Tracking the Impact of the Covid-19 Pandemic with the Use of High-Frequency Geo-Located Bank Transaction Data

Nora Neuteboom
Kings College London
nora.neuteboom@nl.abnamro.com

Feiko Ritsema
ABN AMRO Bank
feiko.ritsema@nl.abnamro.com

George Kapetanios
Kings Business School
george.kapetanios@kcl.ac.uk

Abstract

Using geo-located transaction data from 2 million customers of ABN AMRO bank in the Netherlands, this paper distinguishes the economic effects of consumers responses to the Covid-19 pandemic from those attributable to non-pharmaceutical interventions (NPIs). We compare municipalities that experienced large Covid-19 outbreaks with municipalities that had few or no cases and find that during the first Covid-19 wave the scale of the outbreak in a municipality has a strong negative effect on physical transactions by consumers in that municipality. This behavioral response function of consumers towards the virus is however not constant over time. During the second Covid-19 wave, the behavioural effect of consumers towards the virus has no real impact on consumption.

1. Introduction

Recent research has shown that during the Covid-19 pandemic, consumption has been hit directly by the non-pharmaceutical interventions (NPIs) and the behavioral response function of consumers towards the virus. The NPIs, or sometimes called lockdown measures, involved closing (sub)sectors or placing supply restrictions on others, reducing to opportunity for consumers to spend. Moreover, as consumer were afraid of contracting the virus, or are infected themselves, they reduced their consumption, especially in crowded places were they can easily contract/spread the virus, such as restaurants, bars and large-scale events. Papers that study the consumer behavior during the Covid pandemic has shown that the behavioral response function of consumers towards the virus has a strong negative effect on consumption during the early stages of the pandemic, concluding that without the NPIs the damage to the economy could have been even bigger.
However, we know from behavioural economics studies that behavior from consumer is not always constant over time. The question therefore is, was the behavioral response function of consumers towards the virus constant as the Covid-19 pandemic evolved, or did it change?

The Dutch government's strategy of imposing nationwide NPIs measures presents us with an unique empirical estimation strategy, as both the incidence of the illness and its timing varied substantially across different municipalities. While the NPIs induced homogeneous expenditure dynamics across all municipalities in the Netherlands, spatial heterogeneity remains because the Covid-19 virus may have had different effects in different municipalities across the country. While similar empirical studies have been done in other countries, for instance in China and France (Chen, Qian and Wen (2020) [1] and Landais et al. (2020) [2]), the Netherlands was one of the only countries that saw a NPI policy that was consistent among regions in terms of strictness and imposition date. Therefore, contrary to other studies so far, we don't have to manually control for these differences between regions within the country.

Moreover, this paper is, to our knowledge, the first empirical paper that investigates the behavioral response of consumers during the full period of the Covid-19 pandemic, until the vaccination roll-out started. Our methodology for investigating the behavioural response is to compare Dutch municipalities which were hard hit by the virus, versus those which were barely affected by Covid.

We support the claim by other authors that in the short-run the behavioral response function of consumers was significant in explaining the drop in consumption. However, we also show that this effect disappears over time and becomes insignificant during the second Covid-19 wave.

1.1. NPIs in the Netherlands

On 12 March, the Dutch government announced an intelligent lockdown. Effectively, the next day, all events (concerts, sports) and all meetings with more than

HǏCSS

100 people were forbidden. Bars, restaurants and other public places or venues where people gather had to close as well. The timing and severity of the measures were generally comparable to most of northern Europe (such as Germany and Norway), but less restrictive than in southern Europe, where the virus spread more rapidly (such as Italy, France and Spain). Similar to other European countries, the Netherlands largely reopened again in the summer. As the amount of hospitalized cases picked-up during the autumn of 2020, on 14 October a partial lockdown came into effect. As this did not help to bring down the Covid-19 cases sufficiently, from 15 December 2020, a hard lockdown was imposed, with strict NPIs. All non-essential shops were closed and a curfew was imposed. As the vaccination pace picked-up in the first months of 2021, measures were slowly being relieved. On 28 April shops and outdoor seating areas at restaurants and cafes were allowed to reopen and the curfew was lifted. On 19 May, most of the NPIs were relaxed.

2. Literature review

Our work builds on and contributes to a rapidly evolving literature on measuring the economic impacts of Covid-19. For a good overview on the economic effect of Covid-19 see our previous work on this topic (Neuteboom et al. (2020) [3]). Various papers use high-frequency transaction data, analogous to the data we assemble here, to analyze aggregate consumer spending during the Covid-19 pandemic. There are a number of empirical studies that have attempted to disentangle the direct effects of NPIs and the behavioral response towards the virus on consumption. In general, there is consensus in the literature that changes in behaviour by consumers towards to virus explain a large part of the fall in consumption during the first Covid-19 wave. Whilst these studies typically find that NPIs have had a short-term effect on consumption, the NPIs only account for a small proportion of the observed declines (See: Chetty et al. (2020) [4]; and Landais et al. (2020) [2]; Neuteboom et al. (2020) [3])

Our research is most closely related to the study by Chen, Qian and Wen (2020) [1]. 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.

2.1. Effects on the longer-run

Whilst the short-run empirical evidence described above provide valuable insights, they don't allow us track consumer behaviour over the entire duration of the pandemic.

From behavioural economics, we know that individuals are not always maximizing their utility. Research has demonstrated that people's judgements and decisions are often subject to systematic biases and heuristics, and are strongly dependent on the context of the decision (See: Reisch and Zhao (2018) [5]). From that perspective it would be naive to assume that the behavioural effect that Chen, Qian and Wen (2020) [1] and other rightly point out during the first stage of the pandemic, will continue to persist throughout the full period.

Moreover, studies from psychology acknowledges that individuals may undergo behavioural changes due to the occurrence of particular events, including natural disasters, healthcare crises and terrorist attacks (see for example Mehta, Saxena and Purohit (2020) [6]). Additionally, as there are empirical papers in consumer behaviour on modern pandemics outside of SARS, which was fairly localised.

There are some theoretical attempts to study the behavioural aspect of consumers towards the virus. Moser and Yared (2020) [7] study the optimal lockdown policy and argue that governments are faced with the choice to limit the extent of future lockdown in order to support more optimistic investor expectations and credibility. Others, such as Alvarez, Argente and Lippi (2020) [8] study the optimal lockdown policy with the SIR 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, covers 60% of the population after a month, and is gradually withdrawn covering 20% of the population after 3 months. The model by Boganni et al. (2020) [9] predicts a significant endogenous reduction in economic activity by agents in response to the spread of the virus.

The problem is that these models are highly stylised and abstract from many real world features of the pandemic that are likely to be important. For instance, few models include a mechanism whereby infection risk lowers consumption through voluntary choices – most models generated a link between the economy and the pandemic by simply assuming infected individuals and those who die don't produce any output and don't consume. Acemoglu et al. (2020) [10] shows that NPIs

can lower deaths and therefore increase output, while Chetty et al. (2020) [4] empirically shows that this mechanisms is relatively unimportant for the Covid-19 pandemic compared to the behavioural effect of people who are not infected. Boganni et al. (2020) [9] model the response of by agents towards the virus, but assume that this response is constant over time while our paper shows that this is actually not the case.

3. Data and methods

In the following section we describe the used data in more detail. Moreover, we describe the methodology for the geo-location model and the regressions that we use in the results section. We use similar data to our previous work on this topic Neuteboom et al. (2020) [3], but than for a much longer time horizon and we also now include unique geolocation data on employment by municipality.

3.1. Transaction data

We use real-time 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 [11]). 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 when costumers pay by card, ATM data from cash machines when costumers draw cash and online transactions by e-banking iDEAL transactions. In 2019, on average, clients paid 57% by card (PoS), 24% online by iDEAL and 19% by cash. The Point of Sale (PoS) data include a timestamp, the amount in euros, the corresponding account number, the counter party description and the counterparty account number. We have used a labeling function based on keywords in the transaction description to identify the category the transaction belongs to.

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. [1]

In order to measure the total consumer expenditure, we use an aggregate of all pin card transactions by municipality in euros. To investigate whether people may have changed their behavior, due to a preference switch toward more online goods, we use credit card and iDeal (internet banking) supermarket expenditures by municipality in euros.

3.2. Other variables

Local unemployment - We compiled data from the Dutch UWV (Employee Insurance Agency), an autonomous administrative authority that implements employee insurances and social benefits. We use the weekly figure of the unemployment benefits statistics, based on the number of unemployment benefits registered in the UWV administration for that week by municipality.

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.[2] The National Institute for Public Health and the Environment provides the cumulative number of hospitalzed Covid-19 cases on a daily basis by municipality. We use the seven-day moving average of new daily Covid-19 patients that have been admitted to the hospital. Patients that are geolocated by their home address, i.e. their individual residential location. For example, if a Covid-19 patient is admitted to a hospital in Amsterdam, but lives in Rotterdam, the new hospitalized Covid-19 patient will included in the amount of hospitalized covid-19 patients in the municipality of Rotterdam.[3]

Stringency index - We make use of the Oxford COVID-19 Government Response Tracker from the Blavatnik School of Government and the University of Oxford (Hale et al. [13] (2021)). Specifically, we look at the Stringency Index. This tracker collects information on several different common policy responses that governments have taken to respond to the pandemic, using 9 sub indicators: school

---

[1]Note that this is a drawback of the Carvalho et al. (2020) [12] 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.

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

[3]We prefer the amount of hospitalized cases over the confirmed Covid-19 infections as over time the testing capacity has been scaled up. If we would use the total covid-19 infections measured by the amount of positive test outcomes, we would drastically underestimate the real amount of Covid-19 patients in the early months of the pandemic, as during these months testing capacity was still insufficient.

closing, workplace closing, cancellation of public events, restrictions on gatherings, public transport closing, stay at home requirements, restrictions on internal movement, internal travel controls and public info campaigns. The data from the 9 indicators are aggregated into a set of four common indices, reporting a number between 1 and 100 to reflect the level of government action. While not all sub indicators in the Stringency Index do necessarily hampering consumption, it is the best proxy we have of the NPI policy.[4]

3.3. 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.

- *The lagged variable of the dependent variable* is denoted Y_{it-n} for municipality i at time $t - n$, where $n = (1, 2, \ldots, N)$.

- *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

.

3.4. Data transformations

Our variable of interest is Y_{it}, the total spending of ABN AMRO clients in euros for municipality i, where t can be aggregated over day or week. We calculate $Y_{it} = V_{it} \times P_{it}$. Where V_{it} is volume of transactions and P_{it} is the price of every transactions in euros.

For the aggregations by day, in order to control for seasonality trends, we pair every day with its equivalent day in the previous year and calculate the year-on-year differences. For the weekly data, we pair every week with its equivalent week number in the previous year. In addition to adjusting for seasonal patterns, we also manually adjust for calendar effects. The data depict large periodic fluctuations across days. For the daily aggregation, we report 7-day moving averages to smooth daily fluctuations.

3.5. Geo-location model

We have for every account-holder the zipcode of their registered home address. These zipcodes are merged with an external dataset featuring all Dutch zipcodes and their latitudes and longitudes[5] 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.

We use transaction-weighted density-based clustering to predict POS geolocations. 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 large and clustering will need to happen for many sets of points, a less computationally expensive algorithm is preferred, and hence DBSCAN (Density-based spatial clustering of applications with noise) was selected as the clustering algorithm for this problem.

3.6. Parameter selection

The DBSCAN clustering algorithm[6] 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. 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. 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.

3.7. 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 = $\sum_{i=1}^{N} location_i * weight_i \frac{1}{\sum_{i=1}^{N} weight_i}$ where N is the amount of samples after clustering. Although this is based on analysis of a smaller sample, Van der Cruijsen (2018) [14] 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. See figure 1 for the overview of the algorithm used.

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.
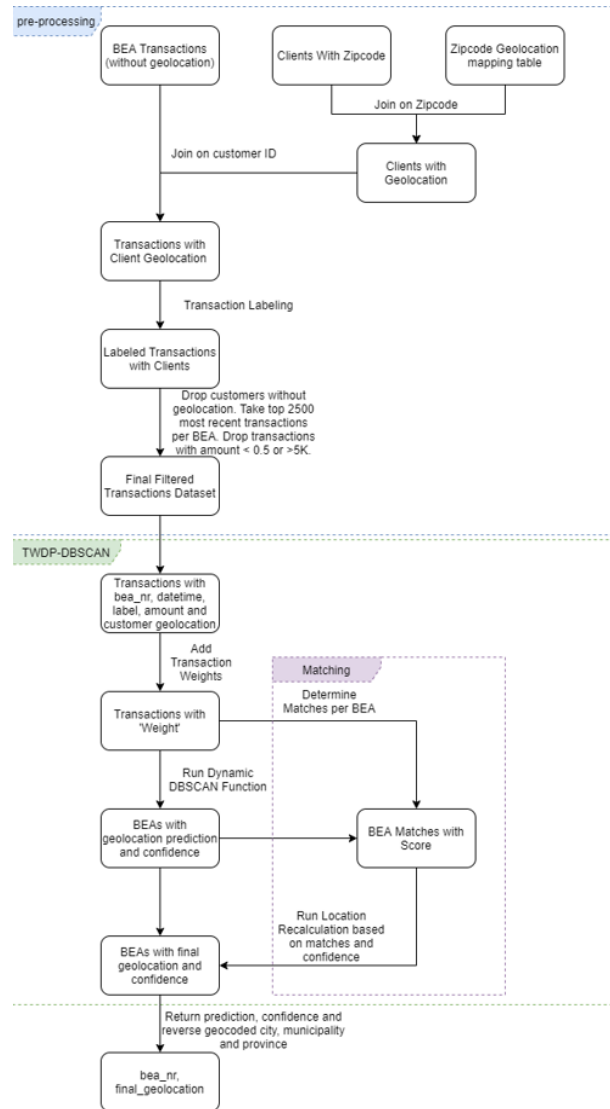

Figure 1. Transaction-weighted dynamic-parameter DBSCAN

3.8. Validation

Table 1 describes the test sets that were used for the final testing, the size of the datasets and how the label (location) was obtained. The ground-truth dataset for validation was a dataset from 2015 from a company collecting PoS locations. 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 preprocessing 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.

Table 1. The two test sets used for evaluation
.

Year	Trans	PoS	Clients	Data
2015	64M	54,700	2.7M	Ground-truth dataset
2020	393M	1.88M	3.24M	Description

The results on the 2015 test data can be found in Table 2. 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 1 are reported as these are the only payment points which have an available ground-truth location.

Table 2. Results 2015 test set (N=64,000,000)
.

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

As the method has been developed primarily with exploratory data analysis on the 2015 dataset (as specific geolocations of payment points were available), 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 3. 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.[7] This methodology predicts the correct province with 95% accuracy, the correct municipality with 86% accuracy and the correct city with 67% accuracy.

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 from

[7]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. Therefore, we have manually created more validation data where we extracted the geolocation from the transaction description. By matching this to our geolocation algorithm, we found a 97% overlap.

Table 3. Results on 2020 test set (N=393,000,000)
.

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

our dataset. We also delete municipalities for which we have incomplete data, which leaves us with 332 municipalities out of 355 for the main regressions. 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.

3.9. Regressions

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. We use the following equation:

$$Y_{it} = \alpha_i + \gamma_t + \beta_1 X_{it} + \beta_1 Y_{it-n} + \varepsilon_{it}, \qquad (1)$$

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 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. This also eliminates any variability between municipalities, for example the size of the municipality and whether it is urban or rural area. Spatial autocorrelation is accounted for by including an autoregressive parameter in the model specification. Moreover, the unobserved factors α_i could be correlated with X_{it} and therefore errors are correlated contemporaneously across municipalities. Pesaran (2006) [15] shows that by including the

cross-sectional averages in the regression the differential effects of unobserved common factors are eliminated. We therefore run the all the main regression from table in two stages to eliminate any cross-sectional dependence.[8] Due to space limitations we cannot show all the results, but all regression were significant at a 1% interval level.

We run the regression for two distinctive periods: the first Covid-19 wave and corresponding NPIs, ranging from the first of January until the beginning of June and the second Covid-19 wave and corresponding NPIs, ranging from September to mid-March. We have chosen mid-March as cut-off date as the vaccination pace starts to increase rapidly from that date onward, which may skew the results as we want to measure the behavioural response of consumers to the virus. At mid-March around 11% of the population have had its first doses.

4. Results

Table 4 shows the regression results for the panel regression of equation (1) for the first Covid-19 wave. The dependent variable (Y_{it}) is the year on year change in total volume of transactions by municipality i by week t, scaled by the amount of ABN Amro clients in municipality i. The key explanatory variable is the new amount of hospitalized Covid-19 cases by week t for municipality i, scaled by the amount of inhabitants of municipality i. We cluster standard errors at municipality level. Regression (1) includes both municipality and time fixed effects. Given that in the Netherlands the strictness and timing of the NPIs were identical for all municipalities, this effect is captured by time fixed effects. We find a strong statistically significant negative coefficient on the explanatory variable (X_{it}). For every new hospitalized Covid-19 per 1000 citizens, transactions drop by about 0.11%. During the peak of the first Covid wave, around 23 new cases were admitted to the hospital everyday in Amsterdam. On a population of around 870.000 people, that would result in a drop of total pin transactions of 2% in a week.[9] In regression (2) in table 4 we add

two lags of the dependent variable (Y_{it}). Including lagged dependent variables can reduce the occurrence of autocorrelation arising from model misspecification. The result is robust to the addition of the lags, as the variable of interest - the amount of hospitalized Covid-19 cases - is still significant. However, the beta coefficient has become slightly smaller.[10]

Another possible explanation for the effect of the Covid-19 variable that is picked up in regression (1) and (2) is that the local virus outbreak adds to the negative sentiment and uncertainty, or reminds people that tough economic times are ahead and therefore they increase their precautionary savings at the expense of spending. This explanation cannot be fully attributed to the virus alone, as normally during economic downturn consumers increase precautionary savings. To further eliminate this effect, we include a variable for unemployment in regression (3), which captures the macro-economic effect of loss of income due to lay-offs. This variable is also significant and shows for every new unemployed worker per 1000 citizens, transactions drop by about 0.25%. However, this result is not robust to the inclusion of the additional lags of the dependent variable (see regression (4)). The autocorrelation of the dependent variables absorbs the significance of unemployment.

Table 4. First Covid-19 wave - Dependent Variable: Transactions$_{it}$

Regressors	(1)	(2)	(3)	(4)
Covid-19$_{it}$	-0.1094***	-0.0869***	-0.1014***	-0.0891***
	(0.0406)	(0.0312)	(0.0423)	(0.0332)
Transactions$_{it-1}$		0.5993***		0.6026***
		(0.0105)		(0.0107)
Transactions$_{it-2}$		0.0995***		0.0993***
		(0.0054)		(0.0055)
Unemp$_{it}$			-0.2598***	-0.0785
			(0.0669)	(0.0527)
N	332	332	332	319
T	30	30	30	30
FE	Both	Both	Both	Both
R^2	0.0010	0.3800	0.0031	0.3840

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

To summarize, during the first Covid-19 wave, we find a strong negative effect of the amount of hospitalized Covid cases on consumption. This conclusion changes drastically as soon as we look at the second Covid-19 wave, taking place end-2020 and beginning of 2021. Table 5 shows the regression results for the panel regression of equation (1) for the second Covid-19 wave.

---

[8]Under the assumptions explained in Pesaran, 2006 [15], for any fixed m in α_i these residuals provide consistent estimates of ε_{it} in the multifactor model (1) 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 **??**. 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.

[9]For our fixed effect regression (1), the value of R^2 is relatively low. This is contrary to findings of Chen, Qian and Wen (2020) [1] and Goolsbee and Syverson (2020) [16], 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 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.

[10]We have tested the model with the inclusion of more lags, but after the t_{-2} the included lags are not significant under a 10% confidence interval.
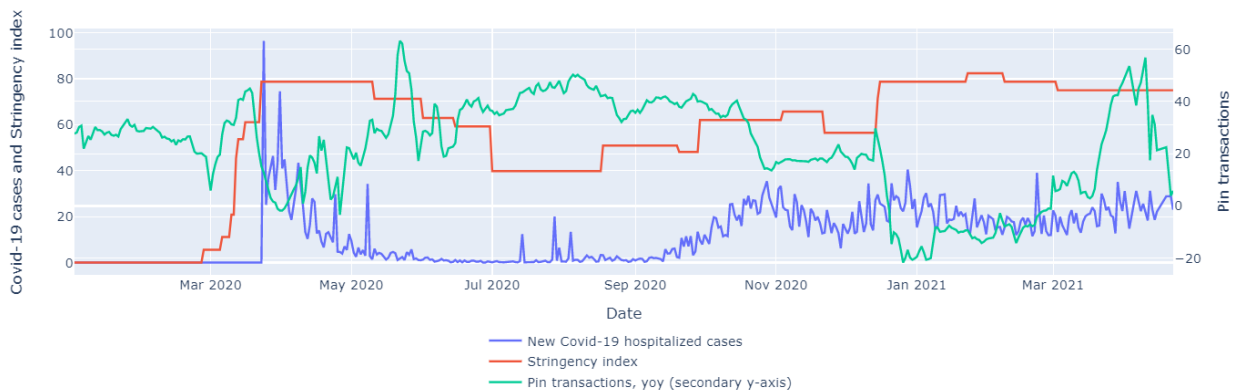
Figure 2. Year on year transactions by ABN Amro clients (euros), new hospitalized Covid-19 cases and the proxy of the NPIs measured by the stringency index

Table 5 regression (1) shows that the amount of hospitalized Covid-19 cases is not significant. This suggests that during the second Covid-19 the behavioral response of consumers changed and had no effect on physical spending. What is striking is that unemployment is significant in regression (3) and (4) and hence during the second Covid-19 wave robust to the introduction of the lags of the dependent variable. This results show that municipalities that saw larger increases in unemployment also saw a larger drop in transactions. For every new unemployed worker per 1000 citizens, transactions drop by about 0.05% (regression (3)). This suggests that unemployment had a strong effect on physical spending during the second Covid-19 wave. We investigate this claim further in the next section 4.1.

Table 5. Second Covid-19 wave - Dependent Variable: Transactions$_{it}$

Regressors	(1)	(2)	(3)	(4)
Covid-19$_{it}$	-0.0009	-0.0244	0.0019	-0.0219
	(0.0390)	(0.0372)	(0.0403)	(0.0384)
Transactions$_{it-1}$		0.1024***		0.1007***
		(0.0110)		(0.0113)
Transactions$_{it-2}$		0.2488***		0.2488***
		(0.0107)		(0.0110)
Unemp$_{it}$			-0.0796***	-0.0505**
			(0.0285)	(0.0267)
N	332	332	332	319
T	30	30	30	30
FE	Both	Both	Both	Both
R^2	4.246e-05	0.1210	0.0010	0.1210

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

4.1. Unemployment

How are the amount of Covid-19 cases and unemployed interlinked? Table 6 show the regression

results of the target variable hospitalized Covid-19 cases on unemployment. Regression (1) show the results for the first Covid-19 wave. In the fixed effects regression, Covid-19 is not significant for explaining the drop in unemployment. In other words, the variance of Covid-19 outbreak in different municipalities did not lead to higher unemployed in these same municipalities. This would support the claims made by Coibion et al. (2020) [17] and Landais et al. (2020) [2] that the drop in unemployment is mainly caused by the NPIs, instead of the Covid-19 cases.

Regression (2) shows the results from the second Covid-19 wave. This time the Covid-19 variable is significant, albeit under a 10% confidence interval. The coefficient is positve, suggesting that an increase in local Covid-19 cases results in an increase in local unemployment. Regression (3) shows that this results in robust to the inclusion of lagged dependent variables. Regression (4) adds the Stringency Index. This variable is not significant. This is intuitive, given that the Stringency Index does not change between municipalities (only over time) and with the knowledge that unemployment started the rise structurally in the summer, when the Stringency Index was low.

The relationship between unemployment and consumption is well-established (see ref 2). However, the relationship on local level is less well-researched. In other words, does an increase in local unemployment lead to a decrease in local spending? According to our data is does. Table 7 show the regression results of regression unemployment on the total transactions. Also in this regression we look at the year-on-year differences scaled by population. It shows that throughout the sample period, containing the first and second Covid-19 wave, unemployment is significant in explaining changes in transactions. From this

Table 6. First and Second Covid-19 wave - Dependent Variable: Unemployment$_{it}$

Regressors	1st wave	2nd wave	2nd wave	2nd wave
Covid-19 cases$_{it}$	-0.0102	0.0254*	0.02404**	0.0250**
	(0.0077)	(0.0154)	(0.0119)	(0.0110)
Unempl$_{it-1}$			0.5076***	0.5074***
			(0.0117)	(0.0117)
Unempl$_{it-2}$			0.2085***	0.2083***
			(0.0120)	(0.0120)
Stringency$_t$				-0.0081
				(0.0067)
N	332	332	332	319
T	30	30	30	30
FE	Both	Both	Both	Both
R^2	0.0002	0.0003	0.4006	0.4006

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

we can conclude that during the second Covid-19 wave, Covid-19 may not have had a direct effect on consumer spending, it did had an indirect effect by the unemployment variable.

Table 7. Whole sample period - Dependent Variable: Transactions$_{it}$

Regressor	(1)
Unemployment$_{it}$	-0.0616**
	(0.0312)
N	319
T	65
FE	None
R^2	0.0002

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

4.2. Changing consumer behaviour

Can we conclude that the consumers showed no behavioural response whatsoever during the second Covid-19 wave? We find that differences in Covid-19 hospitalizations don't explain the differences in total spending behavior by municipalities, but it might be that consumers have updated their beliefs accordingly to experience and information. As the Covid-19 pandemic progressed, more information became available in terms of how the virus is contracted and which (social) activities may increase the chances of attracting the virus. Therefore consumers may have adapted their consumer behaviours, avoiding certain sectors but refrained from complete withdrawal from economic and social activity.

We investigate this further by looking at online supermarket expenditures during the second Covid-19 wave. 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. Hence, they will be more inclined to order groceries online and get them delivered to their home address. Moreover, we know that total supermarket expenditure was not impacted by the NPIs as during both Covid-19 waves supermarkets, convenience stores and other vital food retailers were allowed to stay open. Supermarkets are also likely to be situated closely to consumers, which minimizes any disturbing spatial effects that may occur.

While the total spending on groceries (and hence online groceries) have increased substantially during the Covid-19 crisis, we measure whether differences in the Covid-19 outbreak by municipality explains the differences in online grocery expenditure with our fixed effects regression.

The dependent variable (Y_{it}) is the year on year change in total online supermarket spending by municipality i by week t, scaled by the amount of ABN Amro clients in municipality i. We run the regression on the second Covid-19 wave (see table 8). As the coefficient is positive, the results show that if there is one additional hospitalized Covid-19 cases in a municipality per 1000 inhabitants the online spending on groceries increases by over 2%. Regression (2) shows that this conclusion is robust to the addition of lagged variables of the dependent variable. Regression (3) and (4) includes the Stringency Index and therefore excludes time effects. The Stringency Index is also significant and shows that as the strictness of the NPIs increases, the amount of spending at online groceries also increases. It is interesting that, when the Stringency Index is included in the regression, the coefficient of the Covid-19 cases increases substantially. Obviously, these two parameters interact with each other; the Stringency Index moves along with the amount of infections (see figure 2. As the Stringency indicator is high there were more soft restrictions such as a maximum amount of costumers in supermarkets, more supervision on the 1.5 metre distance rules and forced disinfection of hands at the entrance. This results shows that whether people decided to buy their groceries online depends on both the NPIs and the severity of the local Covid-19 outbreak.[11]

These results suggest that the change in the within-municipality Covid-19 cases during the second wave may not have had a big impact on total consumption but did affect people's individual economic behaviour, beyond the effect of physical

---

[11]We have also ran the same regression on the first wave, and it shows the same results as during the second Covid-19 wave. We do not include this table because of space limitations

Table 8. Second Covid-19 wave - Dependent Variable: Online groceries transactions$_{it}$

Regressors	(1)	(2)	(3)	(4)
Covid-19 cases$_{it}$	2.0811**	2.0248**	6.9378***	7.1763***
	(0.9012)	(0.9012)	(0.9287)	(0.9276)
Transactions$_{it-1}$		0.0248**		0.00617***
		(0.0114)		(0.0115)
Transactions$_{it-2}$		0.0167		0.0233**
		(0.0114)		(0.0115)
Stringency$_t$			5.3330***	4.6276**
			(0.5565)	(0.05705)
N	274	274	274	274
T	30	30	30	30
FE	Both	Both	Entity	Entity
R^2	0.0007	0.0016	0.0243	0.0286

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

constrains imposed by the government's NPIs. Macro economic effects do not impact consumers' choices to buy their groceries online, but Covid-19 did.

5. Conclusion and discussion

By geo-locating transactions by municipality, we find that, during the first Covid-19 wave, 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. This suggests that people living in badly affected areas altered their economic behavior differently to people living in little affected areas. Our paper therefore supports the claim by other studies, that the behavioural response is important in explaining the decline in consumption.

But this reaction function of consumers toward the virus is not static. We do not find the same effect of the local Covid-19 outbreak on transactions during the second wave. This suggest that people may have adopted their behavior, this may have been due a certain amount of fatigue with the NPIs rules, or may have been caused by the fact that as more information on the Covid-19 virus was available, consumers could adapt their behaviour is such as way that it didn't impacted their consumer behaviour to a large degree. We did however find a significant effect from unemployment during the second Covid-19 wave, suggesting that Covid-19 may not have had a direct effect on consumer spending, it did have an indirect effect due to the rise in unemployment. More research is needed in order to understand how this relationship exactly works.

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. Also, our regression results 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 in 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, this study was done specifically using Dutch transaction data, wereas the reaction function of consumers toward the virus may be different among different cultures and countries. Moreover, the reaction of consumers may depend on the trust that the population has in the healthcare system and/or the government response toward the pandemic.

Although the vaccination programs are speeding up, and many (mainly developed) countries are relaxing NPIs, the Covid-19 pandemic has not ended yet. New strains of the virus and/or low efficacy rates for the vaccines can still prolong the pandemic. Our study incorporates a medium-term horizon, including two distinctive Covid-19 waves. But new waves could take a different shape and therefore the conclusion from this study may again change as well. Therefore we are certain that this will not be the last study on the economic effect of Covid-19.

References

[1] H. Chen, W. Qian, and Q. Wen, "The impact of the covid-19 pandemic on consumption: Learning from high frequency transaction data," working paper, SSRN.

[2] D. Bounie, Y. Camara, E. Fize, J. Galbraith, C. Landais, C. Lavest, T. Pazem, and B. Savatier, "Consumption Dynamics in the COVID Crisis: Real Time Insights from French Transaction & Bank Data," CEPR Discussion Papers 15474, C.E.P.R. Discussion Papers, Nov. 2020.

[3] N. Neuteboom, G. Kapetanios, A. Ventouri, P. Golex, and F. Ritsema, "Disentangling the effect of government restrictions and consumers' reaction function to the Covid-19 pandemic: Evidence from geo-located transactions data for the Netherlands," cepr discussion papers, C.E.P.R. Discussion Papers, Nov. 2020.

[4] R. Chetty, J. N. Friedman, N. Hendren, M. Stepner, and T. O. I. 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.

[5] L. Reisch and M. Zhao, "Behavioural economics, consumer behaviour, and consumer policy: State of the art," *Behavioural Public Policy*, vol. 1, no. 2, p. 190–206, 2017.

[6] S. Mehta, T. Saxena, and N. Purohit, "The new consumer behaviour paradigm amid covid-19: Permanent or transient?," *Journal of Health Management*, vol. 22, no. 2, pp. 291–301, 2020.

[7] C. A. Moser and P. Yared, "Pandemic lockdown: The role of government commitment," Working Paper 27062, National Bureau of Economic Research, April 2020.

[8] F. E. Alvarez, D. Argente, and F. Lippi, "A simple planning problem for covid-19 lockdown," Working Paper 26981, National Bureau of Economic Research, April 2020.

[9] M. Bognanni, D. Hanley, D. Kolliner, and K. Mitman, "Economics and Epidemics: Evidence from an Estimated Spatial Econ-SIR Model," Finance and Economics Discussion Series 2020-091, Board of Governors of the Federal Reserve System (U.S.), Oct. 2020.

[10] D. Acemoglu, V. Chernozhukov, I. Werning, and M. D. Whinston, "Optimal Targeted Lockdowns in a Multi-Group SIR Model," NBER Working Papers 27102, National Bureau of Economic Research, Inc, May 2020.

[11] N. V. van Banken, "Betalingsverkeer (https://www.bankinbeeld.nl/thema/betalingsverkeer/)."

[12] V. M. Carvalho, J. R. Garcia, S. Hansen, A. Ortiz, T. Rodrigo, J. v. R. Mora, and p. ruiz, "Tracking the covid-19 crisis with high-resolution transaction data," April 2020.

[13] R. G. B. K. A. P. T. P. S. W. E. C.-B. L. H. S. M. Thomas Hale, Noam Angrist and H. Tatlow, "A global panel database of pandemic policies (Oxford COVID-19 Government Response Tracker)," finance and economics discussion series, Nature Human Behaviour, Oct. 2021.

[14] C. van der Cruijsen and J. Knoben, "Ctrl C Ctrl Pay: Do People Mirror Payment Behaviour of Their Peers?," *SSRN Electronic Journal*, no. 611, 2018.

[15] M. H. Pesaran, "Estimation and inference in large heterogeneous panels with a multifactor error structure," *Econometrica*, vol. 74, no. 4, pp. 967–1012, 2006.

[16] A. Goolsbee and C. Syverson, "Fear, Lockdown, and Diversion: Comparing Drivers of Pandemic Economic Decline 2020," NBER Working Papers 27432, National Bureau of Economic Research, Inc, June 2020.

[17] O. Coibion, Y. Gorodnichenko, M. Weber, and M. Weber, "The Cost of the Covid-19 Crisis: Lockdowns, Macroeconomic Expectations, and Consumer Spending," tech. rep., 2020.