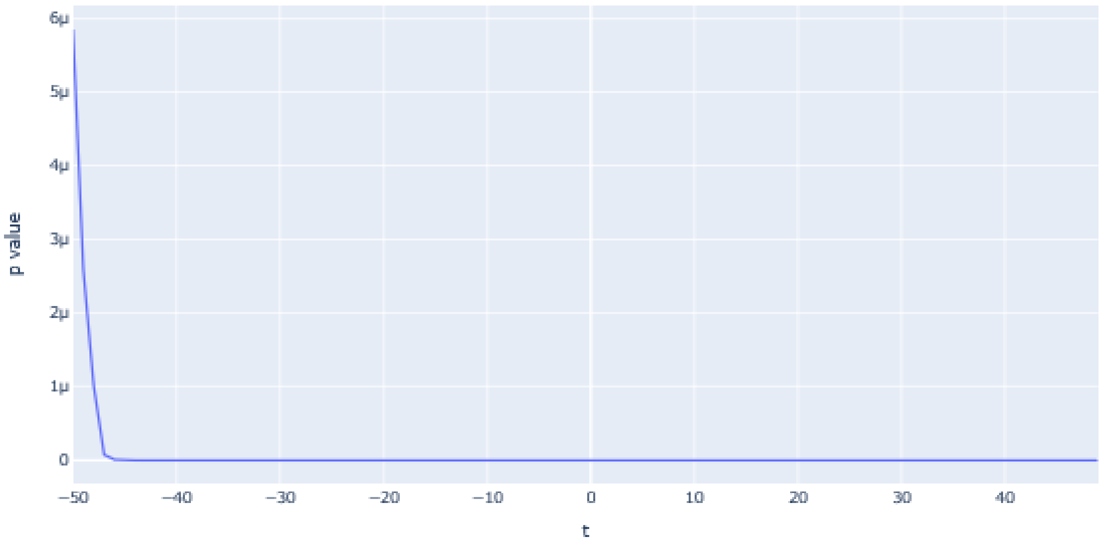

Figure 24: The development of the p-value of the explanatory variable (X_{it}), the new amount of hospitalized Covid-19 cases, shifted over time. We use the second-stage regression (the regression of the residuals) to show the evolution of the beta coefficient. The horizontal axis t indicates the time in days and the vertical axis shows the p-value of the explanatory variable (X_{it}).

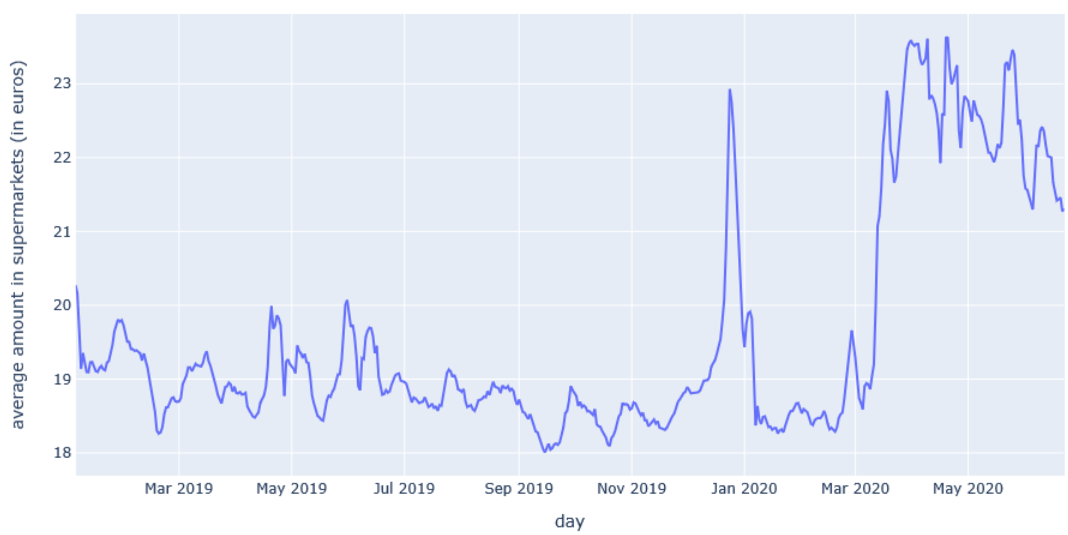

Figure 25: Average amount in euros per transaction spent with debit cards in supermarkets in the Netherlands, 7-days rolling average

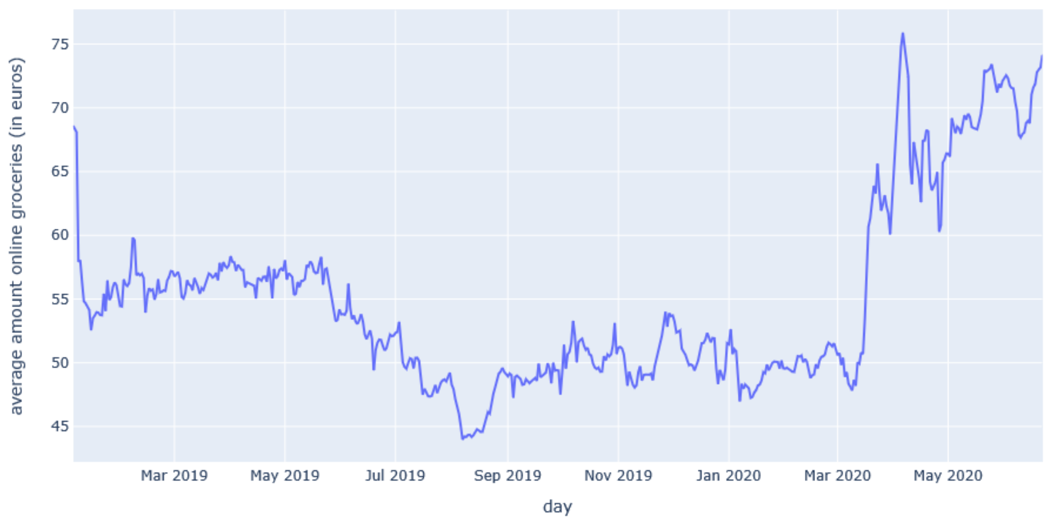

Figure 26: Average amount in euros per order spent on online groceries in the Netherlands, 7-days rolling average

Appendices

A Appendix

To determine geolocations for physical pin transactions we use an unsupervised model. Given that we want to determine dense clusters and classify all other points as outliers, a density-based method is the most suitable approach. As the dataset is rather large and clustering will need to happen for many sets of points, a less computationally expensive algorithm is preferred, and hence DBSCAN was selected as the clustering algorithm for this problem.

Parameter selection The DBSCAN clustering algorithm requires two main parameters to be set. The first, ϵ, is the maximum distance that two points can be from each other whilst still belonging to the same cluster. The second, $min\_samples$, is the minimum number of samples required for a group of points to become a cluster in the final results. As there are many varieties of payment points in the data, there is no 'one size fits all' approach. The goal with the setting of these parameters is to find the optimal (most dense whilst including the most samples) cluster for a set of home locations associated with a payment point, in order to be able to remove points which are outliers or do not improve the location prediction. For this reason, the initial parameters are set at very 'ambitious' levels based on the label of the payment point and the total amount of samples there are. However, when no cluster is found, the parameters are widened until either a cluster has been found or the parameter limit has been reached, and thus, hypothetically, no dense enough cluster exists in the samples to make a good prediction for the payment point.

Transaction weighting The distance between a customer's home and the payment point can differ according to the time of the day, month of the year and total amount spent. Therefore we introduce heuristic transaction weighting, also called transaction-weighted dynamic-parameter DBSCAN (TWDP-DBSCAN). These heuristics determine the 'weight' of a transaction, which can vary between 0 and 2. This weight is used to determine how strongly the transaction influences the final predicted location of the payment point. The final location (after clustering) is calculated as follows (for latitude and longitude separately):

$$location = \frac{\sum_{i=1}^{N} location_i * weight_i}{\sum_{i=1}^{N} weight_i}$$

where N is the amount of samples after clustering. Although this is based on analysis of a smaller sample, Van der Cruijsen (2018) [35] suggests that payment behaviors tend to be consistent across different groups of people and thus this sample is assumed to generalize well to the full dataset.

After the implementation the transaction-weighted dynamic-parameter DBSCAN (TWDP-DBSCAN), it was found that certain labels (categories) of payment points, such as hotels, did not perform well. This can be explained by the fact that consumers usually stay in hotels far from their home address. In order to account for the harder-to-predict categories, payment point matching was introduced, where each prediction was assigned a confidence based on the heuristics of the category, number of samples and cluster size. Payment points were then matched according to the following routine:

1. Partition the transactions by customer
2. Sort the transactions by transaction date and time
3. For each transaction, calculate the time delta between it and the customer's previous transaction
4. If this time delta is smaller than the threshold level, add the previous payment point with its predicted location and confidence as a 'match'

Finally, the lower confidence payment point locations are recalculated based on their own location and the locations of the matched payment points using an algorithm that takes into account the confidence of all predictions, and the time delta of all the matches.

Implementation details The TWDP-DBSCAN algorithm is implemented as a class in PySpark for parallel processing and it implements two public methods, *preprocess* and *predict*. After performing pre-processing, a column with the transaction weights is added and matches are calculated. Then TWDP-DBSCAN is executed for the first location prediction and if selected, the location is recalculated based on the matches. Next the locations are returned or city, municipality and province columns are added and returned. Finally, the predicted geolocations are coded to the respective cities/municipalities/provinces using the reverse geocoder library[18], which takes latitude/longitude pairs and returns the nearest respective town/city.

[18] `https://pypi.org/project/reverse_geocoder/`

The GeoNames[19] dataset was used for zipcode geolocation mapping. Although Ahlers (2013) [1] has determined that GeoNames is not perfect, and certainly has limitations regarding shortening of latitudes/longitudes and imperfections in the data, this is the main widely accessible resource able to map almost all of the Dutch zipcodes to a geolocation. Also, the method used (DBSCAN) is able to filter out most of the major problems of the dataset and the sheer amount of transactions used make the small imperfections relatively low-importance.

Validation The testing results displayed in Table 16 and Table 17 are on a transaction level and not on a payment point level. This is done for the reason that payment points with very few transactions are less relevant in the final result, as the number and value of transactions are what drive the correctness of the measures of consumption. The run-time was measured in a distributed PySpark environment with 128 cores for parallel processing.

Year	**Transactions**	**Payment Points**	**Customers**	**Label Source**
2015	64M	54,700	2.7M	Ground-truth dataset
2020	393M	1.88M	3.24M	Description (fuzzy city matching)

Table 15: The two test sets used for evaluation

Table 15 describes the test sets that were used for the final testing, the size of the datasets and how the label (location) was obtained. As there is no training with unsupervised learning (clustering), there is no need to split the dataset into training and testing samples, and the full dataset can be used for final tests. Also, both test sets went through the same pre-processing function where payment points with less than 10 or more than 50,000 transactions were removed and transactions with a value under 0.5 euros or over 5,000 euros were removed. This is done to remove non-significant transactions and outliers.

[19] http://geonames.org

Model	**MHE**	**CA**	**PA**	**Run-time**
Baseline	4,213 m	34.88%	85.05%	4 mins
TWDP-DBSCAN	981 m	71.03%	91.71%	39 mins
TWDP-DBSCAN with matching	981 m	71.03%	91.71%	94 mins

Table 16: Results on 2015 test set (N=64,000,000)

The results on the 2015 test data can be found in Table 16. In this table, MHE represents the *Mean Haversine Error*, CA is the *City Accuracy* and PA is the *Province Accuracy*. Also, the run-time of each model is included. The algorithm was run with all transactions, but only the results for the smaller subset reported in Table 15 are reported as these are the only payment points which have an available ground-truth location.

Model	**MHE**	**CA**	**PA**	**Run-time**
Baseline	N.A.*	33.66%	85.0%	8 mins
TWDP-DBSCAN	N.A.*	67.63%	94.56%	30 mins
TWDP-DBSCAN with matching	N.A.*	67.61%	94.55%	104 mins

Table 17: Results on 2020 test set (N=393,000,000)

**Not Available because the test set only contains cities, not exact locations. Cities are as large as 220km^2, and so would yield a very imprecise estimate.*

As the method has been developed primarily with exploratory data analysis on the 2015 dataset (as specific geolocations of payment points were available here), a good method of validation is to test the method on more recent data without changing the implementation or heuristics. The results, for the first five months of transactions of 2020, are presented in Table 17. On this much larger, diverse and high-quality sample, it is found that the baseline and TWDP-DBSCAN (with or without matching) results are similar to the 2015 results. This shows that the method is built in a robust way and generalizes well over time, as the samples are 5 years apart and no modifications to the method have been made prior to testing.

Limitations The validation data used are from 2015 and only represent a limited sub-sample of all payment points. They thus only contain small numbers of observations for certain labels and areas. This means that the validation set is biased towards certain labels and areas in the Netherlands. Nothing can be done about this limitation, except for manually creating more validation data. For this reason, 2020 transaction data has also been used, where the exact geolocation is not known but the city can be extracted from the transaction description.

B Appendix

Using a fixed effects model allows us to control for all time-invariant characteristics of an entity. For instance, in Table 3, we cannot investigate the effect of municipality-specific characteristics such as whether there is a hospital or not, or whether the municipality has a lot of economic activity in the sectors that we forced to close due to the lockdown. As these factors do not change over time, the fixed effects absorb this variation.[20] To investigate whether the above mentioned factors may impact the data on changes in transaction levels, we run the panel OLS model specified in equation (6), allowing time-invariant variables to be included (X_i). Table 5 shows the results. In the regression we include several variables that measure the average proximity of various public services. Proximity figures measure the average distance in kilometres of all citizens in a municipality to a given public service. We have included the proximity to a general practitioner (GP), hospital, supermarket, restaurant, cultural building (museums, theaters and cinemas) and public transport hub. Table 5 shows that target variable (X_{it}) is strong and statistically significant. From the control variables, the average distance to a GP is the only variable that is not statistically significant. However, the beta coefficients on the other variables do differ substantially. Table 5 indicates that in particular, the proximity to a restaurant and supermarket explain our dependent variable. This is probably because these two variables best capture the "big

[20]The fixed effects specification accounts for differences in municipalities that cannot be explained by the control variables. Statistically, time-invariant characteristics are perfectly collinear with the dummy, so their effect cannot be estimated in the model. These differences between municipalities could reveal idiosyncrasies that are hard to capture using additional regressors varying at the municipality level (for instance the fact that some municipalities are near the national border and so are affected by Covid-19 cases and consumers from outside the Netherlands). They could also capture some omitted factors excluded from our relatively concise specification.

city" effect. In general, big cities have seen larger outbreaks of the Covid-19 virus and consumers have withheld spending in these places because they are generally crowded and a potential source of contamination. Moreover, the proximity of restaurants may partly capture the lockdown effect, because restaurants were not allowed to let costumers in, but only to serve take-away orders.

References

[1] Dirk Ahlers. Assessment of the accuracy of geonames gazetteer data. In Proceedings of the 7th workshop on geographic information retrieval, pages 74–81, 2013.

[2] David Altig, Scott R Baker, Jose Maria Barrero, Nicholas Bloom, Philip Bunn, Scarlet Chen, Steven J Davis, Julia Leather, Brent H Meyer, Emil Mihaylov, Paul Mizen, Nicholas B Parker, Thomas Renault, Pawel Smietanka, and Greg Thwaites. Economic uncertainty before and during the covid-19 pandemic. Working Paper 27418, National Bureau of Economic Research, June 2020.

[3] Fernando E Alvarez, David Argente, and Francesco Lippi. A simple planning problem for covid-19 lockdown. Working Paper 26981, National Bureau of Economic Research, April 2020.

[4] Fernando E Alvarez, David Argente, and Francesco Lippi. A simple planning problem for covid-19 lockdown. Working Paper 26981, National Bureau of Economic Research, April 2020.

[5] Asger Andersen, Emil Toft Hansen, Niels Johannesen, and Adam Sheridan. Consumer Responses to the COVID-19 Crisis: Evidence from Bank Account Transaction Data. CEPR Discussion Papers 14809, C.E.P.R. Discussion Papers, May 2020.

[6] Asger Andersen, Emil Toft Hansen, Niels Johannesen, and Adam Sheridan. Consumer Responses to the COVID-19 Crisis: Evidence from Bank Account Transaction Data. CEPR Discussion Papers 14809, C.E.P.R. Discussion Papers, May 2020.

[7] Sangmin Aum, Sang Yoon (Tim) Lee, and Yongseok Shin. Covid-19 doesn't need lockdowns to destroy jobs: The effect of local outbreaks in korea. Working Paper 27264, National Bureau of Economic Research, May 2020.

[8] Collin Baker. FrameNet: A knowledge base for natural language processing. In Proceedings of Frame Semantics in NLP: A Workshop in Honor of Chuck Fillmore (1929-2014), pages 1–5, Baltimore, MD, USA, June 2014. Association for Computational Linguistics.

[9] Scott R Baker, R. A Farrokhnia, Steffen Meyer, Michaela Pagel, and Constantine Yannelis. Income, liquidity, and the consumption response to the 2020 economic stimulus payments. Working Paper 27097, National Bureau of Economic Research, May 2020.

[10] Scott R. Baker, R.A. Farrokhnia, Steffen Meyer, Michaela Pagel, and Constantine Yannelis. How Does Household Spending Respond to an Epidemic? Consumption During the 2020 COVID-19 Pandemic. SSRN Electronic Journal, 2020.

[11] Paul Beaudry and Franck Portier. Stock prices, news and economic fluctuations. Working Paper 10548, National Bureau of Economic Research, June 2004.

[12] Abel Brodeur, David M. Gray, Anik Islam, and Suraiya Jabeen Bhuiyan. A Literature Review of the Economics of COVID-19. IZA Discussion Papers 13411, Institute of Labor Economics (IZA), June 2020.

[13] Tomaz Cajner, Leland Crane, Ryan Decker, John Grigsby, Adrian Hamins-Puertolas, Erik Hurst, Christopher Kurz, and Ahu Yildirmaz. The u.s. labor market during the beginning of the pandemic recession. SSRN Electronic Journal, 01 2020.

[14] Vasco M Carvalho, Juan R Garcia, Hansen Álvaro, Ortiz Tomasa, Stephen Hansen, and Tomasa Rodrigo. Tracking the COVID-19 Crisis with High- Resolution Transaction Data. (April), 2020.

[15] Vasco M Carvalho, Juan Ramon Garcia, Stephen Hansen, Alvaro Ortiz, Tomasa Rodrigo, Jose v. Rodriguez Mora, and pep ruiz. Tracking the covid-19 crisis with high-resolution transaction data, April 2020.

[16] Haiqiang Chen, Wenlan Qian, and Qiang Wen. The impact of the covid-19 pandemic on consumption: Learning from high frequency transaction data. Working paper, SSRN.

[17] Raj Chetty, John N Friedman, Nathaniel Hendren, Michael Stepner, and The Opportunity Insights Team. How did covid-19 and stabilization policies affect spending and employment? a new real-time economic tracker based on private sector data. Working Paper 27431, National Bureau of Economic Research, June 2020.

[18] Jerry Coakley, Ana-Maria Fuertes, and Ronald Smith. A principal components approach to cross-section dependence in panels. 10th International Conference on Panel Data, Berlin, July 5-6, 2002 B5-3, International Conferences on Panel Data, 2002.

[19] Olivier Coibion, Yuriy Gorodnichenko, and Michael Weber. The cost of the covid-19 crisis: Lockdowns, macroeconomic expectations, and consumer spending. Working Paper 27141, National Bureau of Economic Research, May 2020.

[20] Selwyn Cornish. George A. Akerlof and Robert J. Shiller, Animal Spirits. How Human Psychology Drives The Economy, And Why It Matters For Global Capitalism (Princeton University Press, 2009). Agenda - A Journal of Policy Analysis and Reform, 16(4):117–122, 2009.

[21] Peter Ganong, Pascal J Noel, and Joseph S Vavra. Us unemployment insurance replacement rates during the pandemic. Working Paper 27216, National Bureau of Economic Research, May 2020.

[22] Austan Goolsbee and Chad Syverson. Fear, Lockdown, and Diversion: Comparing Drivers of Pandemic Economic Decline 2020. NBER Working Papers 27432, National Bureau of Economic Research, Inc, June 2020.

[23] Sumedha Gupta, Thuy D. Nguyen, Felipe Lozano Rojas, Shyam Raman, Byungkyu Lee, Ana Bento, Kosali I. Simon, and Coady Wing. Tracking Public and Private Responses to the COVID-19 Epidemic: Evidence from State and Local Government Actions. NBER Working Papers 27027, National Bureau of Economic Research, Inc, April 2020.

[24] Sinem Hacioglu, Diego R Känzig, and Paolo Surico. Consumption in the time of Covid-19: Evidence from UK transaction data. CEPR Discussion Papers 14733, C.E.P.R. Discussion Papers, May 2020.

[25] Inoue Hiroyasu and Todo Yasuyuki. The Propagation of Economic Impacts through Supply Chains: The Case of a Mega-city Lockdown to Prevent the Spread of COVID-19. Discussion papers 20037, Research Institute of Economy, Trade and Industry (RIETI), April 2020.

[26] Can Hou, Jiaxin Chen, Yaqing Zhou, Lei Hua, Jinxia Yuan, Shu He, Yi Guo, Sheng Zhang, Qiaowei Jia, Chenhui Zhao, Jing Zhang, Guangxu Xu, and Enzhi Jia. The effectiveness of quarantine of wuhan city against the corona virus disease 2019 (covid-19): A well-mixed seir model analysis. Journal of Medical Virology, 92(7):841–848, 2020.

[27] Nicole Jonker, Lola Hernandez, Renate de Vree, and Patricia Zwaan. From cash to cards: how debit card payments overtook cash in the Netherlands. 2017.

[28] George Kapetanios and Hashem Pesaran. Alternative approaches to estimation and inference in large multifactor panels: Small sample results with an application to modelling of asset returns. The Refinement of Econometric Estimation and Test Procedures: Finite Sample and Asymptotic Analysis, 06 2005.

[29] Christian A Moser and Pierre Yared. Pandemic lockdown: The role of government commitment. Working Paper 27062, National Bureau of Economic Research, April 2020.

[30] M. Hashem Pesaran. Estimation and inference in large heterogeneous panels with a multifactor error structure. Econometrica, 74(4):967–1012, 2006.

[31] Stanislav Sobolevsky, Emanuele Massaro, Iva Bojic, Juan Murillo Arias, and Carlo Ratti. Predicting regional economic indices using big data of individual bank card transactions. Proceedings - 2017 IEEE International Conference on Big Data, Big Data 2017, 2018-January:1313–1318, 2017.

[32] Stuart N. Soroka. Good news and bad news: Asymmetric responses to economic information. Journal of Politics, 68(2):372–385, 2006.

[33] PAUL C. TETLOCK. Giving content to investor sentiment: The role of media in the stock market. The Journal of Finance, 62(3):1139–1168, 2007.

[34] Nederlandse Vereniging van Banken. Betalingsverkeer (https://www.bankinbeeld.nl/thema/betalingsverkeer/).

[35] Carin van der Cruijsen and Joris Knoben. Ctrl C Ctrl Pay: Do People Mirror Payment Behaviour of Their Peers? SSRN Electronic Journal, (611), 2018.

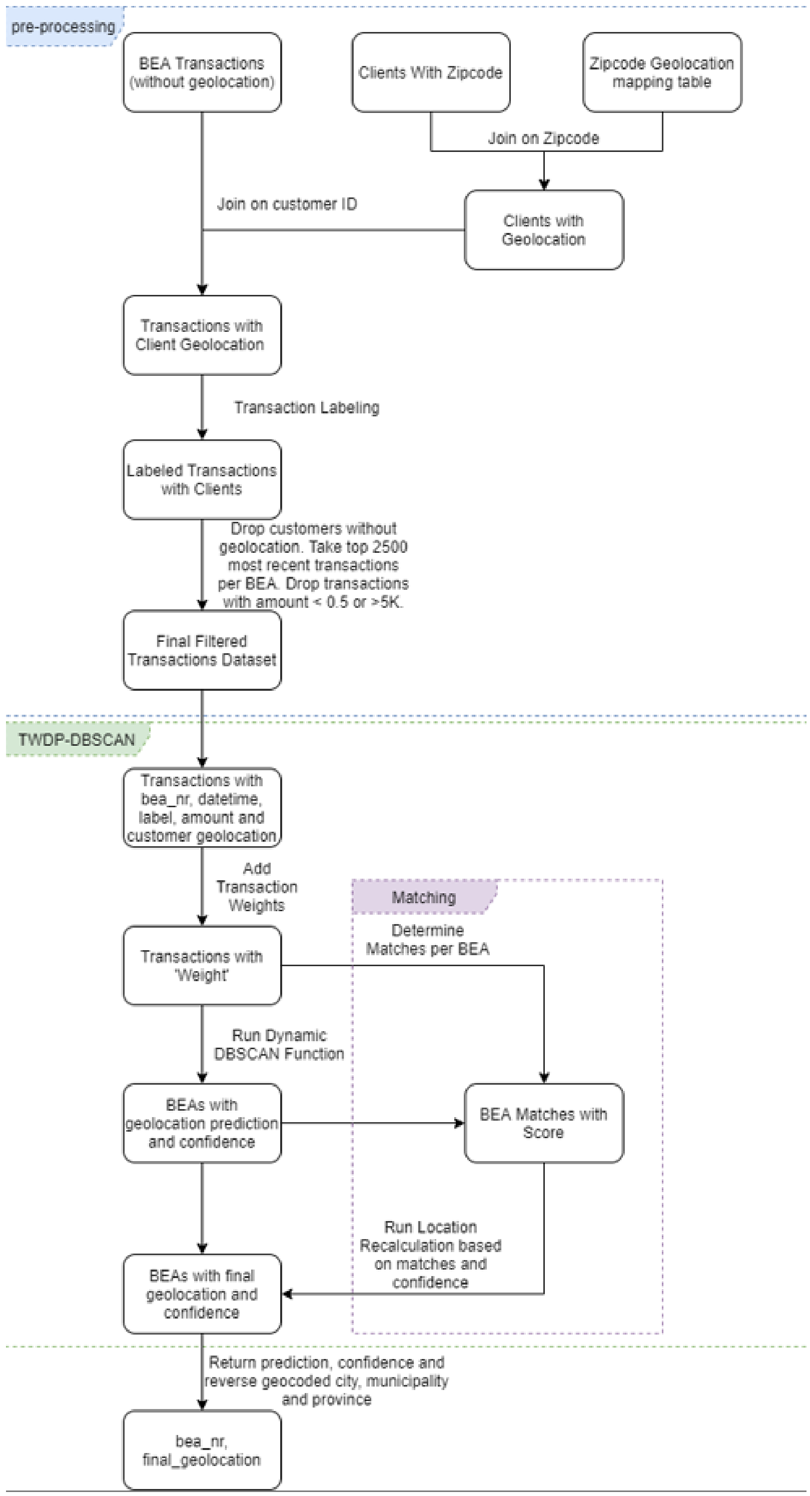

Figure 27: Transaction-weighted dynamic-parameter DBSCAN

The Data Analytics for Finance and Macro (DAFM) Research Centre
Forecasting trends is more important now than it has ever been. Our quantitative financial and macroeconomic research helps central banks and statistical agencies better understand their markets. Having our base in central London means we can focus on issues that really matter to the City. Our emphasis on digitalisation and the use of big data analytics can help people understand economic behaviour and adapt to changes in the economy.

Find out more at kcl.ac.uk/dafm

KING'S BUSINESS SCHOOL

Data Analytics for Finance & Macro Research Centre