Tracking the COVID-19 Crisis with High-Resolution Transaction Data*

Vasco M. Carvalho
University of Cambridge,
Alan Turing Institute and CEPR.
vmcarvalho.web@gmail.com

Juan R. Garcia
BBVA Research.
juanramon.gl@bbva.com

Stephen Hansen
Imperial College Business School and CEPR.
stephen.hansen@imperial.ac.uk

Álvaro Ortiz
BBVA Research.
alvaro.ortiz@bbva.com

Tomasa Rodrigo
BBVA Research.
tomasa.rodrigo@bbva.com

José V. Rodríguez Mora
University of Edinburgh,
Alan Turing Institute and CEPR.
sevimora@gmail.com

Pep Ruiz
BBVA Research.
ruiz.aguirre@bbva.com

July 8, 2020

*We thank Raman Singh Chhina for superb research assistance, and Gary Koop for inspiring comments on a previous draft.

This is a live document and subject to ongoing changes. All analysis is preliminary. The most current version of the paper can be downloaded here.

All data has been anonymized prior to treatment and aggregated at BBVA before being shared externally.

Abstract

Financial and payments systems throughout the world generate a vast amount of naturally occurring, and digitally recorded, transaction data, but national statistical agencies mainly rely on surveys of much smaller scale for constructing official economic series. This paper considers 2.1 billion transactions from credit- and debit-card data from BBVA, the second largest bank in Spain, as an alternative source of information for measuring consumption, a key component of GDP. While card spending growth is more volatile than non-durable consumption growth, normalized spending correlates strongly with official consumption measures. In the cross section, patterns in card spending match those in official household budget surveys very closely. The implication is that card spending can stand in for consumption surveys in environments where official data is not available, for example due to reporting delays or to insufficient geographic or household detail. We apply the idea of card spending as a consumption survey to the COVID-19 crisis in Spain, where we present four findings: (1) a strong consumption reaction to lockdown and its easing at the national and regional levels; (2) a rapid, V-shaped consumption recovery in the aggregate; (3) an adjustment to the average consumption basket during lockdown towards the goods basket of low-income households; (4) a divergence in mobility patterns during lockdown according to income in which poorer households travel more during the workweek. Our main conclusion is that transaction data provides high-quality information about household consumption, which makes it a potentially important input into national statistics and research on household consumption.

1 Introduction

Every day, banks, payments systems providers, and other financial intermediaries record and store massive amounts of individual transaction records arising from the normal course of economic life. As more and more of the world's trade and exchange activity is intermediated on platforms underpinned by digital technology, transaction data is likely to continue to grow rapidly.

Until recently, though, the use of transaction data by national statistical agencies and academics has been fairly limited.[1] In part the reason is logistical, as data sharing is subject to numerous legal and regulatory requirements. The holders of transaction data also increasingly realize its commercial potential and vary in their inclination to allow its use in research and policy analysis. Still, a reasonable expectation is that economists will have access to large-scale transaction datasets in the near future. The COVID-19 pandemic of 2020 has acted as a major stimulus for movement in this direction, and in a short space of time an entire new literature has grown that uses indices derived from transaction data to track the impact of virus spread and lockdown.[2]

While attention is naturally focused on assessing the impact of the pandemic-induced economic shock in the short run, the potential areas of application of transaction data are much broader, but also require more extensive validation against known datasets to be fully convincing. Traditional economic measurement relies heavily on centrally administered, carefully-designed surveys conducted with representative subsamples of the population. In contrast, transaction data arises through the decentralized activity of millions of economic agents. How then do such data compare to national accounts? Which potential biases and distortions exist in indices built from transactions, and what additional insights can they bring? These are ultimately questions that must be answered for transaction data to enter the mainstream, and to continue to be of importance as the pandemic fades from memory.

The primary contribution of this paper is to discuss these issues in the context of the universe of credit and debit card transactions mediated by Banco Bilbao Vizcaya Argentaria, S.A (BBVA), the second-largest bank in Spain, but also with a major market presence in numerous other countries. Our overall dataset contains 6 billion transactions collected from BBVA cardholders and BBVA-operated point-of-sale terminals from seven

[1] Important exceptions include Gelman, Kariv, Shapiro, Silverman, & Tadelis (2014), Baker (2018) and Olafsson & Pagel (2018), which use data from financial apps to track household spending and income.

[2] A selection of papers, besides the April 14th 2020 draft of this paper, includes Andersen, Hansen, Johannesen, & Sheridan (2020a), Andersen, Hansen, Johannesen, & Sheridan (2020b), Alexander & Karger (2020), Baker, Farrokhnia, Meyer, Pagel, & Yannelis (2020a), Baker, Farrokhnia, Meyer, Pagel, & Yannelis (2020b), Bounie, Camara, & Galbraith (2020), Chetty, Friedman, Hendren, & Stepner (2020), Chronopoulos, Lukas, & Wilson (2020), Cox, Ganong, Noel, Vavra, Wong, Farrell, & Greig (2020), Surico, Kanzig, & Hacioglu (2020).

countries, 2.1 billion of which arise in Spain. We explore the properties of the data along three different dimensions: as a coincident indicator for aggregate and subnational consumption; as a detailed household consumption survey; and as a mobility index.

In each case, card spending captures some but not all of the relevant information in the analogous official data series, but nevertheless acts as an informative proxy along comparable cuts of data. This then allows one to make further cuts into the spending data to obtain insight unavailable using external series alone. We apply this idea to the experience of Spain during lockdown and lockdown easing, which are of independent interest to the literature on the effects of the COVID-19 crisis and its aftermath.

The first validation exercise we conduct establishes the properties of card spending for *nowcasting aggregate consumption*, the dominant component in GDP. An initial observation is that the raw data series features excess volatility relative to national accounts. Part of this derives from card spending omitting stable and significant components of consumption - such as rents, school fees, some utilities and subscription services - which tend to be settled through recurrent bank transfers. Second, over long spans of time there is likely extensive margin movements, reflecting entry and exit of clients, cards and points of sale in the BBVA data. Third, if online expenditures are included, this necessarily - because all online expenditures are card based while offline is a mix of cash and card expenditure - biases the card expenditure aggregates to a relatively higher growth/higher volatility bundle.

This not withstanding, we also show that, when rescaled, card expenditure series are a high quality coincident indicator of consumption, both in the aggregate and within narrowly defined consumption categories. Further, we show that aggregate card expenditure series are highly correlated with income at subnational levels of aggregation, such as provinces and postal codes. Taken together, this implies that card expenditure data can be usefully deployed at high frequency and high geographic resolution, allowing researchers to proxy accurately the differential evolution of consumption across narrowly defined geographical units.

We apply this high frequency proxy to document the evolution of aggregate expenditure across countries, provinces within Spain and postal codes within Madrid during the COVID-19 crisis. Both the cross-country and the cross-province analysis are informative of the large and abrupt declines in expenditure that ensued upon lockdown measures being enacted. Further, both international and subnational data suggest a V-shaped pattern of recovery of expenditure. Finally, we additionally exploit the high geographical resolution of our dataset to establish (i) how province-specific timings in lockdown-easing policies were important in explaining differences in expenditure growth across provinces and (ii) how the residents of richest zip codes (in Madrid) suffered the largest declines in expenditure.

The second comparison exercise we conduct benchmarks card spending against the official household budget survey conducted by the national statistics institute. One limitation of card data is that it only applies to a subset of purchases present in the survey: we capture 50% of total average household spending in Spain on the categories we can match. Within these matched categories, the relative spending shares on cards and in the survey correlate well. Moreover, aggregate spending shares across household covariates in the survey and in the card data line up closely, which suggests that relative spending differences in the unmatched survey categories relates strongly to the relative spending differences observes in the matched ones. These results lend confidence that the card data acts as a consumption survey, albeit within a restricted set of spending categories. We then use variation in card spending shares across Madrid postal codes to characterize the consumption baskets of the rich and the poor during normal times, and find patterns broadly consistent with the rich engaging in more market production and the poor in more home production. During the lockdown, average spending in Spain moved in the direction of the low-income consumption basket, which suggests that savings rates for high-income households increased (which is consistent with the evidence in Surico et al., 2020 and Cox et al., 2020). We also show that, in Spain, increases in online spending during lockdown are proportionally higher for cardholders with a lower online spend share in 2019, in contrast to Watanabe, Omori et al. (2020) findings for Japan.

The third comparison exercise relates to mobility, whose measurement has taken on increased prominence during the pandemic. The official series we use in this application is the Google Mobility Report for Spain, which tracks changes in mobility from users' sign-ins to Google's services. We show that changes in transportation spending on card data track changes in personal mobility at the onset and during lockdown very closely. We then disaggregate transportation spending by Madrid zip code, and document that during lockdown, cardholders in lower-income areas travel more than those in higher-income areas during the workweek, but that travel spending on the weekends is more similar across income groups. This is consistent with low-income households traveling to work during lockdown while higher-income households are more able to shelter in place due to working in occupations more amenable to telecommuting. (See Coven & Gupta, 2020 for evidence that directly uses mobility data by neighborhood in New York). We then fit a simple disease incidence model, and show that spending on urban transportation is correlated with COVID-19 cases several weeks later. This allows us to quantify the disease cost born by lower-income neighborhoods of their excess travel relative to the highest-income neighborhoods.

The main methodological contribution our paper makes is to benchmark card spending data against external series to assess its plausibility to conduct analysis of granular economic activity. Datasets like ours that consist of spending from cards and points-of-

sales terminals are currently and will likely remain one of the most commonly available transaction datasets. The comparison exercises we conduct, and the strengths and weaknesses of the data we identify, are hence more broadly relevant beyond BBVA.

The main applied contribution of the paper is to document expenditure adjustments during the COVID-19 pandemic. Relative to this large and fast-expanding literature, we encounter some common patterns. Thus, like Cox et al., 2020 and Chetty et al., 2020 in US studies, and Surico et al., 2020 for the UK we find that higher income groups witnessed the largest fall in expenditures during the crisis. Our analysis of cross-category expenditure reallocation during the crisis, echoes findings elsewhere in the literature, for example in Bounie et al., 2020 for France, Chronopoulos et al., 2020, for the UK Andersen et al. (2020a) and Alexander & Karger, 2020 for the US. Further, our analysis of the effects of lockdown and its easing complements that in Andersen et al., 2020b. The latter argue for the importance of behavioral adjustments in expenditure patterns, responding to local disease dynamics even in the absense of lockdown policies. Consistent with this, we find local disease incidence to be a driver of expenditure growth, even when controlling for different levels of lockdown restrictions across space. Unlike Andersen et al., 2020b, we are able to additionally document the significant effects of different lockdown restrictions, even when controlling for local disease incidence. Finally, like Coven & Gupta, 2020, we explore the relation between mobility and disease incidence. Relative to that contribution, we show that in the absense of direct mobility proxies, card transactions in transportation categories can be used as a mobility proxy at narrow geographical and socioeconomic status levels of analysis. ,

The rest of the paper is structured as follows. Section 2 provides background on the BBVA dataset. Section 3 provides an analysis of card expenditure as a high frequency consumption proxy. Section 4 exploits the real time consumption survey aspect of our data. Section 5 uses expenditure data and tagged geographical information of residence of cardholders to inform an analysis of mobility patterns. Section 6 concludes.

2 Background on the BBVA transaction dataset

The bulk of our analysis centers on Spanish transaction data. Our data for Spain consists of a join between (a) the universe of transactions at BBVA-operated Point of Sales (PoS) and (b) the universe of transactions by BBVA-issued credit and debit cards (in non-BBVA-owned PoS, to avoid double counting). The time stamps of transactions available to us range from the 1st of January 2019 till the 29th of June 2020.[3]

Panel A of Table 1 presents some summary statistics of this large dataset. In total,

[3]Some of the data vintages we use in the analysis below are shorter than this ending some days before the 19th. This reflects the vintage of our query to the database.

our analysis builds up from 2.1 billion card transactions, with about two thirds of the observations in 2019 and the remainder in 2020. At one end of each transaction is a Point of Sale. We observe 2 (1.6) million distinct PoS in 2019 (2020, respectively).

Table 1 additionally presents the empirical distribution of transaction values in our data. The median transaction in either year is just under 20EUR, with the overall distribution of transactions spanning three orders of magnitude, from 2EUR to 200EUR at the 5th and 95th percentile of transaction values.

Each transaction is tagged with information on whether it was carried out at an online PoS (e.g. an internet purchases) vs. offline, at a physical PoS. In this data, 30% of all 2019 PoS are online accounting for 8.4% of all transactions. Note that all online transactions are necessarily completed with a debit or credit card while offline transactions can occur via either card (which we observe) or cash (which we do not). This implies that our sample of expenditures is biased towards online expenditures.[4]

Further, for each PoS, we have a classification of the principal activity of the firm selling goods and services through that PoS. This classification breaks down the universe of transactions into 76 categories, ranging from Toy-Stores to Funeral Homes.

We are also able to distinguish whether the card initiating each transaction was issued by a Spanish bank or by a foreign bank. Throughout, we mainly focus on national card transactions, which account for 93% of the transactions in the sample. Within the sample of national card transactions we sometimes focus on the subsample of BBVA cardholders. In 2019, there are 6.3 million unique BBVA cardholders. This comprises a 16% sample of Spain's adult population of 39 million.

For BBVA cardholders we observe their home address postal code, their education level and age. Panel B.1 and B.2. compare the age structure and educational attainment of BBVA cardholders to that of Spain's adult population. Overall, our sample is broadly in line with the latter on both dimensions, somewhat undersampling the youngest and oldest in the population while oversampling the middle aged.

For some of the analysis below we are able to leverage from the international presence of BBVA and present cross-country expenditure data from Argentina, Colombia, Mexico, Peru and Turkey. Further, BBVA is present in the southern U.S. states of Alabama, Arizona and Texas, which we aggregate into "Southern U.S".

Relative to the Spanish data reviewed above, this international data is less detailed and we observe only aggregate total expenditure patterns for 2019 and 2020. These series result from the aggregation (across both years) of 3.8 billion transactions, with a country

[4]Note that, by covering only card transactions, we are unable to speak to the dynamics of expenditures backed by cash. In particular, it is unclear whether the share of transactions in cash has remained stable throughout the COVID-19 crisis. It is likely that merchants and customers switched away from cash due to fears of viral infection through bank notes and coins. If this is true, then aggregate (cash and electronic) expenditure declines during the pandemic are likely to be larger than what we document.

PANEL A: Transaction Sample Statistics		
	2019	2020
Number of Transactions	1.4 Billion	0.7 Billion
% Offline	92	95
Transaction Values		
5th Percentile	1.6€	1.9€
25th Percentile	8.5€	8.4€
50th Percentile	19.8€	19.3€
75th Percentile	45.4€	44.0€
95th Percentile	191.2€	176.6€
Number of Points of Sale	2 Million	1.6 Million
% Offline	70	65
BBVA Cardholders	6.3 Million	5.9 Million

PANEL B.1.: Age structure		
Age	Spain Population (%)	BBVA Cardholders(%)
<25	0.096	0.067
25-34	0.135	0.150
35-44	0.184	0.250
45-54	0.191	0.217
55-64	0.159	0.150
>65	0.235	0.167

PANEL B.2.: Education structure		
Education	Spain Population (%)	BBVA Cardholders (%)
Secondary or less	0.67	0.65
Tertiary	0.33	0.35

Table 1: Panel A: Summary statistics for BBVA transaction data, by year. Panel B.1.: Age structure of BBVA cardholders vs. Spain's over-18 population; Source: Instituto Nacional de Estatistica (INE). Panel B.2.: Educational attainment of BBVA cardholders vs. the Spanish population (25-64 years old); Source: OECD Education at a Glance 2014, Education Indicators.

breakdown of: 161 million individual transactions for Argentina, 82 million for Colombia, 2.8 billion for Mexico, 159 million for Peru, 359 million for Turkey, and 254 million for the southern US states.

Finally, note that expenditures are measured in nominal terms throughout and our data does not include any price-level information. Particularly for our Covid-19 applications, note that it is likely that the relevant deflators are changing substantially as the crisis unfolds.

3 Transaction data as a high frequency consumption proxy

3.1 Time Series Validation and Subnational Correlations

3.1.1 Tracking Macro and Micro Series Over Time

We start by comparing the time-series properties of our transaction data to official measures of economic activity in Spain. In our first exercise, we deploy a quarterly aggregate of the same universe of transactions reported above and compare with national account (nominal) aggregate series. This lower frequency allows us to track expenditure back to the first quarter of 2016. To account for seasonal patterns, both in our expenditure series and in the national accounts, we compute Year-on-Year (Y-o-Y henceforth) growth rates, i.e. the growth rate between the current quarter and the same quarter in the previous year. Finally, recall from our earlier discussion that online transactions are over-represented in our card series. Furthermore, in our dataset, online transactions have a significantly larger and more volatile growth rate than their offline counterpart. To avoid the bias that this may impart, our comparison to national aggregates is based on offline, physical purchases only.

We find that our measure correlates well with national accounts' "Household Domestic Final Consumption", for a time series correlation of 0.859. The correlation improves further when we compare it to "Non-Durable Household Domestic Final Consumption" for a correlation of 0.874. This is as it should be: by covering only debit and credit card transactions at PoS, we do not cover large durable purchases (e.g. the purchase of a car) via wire-transfers between bank accounts.[5] Finally, we note that the coverage of our data improves slightly over time and so do these overall correlations. Looking only at correlations computed from the first quarter of 2017 onwards, the correlations above increase to 0.889 and 0.956, respectively.

[5]Our proviso regarding online expenditures not withstanding, it is still the case that these correlations remain high when we compute total expenditure growth in BBVA, rather than just offline. They are, respectively, 0.739 and 0.863.

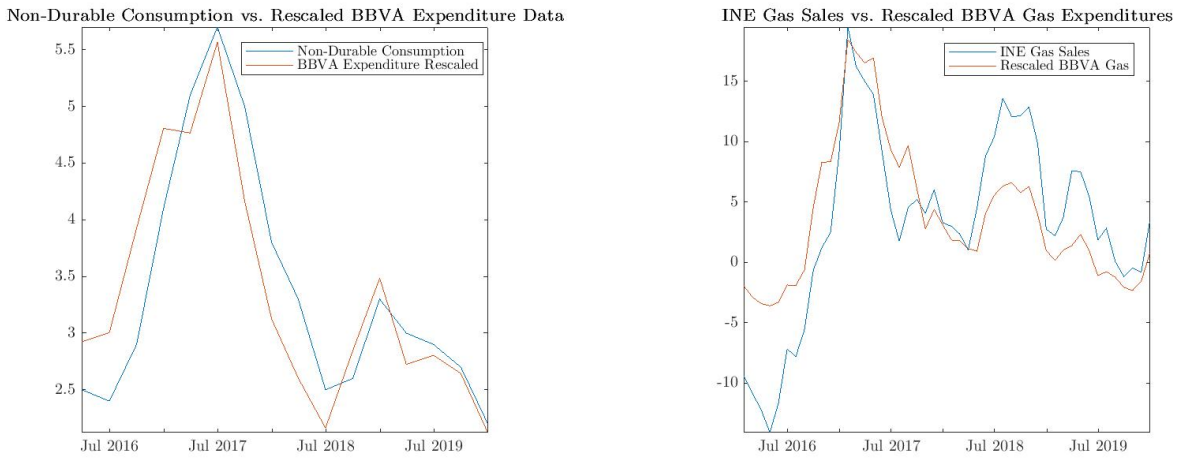


Fig. 1: Panel A (LHS): Quarterly year-on-year growth rate of BBVA aggregate offline expenditures series vs. Quarterly year-on-year growth rate of Household Non-Durable Domestic Consumption in Spain. BBVA offline expenditure series are rescaled by 0.401, the elasticity of national accounts non-durable consumption growth to BBVA offline expenditure growth. The quarterly consumption series is sourced from the Spanish national accounts. Panel B (RHS): Monthly year-on-year growth rate of BBVA expenditures at gas stations vs. monthly year-on-year growth rate of Spain's National Statistics Institute (INE) Retail Trade Index for gas stations. BBVA's gas expenditure series is rescaled by 0.346 All source data is nominal and not deseasonalized.

While highly correlated with national accounts consumption series, our offline series is nevertheless still more volatile than, say, non-durable domestic consumption. To aid interpretation of the magnitudes of expenditure adjustment presented below, we can re-express our series in implied non-durable domestic consumption growth by calculating the elasticity of growth rates across two series. To do this, we perform a simple regression of non-durable domestic consumption Year-on-Year quarterly growth on Year-on-Year quarterly growth in the BBVA expenditure data. We obtain an elasticity of 0.401 (with 95% confidence interval of [0.368,0.433]).

The reasons for this excess volatility in offline spending are at least twofold. First, by tracking only card expenditures, our data does not cover stable expenses - such as rents, school fees, some utilities and subscription services - which tend to be settled through recurrent bank transfers. Second, over long spans of time there is likely extensive margin movements, reflecting entry and exit of clients, cards and PoS in the BBVA sample. Finally, note that over longer horizons and due to growth in the bank's client base, the mean growth is also overstated relative to national accounts.

Panel A of Figure 1 plots the Y-o-Y quarterly growth of Spanish (nominal) national accounts non-durable consumption series quarterly growth rate against our nominal BBVA expenditures series, with the latter rescaled value by the above mentioned elasticity. In line with the high correlations described above, the Figure shows that our series is a good

coincident indicator for non-durable consumption growth.

Going from the macro to the micro, we can additionally validate the dynamic properties of our data against high frequency data on narrow consumption categories. In particular, in this second exercise, we build on previous work by BBVA research in Bodas, López, López, de Aguirre, Ulloa, Arias, de Dios Romero Palop, Lapaz, & Pacce (2019) which develops and benchmarks a subset of this data covering retail sales. Here, we compare the properties of expenditures at a narrowly defined sector - gas stations. We compare the dynamics of expenditure in the BBVA data relative to the highest frequency comparable index available from the Spain's National Statistics Institute (INE), covering monthly retail trade sales in gas stations.

As before we focus on (now monthly) Y-o-Y growth rates. The raw correlation between these two series is also high, at 0.784. The corresponding elasticity of growth rates across the two series is similar to that of the aggregate, at 0.346. This implies that the BBVA series is again more volatile than that compiled by INE. Panel B of Figure 1 plots the INE Gas Retail Sales series against that the corresponding BBVA series, rescaled by the estimated elasticity.

Overall, we conclude that the BBVA expenditure data is a valid proxy for consumption, both in the aggregate and at the micro level, for narrowly defined categories of consumption. Further, this coincident indicator nature of the BBVA series holds both at the quarterly and monthly frequencies. Nevertheless, the higher volatility of the BBVA series (at either level of aggregation) also suggests that care should be taken when interpreting the dynamics relative to standard consumption series.

3.1.2 Income and Expenditure Patterns with High-Resolution Geography

One possibility enabled by the geo-tagging of transaction data is to observe high frequency consumption proxies at various subnational levels of geography. This is particularly valuable in settings, such as Spain, where, say, quarterly subnational time series of consumption, are not available to researchers and policy-makers. This also implies that we cannot validate our high frequency subnational time series against officially released data. Instead, here we show that the BBVA expenditure data also correlates well with cross-sectional measures of regional income.

Throughout the rest of the paper we will exploit two different subnational regional units of aggregation in Spain. The first, more coarse, unit is the province. Spain is divided in 50 provinces and two autonomous cities (taken here as a province). This administrative unit is of particular interest in the present exercise as the policies of lockdown and its subsequent were taken using the province as the unit of implementation. For example, during the COVID-19 crisis, albeit all provinces went into lockdown on the same day, they eased it at different dates.

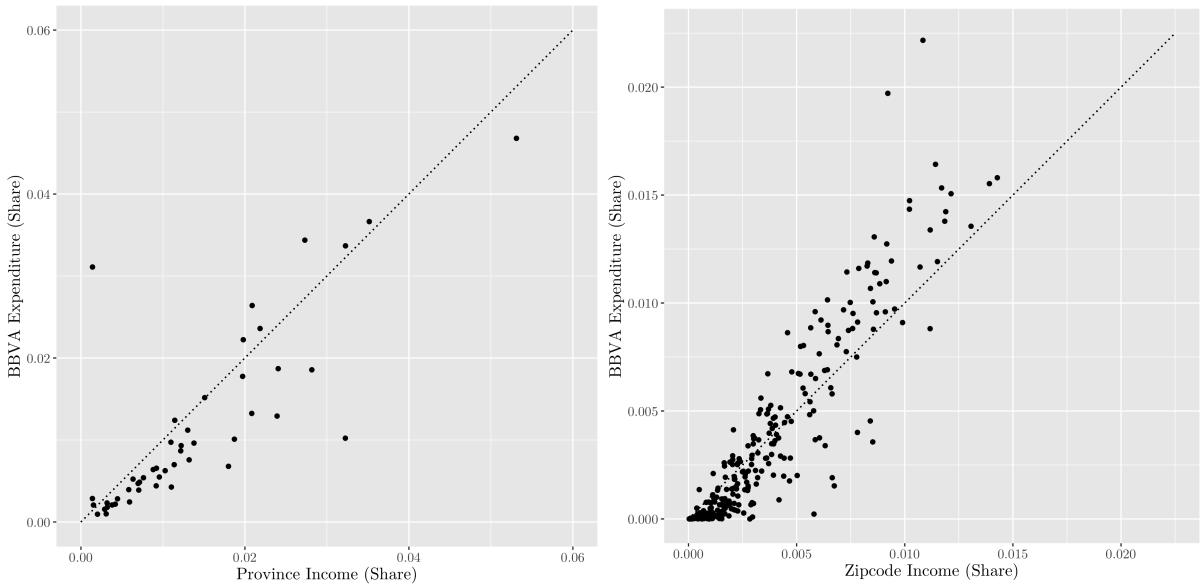


Fig. 2: Panel A (LHS): Share of income of each Province in Spanish GDP (INE) vs. the Share of Expenditure in each Province, 2019 (BBVA data) Panel A does not include, two outlying provinces, Barcelona and Madrid (the correlation is higher when we include them). Panel B (RHS): Share of income of each Madrid postal code in Madrid's aggregate income (INE) vs. the share of expenditure by BBVA cardholders residing in each Madrid postal code over total expenditures by BBVA cardholders residing in the Madrid province (BBVA data), 2019. For postal code income we aggregate up from "secciones censales" level-data from INE.

The left panel of figure 2 plots the share of total 2019 expenditure in BBVA data spent in each Spanish province against the share of that province's GDP in Spain's GDP. The latter data is sourced from Spain's national statistics institute and refers to 2018, the latest year available. To preserve readability, the plot does not display two outliers in income and expenditure shares, Madrid and Barcelona. Including these two provinces, the Pearson correlation is 0.975 (if not including them, as shown, the Pearson correlation is 0.9).

The granularity of geo-tagged transaction data also allows us to observe economic activity across more narrow spatial definitions. In particular, we explore how BBVA expenditure data correlates with activity within 5-digit postal codes in the Madrid province. Madrid postal codes are relatively homogeneous units of around 20000 individuals on average. We observe a total of 296 postal codes income in year 2017.

To obtain a measure of income at the postal code level, we build up from a granular cross-section of data available from the Spanish Statistical Office (INE) referring to "secciones censales". These are small spatial divisions (equivalent to US Census tracts) and homogeneous in size, forming groups of around 1500 individuals each. For each of these groups we know their aggregate taxable income (from tax returns of residents in

each "seccion censal").[6]

On the side of BBVA's expenditure data, we exploit the fact that we have postal code information on the place of residence of BBVA clients. This allows us to calculate the 2019 total offline expenditure by BBVA clients residing in each of the Madrid postal codes. We then compute the respective shares of offline expenditures by postal code residents in BBVA's offline aggregate Madrid expenditure (by all BBVA clients residing in the province of Madrid).

We then proceed in an analogous fashion to the province-level correlations discussed above. Thus, in the right-hand panel of figure 2 we correlate the (official) income share of postal codes in Madrid with the share of BBVA consumption expenditure by BBVA clients living in the corresponding postal code. As it is apparent they also correlate well at this level of dissagregation, for a correlation of 0.923.

Overall, we find that, in the cross-section, subnational BBVA expenditure data correlates well with the available official income data, available at either province or postal code level. Throughout the rest of the paper, we will be exploiting the dynamic aspect of this subnational expenditure data (unavailable from official data).

3.2 Tracking the COVID-19 Crisis in Real Time

3.2.1 A Global Expenditure Contraction

BBVA is a large global corporation, ranked within the 50 largest financial institutions worldwide by total assets, as of April 2020 (Standard and Poor's Global, (2020)). Through national affiliates, BBVA is present not only in Spain but also in Turkey, Mexico, Colombia, Peru, Argentina. Further, BBVA is present in the southern U.S. states of Alabama, Arizona and Texas. Here we leverage from the availability of daily card expenditure data for BBVA clients (or their affiliates' clients) in each of these countries and present a first global view of the impact of COVID-19 on expenditure patterns.

For each of these countries, we aggregate all transactions at the daily frequency and compute year-on-year (Y-o-Y henceforth) growth rates. Since we do not observe a consistent set of disaggregated expenditure categories for all countries, we focus on total expenditure reported by each BBVA affiliate. Further, for U.S. data, we can only compute Y-o-Y growth rates from the 1st of March 2020 onwards. Finally, because our data is nominal and different countries in our sample experience very different inflation rates, we present all expenditure growth Y-o-Y growth rates in percentage point deviation from their pre-March 8th mean. Our findings should therefore be interpreted as the total percentage point decline in daily Y-o-Y growth relative to the average growth observed in pre-COVID-19 times.

[6]This information is available from the web-page of the "Instituto Nacional de Estadística"

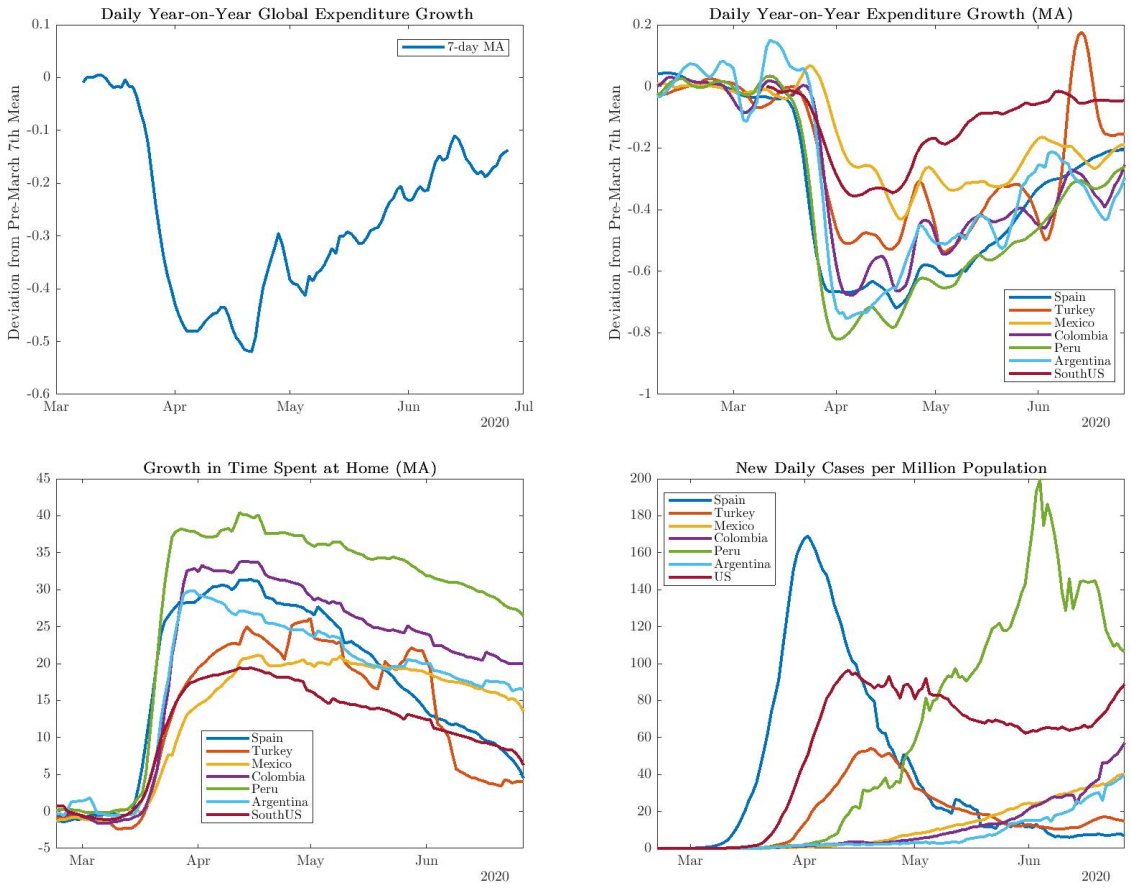


Fig. 3: Cross-Country Evolution of Expenditures, Mobility and Disease. Panel A (Top left): 7-day moving average of global index of Y-o-Y daily expenditure growth, in deviation from pre-March 8th mean growth; based on BBVA national expenditure series aggregated with GDP weights. Panel B (Top right): Cross country evolution of Y-o-Y growth of expenditure, in deviation from pre-March 8th mean growth; 7-day moving average. Panel C (Bottom left): Google Mobility Index growth in time spent at home, in deviation from pre-March 8th mean; 7-day moving average. Panel D (Bottom right): Daily new cases per million inhabitants; 7-day moving average. Source for top row graphs: BBVA; Source for bottom row graphs: DELVE (2020), Royal Society.

Figure 3 reports the results. The top left panel plots a centered 7-day moving average for a global index of Y-o-Y daily expenditure growth throughout the pandemic. To obtain it, we weight each country series with their relative weight in world GDP. In total, our data accounts for a non-negligible 8% of world GDP. [7] We observe a large and abrupt 50 percentage point decline in global Y-o-Y growth of expenditure starting mid-March. A global recovery seems to start in mid-to-late April such that by the 21st of June - the last available day in this international panel - the global Y-o-Y decline is at only -12p.p.. Through the end of June, the apparent global dynamics of expenditure suggest a V- rather than L-shape recovery.

This global picture conceals substantial cross-country heterogeneity as the top right panel of figure 3 depicts. Thus, through early May, Mexico and the southern US states saw a relatively mild decline of 30 p.p. and 20 p.p. respectively. Conversely, Peru and Spain endured very large declines: by early May expenditure growth was still 60 p.p. below that observed pre-COVID. It is also interesting to note that cross-country dynamics of recovery have since changed somewhat. While it is still the case that, by late June, Southern US states all but returned to the pre-March levels of expenditure growth, we also see fast recoveries in Turkey and Spain. In contrast, by the end of June, the slowest recoveries are in South America, with Argentina, Colombia and Peru, still 30p.p. below their pre-March growth averages.

This large decline is, at least in part, the result of decreased market interactions which, in turn, stem from lockdown and confinement orders imposed by governments worldwide. The bottom left panel reports the average growth in time spent at home in each of these countries, as derived from Google's COVID-19 Mobility reports (Google, 2020). The latter exploit accurate "Location History" metadata associated to Google account holders' logins as they move through space.

Again, we observe substantial cross-country variation. During April, Peru, Colombia and Spain report the highest average increases in the time spent at home and the Southern US, Mexico reporting the lowest. Further, during May and June, Spain and Turkey quickly return to their pre-March average. By late June, these two countries, together with the US are closest to a return to normal. In contrast, Peru, Colombia and Argentina are the countries further from it.

Overall, the joint patterns expenditure and mobility are coherent across countries: larger declines in expenditures are associated with greater home confinement. A pooled cross-country, cross-time correlation of these two measures returns a highly significant -80% correlation. For completeness, the bottom right panel of figure 3 gives one measure of disease dynamics, the number of new daily cases per million people in each of these

[7] We weight the U.S. states in our sample appropriately, based on their GDP share in the U.S. and then multiplied by the U.S. share in world GDP.

countries.[8] A pooled cross-country, cross-time correlation of these this measure with our BBVA expenditure series across countries measures now returns a -35% correlation, reflecting less coherence between disease dynamics and expenditure patterns over time and countries.

3.2.2 In and Out of Lockdown: Province-level evidence for Spain

The cross-country patterns above are suggestive of a link between lockdown policies enacted as a response to the pandemic and large expenditure adjustments on part of households. To investigate further this link, we now proceed to exploit subnational variation in the timing and duration of these lockdown policies, using Spain as a case-study.

The Spanish COVID-19 pandemic has been playing out dramatically over the last five months. The first confirmed COVID-19 infection in Spain dates from the 31st of January 2020 (in the Canary Islands). During the month of February, gradual spatial diffusion of the disease ensued such that, by the 9th of March every province in Spain reported at least one confirmed case. March was to witness the pandemic intensify throughout Spain, with 94,417 confirmed cases and 8,189 confirmed deaths by March 31st.[9].

The Government of Spain announced a State of Emergency (*"Estado de Alarma"*) on March 13th, allowing the Government to restrict mobility and activities within the country. Thus, effective March 15th, the country went into a very strict National Lockdown policy which greatly restricted mobility and commercial activity. In particular, this lockdown implied that all citizens were to stay in their residences except for food and medicines, and basic staples, work or deal with emergency situations. Further it implied the temporary shutdown of most leisure and retail spaces, such as bars, cafes, restaurants, cinemas and non-essential commercial and retail businesses. In face of rapid progression of the pandemic, this lockdown was further tightened on the 28th of March, when all non-essential activity was banned.[10] Note that, in spite of the large differences in disease incidence across provinces, this policy was implemented uniformly across the whole country.

The strict lockdown was partially relaxed on May 4th, when some commercial activity resurfaced. As we detail in Appendix A, on that date government decrees allowed for some - still restricted and socially distanced - market interactions with small retail (e.g.

[8]Note that, in the current draft, the US series refers to the country rather than the southern US states. The correlations we report in the text are robust to excluding the US states.

[9]Official numbers form the Spanish Health Ministry: *"Actualización no 61. Enfermedad por el coronavirus (COVID-19)"*

[10]As in many other countries, policy response at initial stages of this pandemic was sluggish. The first set of responses were in place by early March with localized quarantines and lockdowns of five towns and municipalities in the regions of La Rioja (Haro, 7th of March) and Catalonia (multiple municipalities, 12th of March). Between the 9th and 12th of March, multiple regional authorities proceeded to suspend all educational activities and some flight routes were also suspended.

hairdressers taking one client at a time, by appointment only and with protections in place) while keeping large commercial surfaces and superstores shut. This easing of the restrictions was implemented for the whole of Spain simultaneously, again irrespectively of the differential incidence of the pandemic across provinces. This period was deemed "Phase 0" and was the beginning of the process of easing of the lockdown which would be conductive to a smooth normalization of activity.

Thus, in the following weeks the government started to relax further the restrictions imposed by the lockdown. In contrast with the previous policy, this easing of the restrictions across 4 further consecutive "phases" was implemented in different provinces at different times taking into consideration the health and healthcare conditions in each the province.[11] These phases (labeled 1,2,3 and 4) implied successively less restrictions on economic activity. Thus, as we detail in Appendix A, Phase 1 allowed for larger retail (but not superstores and malls) to reopen at restricted capacity and for outdoors commercial activity (including restaurants) to resume. Phase 2 lifted all size restrictions on commercial activity (including malls) and some indoor commercial activities, while still keeping capacity caps in place. Phase 3, in turn, relaxed further these capacity limits somewhat. This easing process has culminated in the current Phase 4, the "new normal",[12] achieved finally on June 21st in the whole of the country. Crucially for our analysis below, note that, in the time between the first set of provinces easing into Phase 1 (May 11th) and the arrival of the "new normality", different provinces were facing different restrictions.

In order to investigate the effects of the restrictions imposed by the lockdown *and* the subsequent easing of these restrictions, we drill down from national aggregate expenditure patterns to the differential timings of lockdown easing by comparing provinces that switch between "Phases" and provinces which did not.

Figure 4 provides a first indication of the expenditure growth effects of lockdown and its gradual and uneven easing. Panel A on the top right hand side of the graph plots aggregate expenditure growth in Spain over this period, normalized by the average Y-o-Y growth[13] occurring pre-March 8th. As it is clear, expenditure growth fell abruptly on the

[11] There were some exceptions to the provincial rule. In some provinces, the lifting of restrictions was applied only to part of the province. In what follows, we have treated the whole province as if there were no change of easing status. Results are qualitatively unchanged if we apportion such provinces to an easing status.

[12] In Phase 4 there are minimal restrictions on activity beyond the imposition of minimal distance between customers and the wearing of masks in some circumstances. There are currently reversals to restrictions of mobility targeting small populations in the case of new incidences of the disease.

[13] In this graph, and in all posterior references to Y-o-Y growth, we use a slightly modified version of it that it allows us to easily control for . In order to control for weekly seasonality in the behaviour of expenditures we proceed as follows: we pair every day following January 8th, 2020 with its equivalent weekday in the equivalent week of the previous year. Thus, given that Epiphany is one of the most important holidays of the year in Spain and we exclude Y-o-Y comparison over the holiday period, we pair the first Tuesday after the Epiphany holiday in 2020 (January 8th) with the first Tuesday after Epiphany in 2019 (January 7th), and we then proceed daily, always pairing days of the week (first

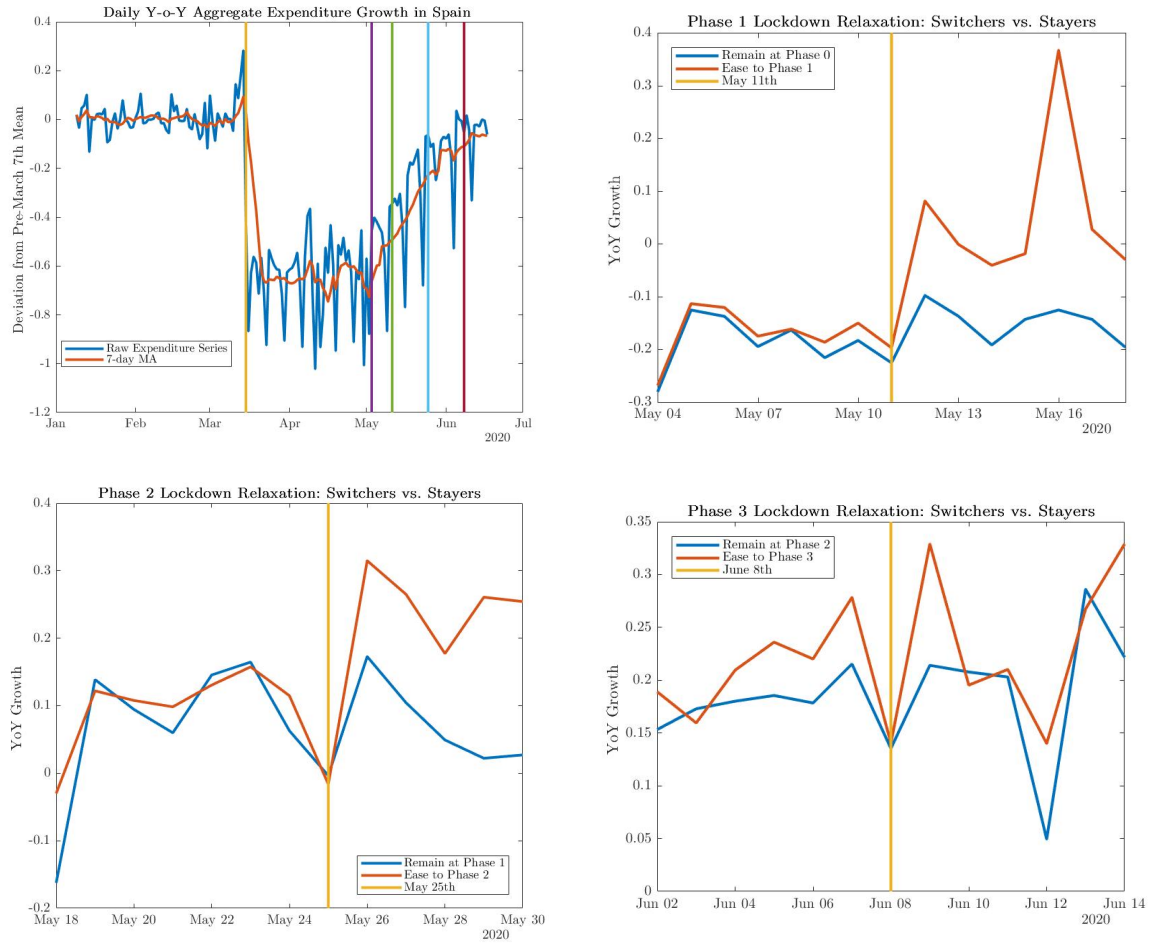


Fig. 4: Panel A (Top left) Time series of Y-o-Y growth of expenditure from BBVA series for Spain (aggregate). The vertical lines indicate the timing of events. The first (green) vertical line is drawn on March 15th, when Lockdown was implemented. The third one is May 4th (start of Phase 0), when easing started nationwide. The fourth vertical line stands for May 11th (start of Phase 1), when provinces started to differentiate in the intensity of the lockdown, some of them easing lockdown faster than others. The remaining lines are drawn on the 25th of May (start of Phase 2 for some provinces) and 8th of June (start of Phase 3). The series is normalized by the Y-o-Y growth before March 8th. Also included (in orange) is the corresponding 7-day centered moving average. Panel B (Top right) plots the average Y-o-Y expenditure growth for the provinces which eased into Phase 1 on May 11th (in orange) plotting it against the average for those provinces that stayed in the more restrictive Phase 0 (in blue). Panel C (Bottom left) and Panel D (Bottom right) plot event-study graphs centered around May 25th and June 8th, when some provinces further ease into Phase 2 and Phase 3, respectively. We plot deseasonalized data, obtained as follows: we first regress our Y-o-Y province-level growth series on a full set of day of the week dummies. We then plot event-study graphs using de-seasonalized daily expenditure growth, centered around lockdown easing announcement days.

day of lockdown, by about 60 p.p. and remained depressed at that level till early May,

Wednesday with first Wednesday, etc.). We then measure the 2019-2020 Y-o-Y growth for the same day of the week.Notice that this strategy additionally deals with the issue that 2020 is a leap year.

when easing of lockdown ensued. The aggregate data is also suggestive of a recovery starting with the nationwide enactment of "Phase 0", on May 4th. By the 21st of June, when our data ends, expenditure growth in Spain is only a few percentage points off its pre COVID-19 average, denoting a near complete recovery (in expenditure).

To further render evident the direct effect of the policy restrictions in place, Panels B, C and D, exploit the province-specific timings of lockdown easing across Spain. Thus, Panel B in the top right of figure 4 plots the average Y-o-Y expenditure growth for the provinces which eased into Phase 1 on May 11th (in orange) against the average growth for those provinces that stayed in the more restrictive Phase 0 (in blue). Plots B and C plot the corresponding event-study graphs centered around May 25th and June 8th, when some provinces further eased into Phase 2 and Phase 3, respectively.[14] Visually, there is evidence of divergence in the expenditure paths of provinces easing lockdown relative to those that do not. This differential recovery pattern is particularly evident for Phase 1 and Phase 2 switchers vs. stayers.[15] While this initial analysis provides suggestive evidence, the fact remains that different Spanish provinces are: (i) selected into treatment based, at least partly, on disease incidence and (ii) differ along a host of observable and unobservable characteristics. To address this issue we now turn to further regression analysis.

In Table 2 we present panel regressions of the daily provincial Y-o-Y growth of expenditure on phase of the lockdown and easing dummies. Throughout standard errors are clustered at the province level.

The first column gives the basic province-level time series pattern in the data, as a function of the stage of lockdown and easing. In particular, we regress province expenditure growth on a series of time dummy variables, where the omitted time category - serving as a reference- is the period before March 8th, one week before lockdown enactment. The reported coefficients can thus be read as the excess percentage point growth of average provincial expenditure, relative to pre-pandemic growth and as a function of the policy adopted at each stage of the pandemic.

It is clear that expenditures increased substantially (an average of more than 8 p.p. across provinces) in the week ahead of the lockdown, most likely in anticipation to it. The period of strict lockdown, with its associated restrictions on commercial activity has a large fall of about 60 p.p. in Y-o-Y growth of expenditure. These patterns are consistent with figure 4 where it can be ascertained that this expenditure contraction coincides with lockdown date and lasts as long as restrictions remain at their strictest level, up until

[14]To abstract from weekly seasonality patterns in the raw data, we first regress our Y-o-Y province-level growth series on a full set of day of the week dummies. We then plot event-study graphs using de-seasonalized daily expenditure growth, centered around lockdown easing announcement days. The differential patterns evident across provinces in 4 are robust to considering raw data.

[15]Since our data stops on June 21st we eschew throughout of analysing the effects of Phase 4. Before this date only a very small number of provinces were on Phase 4, leading to a lack of statistical power.

	Daily YoY Expenditure Growth by Province					
	(1)	(2)	(3)	(4)	(5)	(6)
Week Before Lockdown	0.0844***	0.111***	0.0844***	0.102***		
	(0.00837)	(0.0124)	(0.00839)	(0.00947)		
Lockdown	-0.598***	-0.570***	-0.598***	-0.580***		
	(0.0143)	(0.0190)	(0.0143)	(0.0154)		
Lockdown Easing						
Phase 0	-0.478***	-0.475***	-0.471***	-0.471***		
	(0.0186)	(0.0183)	(0.0174)	(0.0174)		
Phase 1	-0.263***	-0.262***	-0.264***	-0.262***	0.108***	0.109***
	(0.0164)	(0.0166)	(0.0159)	(0.0160)	(0.0181)	(0.0186)
Phase 2	-0.125***	-0.125***	-0.127***	-0.127***	0.210***	0.211***
	(0.0148)	(0.0148)	(0.0143)	(0.0143)	(0.0285)	(0.0295)
Phase 3	-0.0756***	-0.0763***	-0.0815***	-0.0801***	0.242***	0.245***
	(0.0248)	(0.0246)	(0.0207)	(0.0207)	(0.0394)	(0.0408)
Daily Covid Incidence		-0.2802**		-0.183***		-0.0411
		(0.1153)		(0.0494)		(0.0490)
Province F.E.	N	N	Y	Y	Y	Y
Day F.E.	N	N	N	N	Y	Y
N	8,378	8,378	8,378	8,378	8,378	8,378
R^2	0.431	0.434	0.526	0.527	0.753	0.753

Table 2: Panel regressions of daily provincial Y-o-Y growth of expenditure on phase of the lockdown and easing-date province specific dummies. Column (2) controls for daily disease incidence at the provice level. Columns (3) and (5) add provincial fixed effects and provincial and day fixed effects, respectively. Columns (4) and (6) add disease incidence controls. Standard errors are clustered at the province level. BBVA data through June 21st. Daily incidence of COVID-19 in each province obtained from the Spanish Health Ministry: here

May 4th.

Likewise, it is apparent that the initial easing of the restrictions - Phase 0, applied nationally - coincides with a sudden increase of activity. While different provinces remained at this institutional stage (and level of restrictions) for different lengths of time, the average value of Y-o-Y growth of expenditure is on average about 12 p.p. higher than in the preceding, strict lockdown, period.

The point estimates in column (1) indicate that further easing of restrictions is associated with further substantial improvements of expenditure growth, Y-o-Y growth being "only" 8 p.p. lower than its "normal" value by the time a provinces reaches Phase 3. Overall, based on these simple means, Phase 1 and Phase 2 easings on the activities of progressively larger retail spaces and hospitality (albeit still under capacity restrictions) seem to contribute the most to a strong expenditure recovery.

In columns (2), (3) and (4) of table 2 we additionally control, respectively, for differential disease dynamics across provinces, province fixed effects and their combination. Daily provincial incidence of COVID-19 (measured as the number of new cases per 1000 habitants), provides a first attempt at dealing with the basic exogeneity issue: the policy decision to ease restrictions depends on the incidence of COVID-19 at the province level, and provinces with less incidence should be expected to perform better, even in the complete absence of restrictions to activity. Consumption expenditure indeed seems affected by the incidence of the disease, even conditional on the de jure restrictions in place. Province fixed effects additionally control for systematic differences across provinces, such as in income, population density, rural/urban prevalence, which can be assumed to be fixed (or at least slowly varying) at the daily frequency. Across these specifications, the point estimates on the effects of lockdown and subsequent easing phases are essentially unchanged.

Finally, in column (5) of table 2 we present difference-in-differences estimates with province and day fixed effects and stage-of-easing-specific dummy variables. Column (6) additionally controls for the daily incidence of the pandemic at the province-level. Note that, due to the inclusion of time fixed effects our estimates are now identified out of differences in the timing of (the easing of) restrictions at the province-level. Thus, relative to the previous specifications, the omitted category is now "Phase 0", the last common policy baseline across all provinces.

The estimates we obtain are nevertheless not dissimilar from the ones obtained previously. Thus, we again observe that Phases 1 and 2 induce sizeable recoveries in expenditure growth by enlarging the set of market activities available to consumers. At the same time the intensive margin easing of capacity restrictions associated with Phase 3, does not generate a statistically significant differential effect. Further, these conclusions

are unaffected by the inclusion of province-level disease dynamics.[16]

The standard difference-in-differences analysis above exploits variation across groups of provinces that receive treatment - i.e. lockdown easings - at different times. One first concern that arises is that different provinces were on different pre-treatment expenditure trends. We address this concern by focusing in on the differential effects of Phase 1 easing, the largest point estimate obtained. Specifically, there are two groups of provinces that are of interest: the early-easers, switching to Phase 1 on May 11th vs. later easers, coming out of Phase 0 only in the subsequent weeks. We start by noting that, pre-March 8th, there is no statistically significant differential trend in expenditures across these two groups of provinces. Further, the same conclusions obtain when looking at the differential expenditure trends within the lockdown period or within the Phase 0 period, when both sets of provinces were subject to the same nationwide restrictions. [17]

A second concern that arises, as articulated in **?**, is that the treatment effect may not be stable over time. In our context, this means that the expenditure effects of lockdown easing maybe different across early and late switcher-provinces, perhaps indicating that other unobservable time-varying factors are driving the province-level response. To address this concern, we again focus on Phase 1 treatment effects. To do this, we zoom in on the period running through May 25th, by when all provinces were still in either Phase 0 or 1. Thus, within this subsample, we now have three groups of provinces: early-switchers, easing onto Phase 1 on May 11th, late-switchers on May 18th and never-switchers (till May 25th). Based on this classification we can use **?** decomposition theorem to estimate changes in Phase 1 treatment effects across different subgroups. Our estimates imply stable treatment effects. The DD estimate based on the difference between early and late switchers is 0.157. The converse estimate based on effects on late switchers vs those that had already eased previously, gives a DD estimate of 0.139. Finally, the DD estimate formed by the differential growth between ever treated and never treated gives 0.153. We conclude that, at least for the case of Phase 1, the treatment effect is stable with respect to the timing of treatment.

[16] The fact that COVID-19 incidence is not a statistically significant driver of expenditure in this period likely reflects the fact that, towards the end of our sample, disease dynamics are less heterogeneous across provinces and therefore not robust to the inclusion of day fixed effects.

[17] Early switchers daily growth during the pre-lockdown period is, on average, 1.8 percentage points higher than that of late switchers but the associated p-value is 0.195. Alternatively, taking the first ten days of May as the relevant pre-treatment different gives an insignificant 0.01 percentage point difference. Conclusions are unchanged by defining different pre-treatment periods within the joint lockdown and Phase 0 periods.

id	Category	id	Category	id	Category
1	Travel Agencies: Distance Sales & Web	27	Musical Instrument	53	Gas Stations
2	Travel Agency: Physical Location	28	Telephony	54	Parking
3	Food. Small Retail	29	DIY: Chains	55	Tolls
4	Supermarkets	30	DIY: Small Retail	56	Taxi
5	Department Stores	31	Florists: Chains	57	Sea Transport
6	Hypermarkets (Super Stores)	32	Florists: Small Retail	58	Urban Transport
7	Hotels & Lodging	33	Furniture: Chains	59	Mid- & Long-Distance Trains
8	Real State	34	Furniture: Small Retail	60	Tax and Public Administration.
9	Car Wash	35	Books	61	Miscellaneous Goods
10	Car Technichal Inspection	36	Newspapers & Magazines	62	ATM
11	Motor Vehicles Sales, Repair & Spare Parts	37	Jewelry	63	Donations
12	Bars & Coffee Shops	38	Fashion: Chains	64	Duty free
13	Fastfood & at Home Delivery	39	Fashion: Small Retail	65	Education
14	Pubs & Clubs	40	Leather Goods	66	Tobacco
15	Restaurants	41	Shoe Shops	67	Funeral Homes
16	Drugstore & Perfumes: Chains	42	Lotteries & Betting Offices	68	Phonebooths & cibercafes
17	Drugstore & Perfumes: Small Retail	43	Shows & Entertainment	69	Branch
18	Massages & Personal Care	44	Museums & Touristic Visits	70	Others
19	Beauty & Hairdressers	45	Ticket Sales	71	Mail & Parcel Delivery
20	Sports	46	Pharmacy	72	Mobile
21	Sport Equipment: Big Chains	47	Hospitals	73	Insurance
22	Toys & Sport Equipment	48	Opticians	74	Laundry & Dry Cleaning
23	Toys: Chains	49	Airline	75	Veterinary
24	Photography	50	Car rental	76	Video Clubs & TV on Demand
25	Computers, electronics & appliances: Chains	51	Boat & Airplane rental		
26	″ ″ ″ ″ ″ ″ : Small Retail	52	Mid- & Long-distance Bus Trips		

Table 3: Description of Categories of Expenditure. We mark in red those categories that were restricted during the lockdown.

4 Transaction data as a granular consumption survey

In this section, we explore how spending varies across households and good categories, both in normal times and in lockdown. The traditional source of data for measuring household consumption baskets is spending surveys, which are designed by national statistics institutes to be representative of national spending. Although BBVA card data is naturally occurring, we show below that it captures many of the salient patterns in a large-scale household survey run annually by the Spanish National Statistics Institute (INE). In this sense, card spending datasets can be viewed as a complement to traditional consumer surveys, which have many applications beyond measuring consumption during pandemics. For example, they open the possibility of measuring spending across small geographic units and in narrow categories, in addition to providing the near real-time information we emphasize above.

In addition to the postal-code-level income proxies that we discuss above, we also use education and age as household covariates in our analysis. In addition, we study category-level expenditure, where the 76 categories are defined in table 3. These categories are defined by BBVA, do not match exactly into the standard sector definitions, and contain a mix of products as well as store types. For later analysis, we have highlighted in red the expenditure categories that were restricted during lockdown by government decree.

4.1 BBVA card data as a consumption survey

We previously showed that card spending across space is strongly related to income, and that time series variation in card spending correlates with movements in national

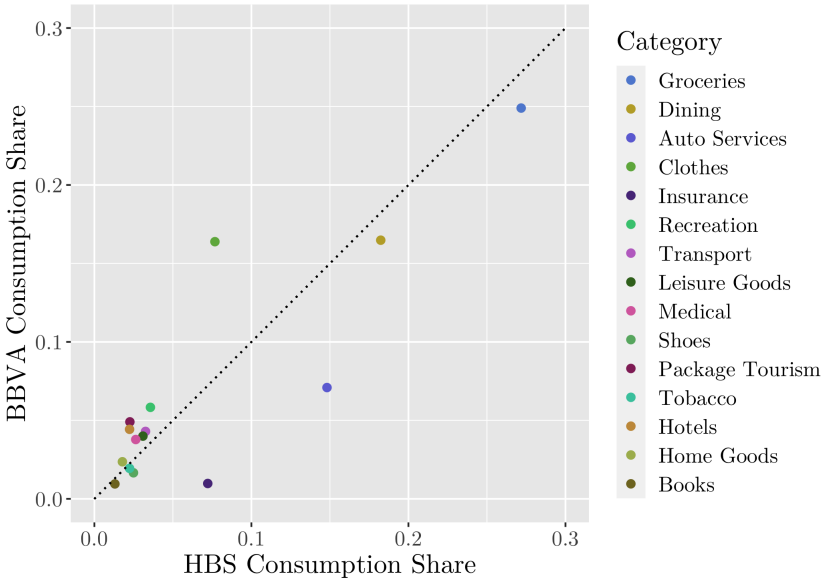


Fig. 5: Consumption Share Comparison in Matched ECOICOP Product Categories

accounting aggregates. We now explore the relationship between card spending and the annual Household Budget Survey (HBS) conducted by INE. We use its most recently available vintage from 2018 (see https://www.ine.es/en/prensa/epf_2018_en.pdf for additional details). The HBS is a national survey that draws on a sample of 24,000 households across the whole of Spain and a number of individual and household characteristics. It is designed to be representative of Spanish spending patterns, and as such presents a natural benchmark for validating the spending patterns in BBVA card data. For this exercise, we use card spending data from BBVA in 2019.

4.1.1 Validation of expenditure categories

The HBS contains spending across 40 separate good categories defined by the European Classification of Individual Consumption by Purpose (ECOICOP). Some of these are in theory not present in card spending (e.g. imputed rental value of owner-occupied housing), others are in practice not present in card spending (utility bill payments), while others have an ambiguous relationship with the categories in table 3 (non-alcoholic beverages). We are able to match 15 categories of ECOICOP with categories from table 3 with some confidence, which in total make up 48.2% of HBS spending in 2018[18] and 65% of BBVA card spending in 2019.

Figure 5 plots the shares in the matched ECOICOP categories, where we have recomputed shares in both datasets so that they sum to one within the set of matched categories. Overall, there is a strong relationship between both measures of spending:

[18]Imputed rental values, actual rental payments, car purchasing, and utility bills make up another 34% of national spending.

		National Statistics	Card Data
Age	<30	0.210	0.207
	30-44	0.269	0.283
	45-64	0.300	0.3304
	65+	0.225	0.206
Education	Secondary and below	0.673	0.637
	Tertiary & Above	0.327	0.363

Table 4: Consumption Share Comparison by Demographics

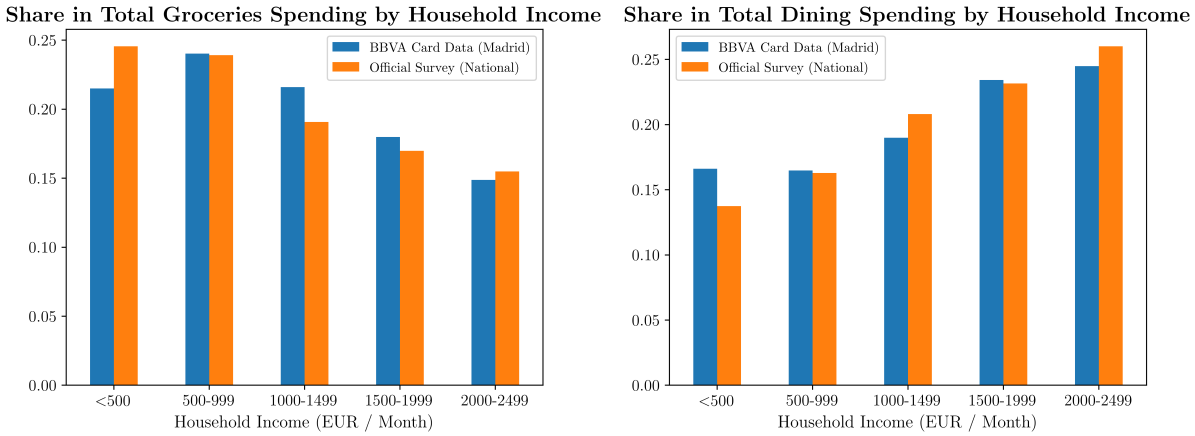
the correlation coefficient is 0.865. Moreover this is not entirely driven by the large consumption categories. When we restrict attention to categories that have a (re-computed) HBS share of less than 5%, the correlation is 0.726. For individual categories, one can observe an over- or under-representation of card spending relative to the official survey. This could be driven either by a non-representative selection of Spanish nationals into the pool of BBVA clients, or else a category-specific preference for card payments versus other forms of payments. Still, the overall message is a strong concordance between BBVA and official survey spending across categories.

4.1.2 Validation of expenditure across household covariates

We now ask whether spending on BBVA cards and official data are comparable across households. The HBS breaks down household spending in four age categories: under 30 years old; between 30-44; between 45-64; and above 65. Here the age corresponds to that of the main earner in the household. Against this we compare spending per age group from BBVA data, divided by the total number of unique cardholders within each age group (breakdown given in data description section above).[19] The HBS also provides a breakdown of spending by education, which we again divide into two categories to yield a mapping into BBVA's education categories.

Table 4 contains total spending shares across the household covariates from our two datasets. There is again a tight relationship between spending patterns in BBVA data and in the HBS. *This is in spite of the two sources of spending being defined on different sets of goods.* This suggests that not only is card data at the household level a good representation of spending, but that heterogeneity across households in omitted spending categories is very similar to the heterogeneity across households that we observe in BBVA categories, at least in terms of total spending. In other words, when a household of particular age and education structure spends more on credit and debit cards, they also appear to spend proportionally more on housing services, utilities, etc.

[19]Our BBVA data has age categories for under 25 and for 25-34. To create a match with HBS, we allocate half the spending and cardholders from the 25-34 to the under 30 category and half to the 30-44 category.



(a) Groceries

(b) Dining

Fig. 6: Comparison of Consumption Shares by Income in Food Categories

4.1.3 Household spending and income

Next we turn to explore spending by income in BBVA data and in the HBS. The HBS records household spending by numerous income groupings, measured as net household income per month. As discussed above, our main income proxy is the income of postal codes within Madrid.[20] To group postal codes into the same income bins as appear in the HBS, we divide annual income per capita at postal code level by twelve. Within each income grouping (groups of Madrid postal codes and households in the HBS), we compute the share of spending across the same 15 ECOICOP categories as appear in figure 5.[21] In general, the within-income-group correlation in consumption shares remains very high: it ranges from 0.83 to 0.95.

Figure 6 takes the two largest matched spending categories—grocery and dining spending—and compares the share of different income groups in the total consumption of both. Both datasets capture very similar spending patterns with respect to income, namely that poorer households make relatively more grocery purchases, and richer household spend relatively more in restaurants. The levels of these shares are also comparable across the two datasets in spite of their being three potential sources of divergence in the BBVA series: (1) it only applies to Madrid rather than Spain as a whole;[22] (2) it comes

[20] Given the size and economic importance of the Madrid region, and the fact that it is one of the areas of Spain with higher incidence of the pandemic (it is the region with the highest absolute number of cases, and close to it in relative numbers), we have opted to concentrate our attention to this region.

[21] The HBS includes net-income-per-month categories—corresponding to income above 2,500 EUR / month—that lie above the maximum average monthly income per capita in Madrid zip codes.

[22] In some categories, one can observe a divergence between the two series arguably related to Madrid not being representative of Spain as a whole. For example, in the HBS auto services spending is increasing in income, but decreasing in income in Madrid. One explanation is that in Madrid, higher-income households are more concentrated in high-density areas that are well-served by public transportation and taxis.

High-Income Categories		Low-Income Categories	
Category	Corr. with Income	Category	Corr. with Income
Taxi	0.67	Gas Stations	-0.48
Sports	0.62	Supermarkets	-0.35
Beauty & Hairdressers	0.58	Car Technical Inspection	-0.35
Restaurants	0.58	Telephony	-0.26
Parking	0.53	DIY: Small Retail	-0.25
Fashion: Small Retail	0.42	Insurance	-0.25
Mid- & Long-Distance Trains	0.41	Tobacco	-0.23
Pharmacy	0.40	Auto Sales/Repair/Parts	-0.23
Travel Agency: Physical Location	0.38	Veterinary	-0.22
Bars & Coffee Shops	0.37	Miscellaneous	-0.18

Table 5: Categories more positively and negatively correlated with average income across Madrid postal codes. In red, categories restricted during the lockdown.

from card spending rather than survey responses; (3) the postal code income measure is not directly observed but constructed from neighborhood-level income units. The results on income thus not only validate the use of card spending as a consumption proxy, but also our income measure.

So far, we have shown results for BBVA data that can also be obtained from a representative consumption survey. But one of the advantages of the card spending dataset is that it allows for splits over space, time, and consumption baskets that cannot be obtained through national statistics. Before turning to the COVID-19 application of this idea, we illustrate how income affects consumption across Madrid. To do so we compute consumption shares in each of the categories listed in table 3 across Madrid postal codes, and then correlate the share of individual categories with income per capita. Table 5 lists the ten categories with the highest and lowest correlation with income.

Across Madrid, higher-income postal codes are associated with more spending on food and drink outside the home, health and wellbeing, travel, and time-efficient transportation (taxis and parking lots). Lower-income postal codes are associated with spending on car-related categories, home production of food and household maintenance (supermarkets and DIY), and consumption of tobacco. This provides an insight into how income difference translate into trade-offs between time and money, investments in personal health, and access to leisure and entertainment. In this sense, the card data also doubles as a time-use survey given sufficiently rich categories of expenditure and the ability to track spending across households. Given the difficulty of collecting representative time-use data, this is another important potential use of card data.

Furthermore, the list of high and low income categories, when looked in relation to the restrictions imposed by the lockdown (table 3) is very suggestive of the consequences for the spending patterns of both groups during the lockdown. We now turn to that.[23]

[23]In Appendix B we show (1) that online spending is correlated with both income and age, (2) that in Spain online spending share during the lockdown was negatively correlated with previous usage of online

4.2 Composition of Consumption in the Lockdown

Having established the adequacy of BBVA data to study consumption patterns, in this section we analyze the change of patterns of consumption across categories and by income (postal code) groups during the lockdown and easing process. We aim to understand how the distribution of consumption across the different goods and services changed in response to the restrictions of movement and trade; and how this adjustment has been different for different income groups. This, in turn, will help understand (in the following section) the manner in which *aggregate* consumption has been affected across different income groups.

4.2.1 Aggregate composition adjustment across categories

We begin our analysis by examining the evolution of the market shares of 18 broad categories,[24] aggregating the categories presented in table 3. This is described in figure 7, where we plot the evolution of the shares in *offline spending* for these categories.[25] The categories are ranked top to bottom by their expenditure shares prior to the crisis (i.e. through 8 March).

Notice that these shares are quite stable up until the week preceding the national lockdown, when a clear re-allocation pattern emerges: spending on food and in hypermarkets (large surface supermarkets) grows considerably, and these two sectors alone make up over half of all expenditure by late March. At the same time, other sectors (such as fashion and leisure and entertainment) collapse entirely. Moreover, in the same manner that aggregate spending recovered fast once the restrictions begin to be eased, the composition of consumption returns steadily to pre-lockdown allocations following the entry in the "phase 0" of the easing period, on May 4th.

In order to understand the change in patterns of consumption during the lockdown, we return to the finer 77-categories classification, and proceed to identify the expenditure categories most altered by the crisis.

We order the categories by the relative change in their average market share, defined by comparing the average share before March 8th with the average share after March 14th and before lockdown easing. They are identified in Table[26] 6 along with the growth

spending , and (3) that the lockdown implied a large substitution towards online spending but that this pattern corrected itself extremely fast once the easing policies took place.

[24] Like the classification shown in table 3, this more coarse classification is determined by BBVA for its own purposes, and it is not perfectly aligned with standard classifications

[25] The evolution of online does not show any dramatic change in composition. Aggregate shares move by the composition of the changes in offline spending and by the changes of the overall share of online spending.

[26] From the 77 categories in which BBVA divides the data we have further eliminated the category ATM.

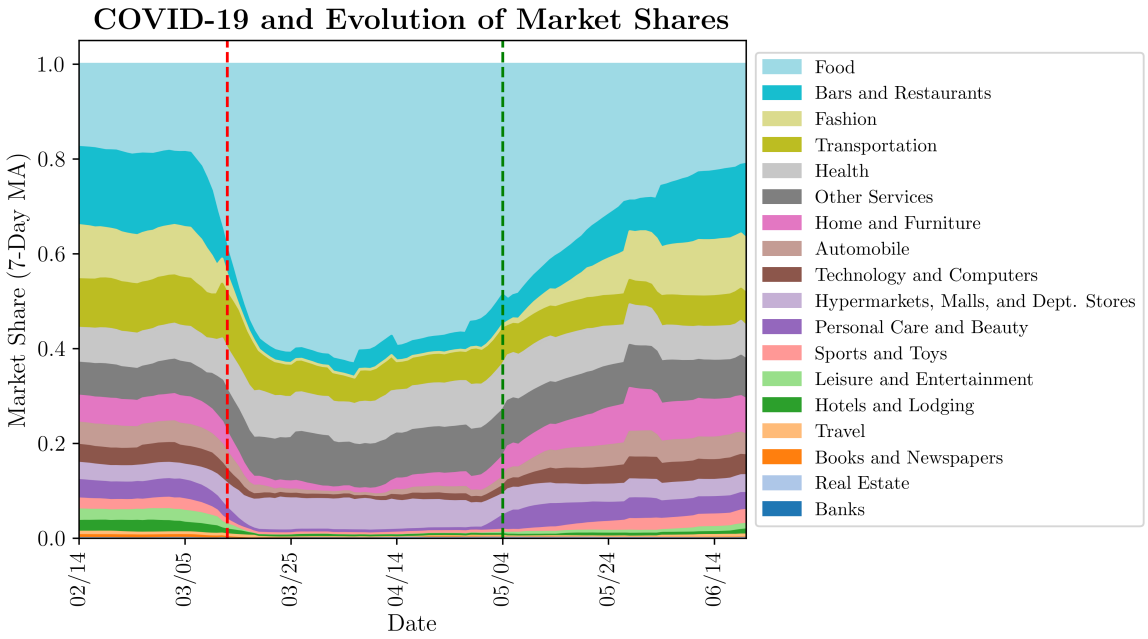


Fig. 7: Evolution of the share of Offline Spending across categories.

Top 10 Sectors in Market Share Growth		Bottom 10 Sectors in Market Share Growth	
(decreasing order of gain)	*Growth.*	(decreasing order of loss)	*Growth.*
Food: Small Stores	2.24853	Fashion	-0.97797
Tobacco Store	2.22432	Pubs and Disco Clubs	-0.93504
Mobile Phone Credit	2.06751	Furniture and Decoration Chains	-0.932594
Supermarkets	1.98371	Leather Shops	-0.93121
Hypermarkets	1.67307	Shoe Shops	-0.928647
Pharmacy and Parapharmacy	1.52951	Toys: Chains	-0.920665
Gifts and Donations	1.12815	Massage and personal Care	-0.894873
Insurance	0.835929	Fashion: small shops	-0.892908
Veterinary and pets	0.719036	Restaurants	-0.883958
Newspapers and Press	0.668963	Automobile Inspection (ITV)	-0.871738

Table 6: Best and Worst performing categories of expenditure by market share post-lockdown growth. In red, categories restricted during the lockdown.

rate of their average share between the periods. Like in the previous category tables, we mark in red those categories directly restricted by the lockdown.

Notice that, as expected, the expenditure categories that suffered most from the lockdown are all in red (being directly forbidden during the lockdown), while none of the ones who performed better were directly restricted. The categories that collapsed are those that either were essentially closed by direct imposition during the State of Alarm (such as Pubs, bars or restaurants), sell goods of scarce utility during the lockdown period (such as leather goods or fashion), or are personal services, such as Massages, of impossible implementation.

The goods and services that coped better in the new circumstances are those attending to basic necessities (such as food), or that cater goods with very low demand elasticity

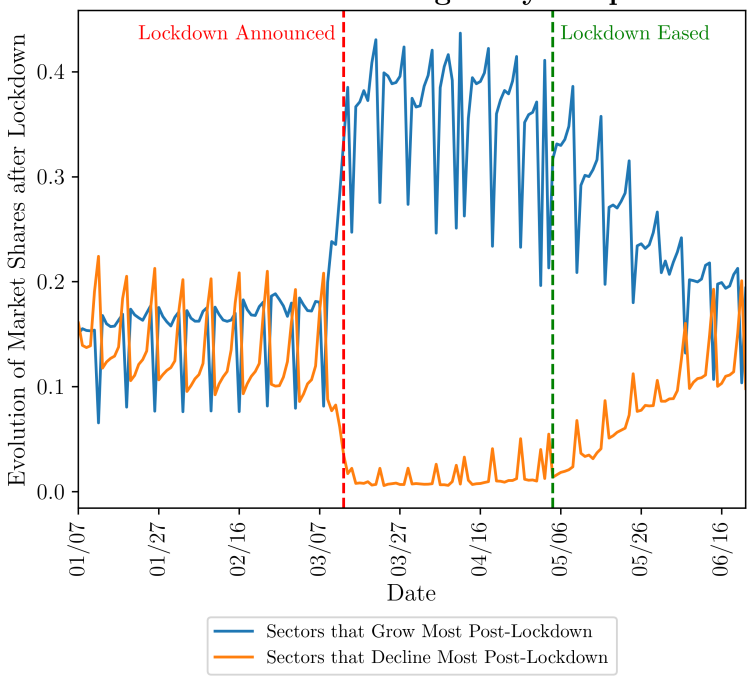


Fig. 8: Evolution of the market shares of the sectors that increased their share the most during the lockdown, and those that decreased their share the most.

(such as Tobacco). For obvious reasons, all of them were deemed critical sectors, and remained open for business during the lockdown (with restrictions on capacity usage, minimal distance between customers and so on). In addition there are categories supplying services to the business industry and that due to them being classified as "strategic" faced few restrictions of activity in the first phase of the lockdown.[27]

Looking at the aggregate evolution of these two sets of expenditure categories is illustrative of the dynamics of the crisis. In figure 8 we show that the evolution of market shares for the two sets of goods. In normal times, expenditure across the two sets is highly negatively correlated. The sectors that grow (decline) post-lockdown are consumed in relatively higher amount during weekdays (weekends), which again re-enforces the distinction between necessities and leisure consumption. Both make up roughly 20% market share prior to the lockdown. During the lockdown, the market share of the best performing categories to an average value of 60%, while the worst performing categories

[27]It is interesting to note that the expenditure category that improved most are small food shops, not only its share has risen even more than that of its larger competitors, Supermarkets and Superstores ('Hipermercados'). This is likely a consequence of the restrictions to movement. Proximity to the customer is of key importance during the lockdown, and by their very nature, small shops and convenience stores do compete favorably versus large sellers that are more sparsely located.In addition a movement away from cash and towards electronic payments during this period may result in a bias towards small shops, as they seem more likely to have depended more on cash previous to the lockdown.We thank Jonathan Thomas for point this out

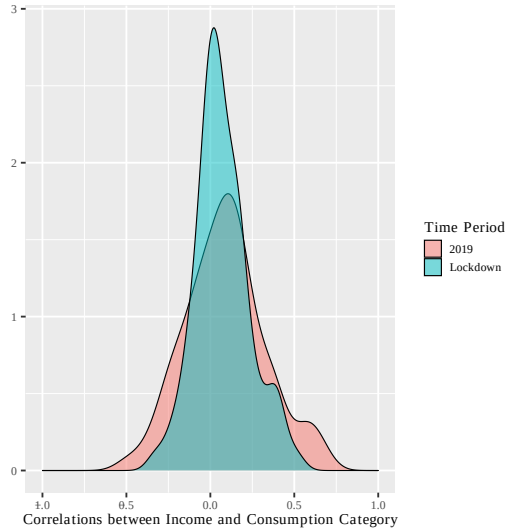


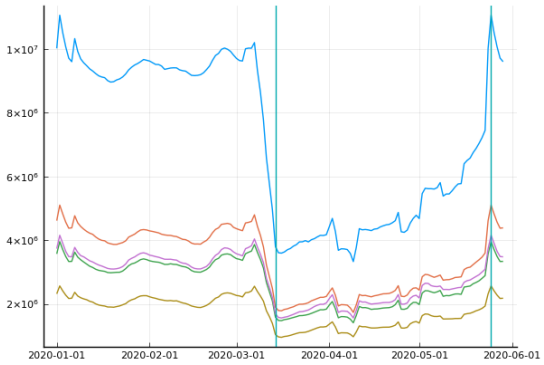
Fig. 9: Kernel distribution of the correlation between income and consumption share across goods. For a given time period the graph at a certain correlation level represents the density of sectors with such correlation.

make up on average just 1.6% of consumption. Again, these patterns quickly reverse after lockdown easing and have now essentially returned to pre-lockdown patterns.

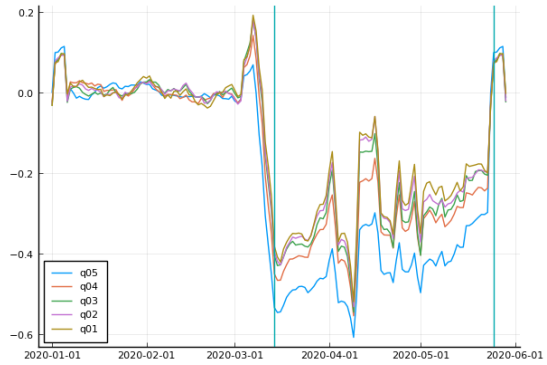
We now explore how the consumption patterns of different income groups were affected by these restrictions. Notice in table 5 that the categories more positively correlated with postal-code income are much more likely to face restrictions than the categories very negatively related to income. In figure 9 we plot the kernel densities of the correlations between income and each of the 77 BBVA categories both in 2019 and during the lockdown. If a category is highly correlated with income (either positively or negatively) income is a determining vector in the consumption of such category. The flatter the curve is, the more information has income on the consumption distribution, as more categories are mostly consumed by either rich or poor agents. The more than the curve is concentrated in zero, the less information has income on the consumption pattern of agents, as it implies a larger share of categories for which income is not a important predictive driver of consumption.

During the lockdown the kernel distribution of the correlations is clearly more concentrated around zero, implying that income is less determinant to explain the composition of consumption of the agents. Moreover, we have computed the change in the Spanish national consumption shares from 2019 to lockdown, and found a negative correlation (-0.34) between the change in a sector's national consumption share and its correlation with income in 2019.

Thus, it is difficult to escape from the conclusion that (i) not only during the lockdown the patterns of consumption of rich and poor became more similar, but also (ii) that they become more similar to the one of the poor. The reason is simple: imposing restrictions



(a) Expenditure levels in Madrid's postal codes during 2020, by quintile of postal code average income.



(b) Year on Year growth rate of expenditure in Madrid's postal codes during 2020 by postal code average income (in quintiles). Normalized by the average Y-o-Y growth before 08/03/2020

Fig. 10: Expenditure across Madrid postal codes during the pandemic.

on access to luxuries alters the consumption pattern of those who consume luxuries much more than to those who don't; and it is not the poor who is characterized by consuming luxuries. By reallocating consumption away from luxuries and towards necessities the representative consumer during the lockdown looked more similar to a relatively poor consumer's consumption basket (on a weekday) before the lockdown.

4.3 Dynamics of Aggregate Consumption across Income Groups during the Lockdown

Our objective here is to learn the manner in which economic conditions (measured as differences in average income across postal districts of Madrid), and/or differences in the incidence of the pandemic across districts affect the behavior of *aggregate* expenditures within these micro-areas of Madrid.

We start by plotting how aggregate spending patterns have been affected during the lockdown to different groups. In the same manner as in the rest of the paper, we use the average income of postal codes as a proxy of the economic conditions of the inhabitants within it. For the remainder of this section, we divide the postal codes of Madrid in quintiles according to their average income per capita in 2017. Unlike Spanish provinces, there was no different institutional treatment across these distinct postal codes. They differ, though, in their socioeconomic composition and in the degree in which they have been affected by the pandemic.

In figure 10 (analogous to figures 3 and 4 for the current level of aggregation) we show the evolution of expenditures during 2019 and 2020 for the postal codes grouped by the five average income quintiles. In figure 10a we present the evolution of the level of the

absolute level expenditure in those five groups during 2020 (in a 7 day moving average). Panel 10b plots a moving average of the evolution of Y-o-Y growth of expenditure during 2020 for each group of districts The two vertical lines indicate (i) the lockdown day (March 15th) and (ii) the day Madrid went into phase 1 of the easing process (May 25th); three weeks after the first (mild) relaxation of the lockdown, and two weeks after the first group of provinces easing into phase "1".

The figures reveal the already familiar pattern of an abrupt fall in Y-o-Y expenditure growth upon lockdown, and the sharp recovery consequent upon implementation of the easing the restrictions: it happened in Madrid as well as in the rest of the country. Our emphasis here, though, resides in the fact that different groups of income were affected differently.

The fall in total expenditures is larger in the richer neighborhoods than in the poorer ones. In a sense that is not surprising, as they started from a higher value. Notice, nevertheless, in panel 10b that the *relative* fall of the expenditure in the richer neighborhoods is also larger than in the poorer ones.

As we have seen, this is easily explained by the fact that the restrictions affected more to the patterns of consumption of the rich than to the poor. The movement away from luxuries and towards necessities affects much more those who spend in luxuries. It has consequential implications: conditional on similar relative changes in income (on which we have no data) we should expect the savings rate of the better of to have increased (and perhaps substantially) with respect to the poorer. Notice that as a consequence the bulk of the fall in expenditures corresponds to the decreases in purchases in better-off neighborhoods. Notice also that the differences across income groups are also manifest in the week previous to the lockdown, when the richer neighborhoods have a substantially smaller increase in Y-o-Y expenditure growth

Finally, notice that (as in the rest of Spain) the recovery on expenditures upon the change in regulation is in Madrid also large and sudden. Moreover, it affects all income groups largely in the same manner.

In order to quantify these effects and to include the possibility that different areas were affected by the pandemic to different degree we turn next to a more formal econometric analysis.

We will sometimes control for daily disease incidence at the postal code level. To this end, we obtain data at the level of health districts in Madrid. The Health authorities of the Autonomous Community of Madrid divide the region in 286 Health Districts of approximately uniform size as their basic unit for the provision of health services, and they report the daily incidence of the pandemic in each of those districts. To account for the differential incidence of the pandemic across the geography of Madrid we use the geographic position of health districts and postal codes to calculate and impute the daily

	Daily Y-o-Y Expenditure Growth at each Postal Code					
	(1)	(2)	(3)	(4)	(5)	(6)
New Confirmed Cases per capita	-6.841**		-6.853**	-3.432*		-3.352*
	(2.411)		(2.520)	(1.660)		(1.675)
quintile 02		0.013	0.014			
		(0.055)	(0.055)			
quintile 03		0.008	0.008			
		(0.055)	(0.055)			
quintile 04		-0.125**	-0.125**			
		(0.041)	(0.040)			
quintile 05		0.075	0.077			
		(0.165)	(0.165)			
lockdown & quintile 02		-0.135	-0.139		-0.109	-0.111
		(0.098)	(0.097)		(0.101)	(0.100)
lockdown & quintile 03		-0.235*	-0.235*		-0.193	-0.194
		(0.099)	(0.098)		(0.104)	(0.104)
lockdown & quintile 04		-0.260**	-0.262**		-0.264**	-0.265**
		(0.082)	(0.081)		(0.085)	(0.085)
lockdown & quintile 05		-0.310**	-0.302**		-0.293*	-0.289*
		(0.113)	(0.112)		(0.119)	(0.118)
day F.E.	Y	Y	Y	Y	Y	Y
Postal Code F.E.	N	N	N	Y	Y	Y
N	39,312	39,312	39,312	39,312	39,312	39,312
R^2	0.018	0.019	0.020	0.120	0.120	0.121

Table 7: Regression of Madrid postal code daily Y-o-Y growth rates on lockdown dummy variable, daily COVID-19 incidence per capita in each postal code, Income quintile of postal code and interactions with lockdown period. Standard errors clustered at the Madrid postal code level.

incidence of confirmed COVID-19 cases within the different postal codes.[28]

In table 7 we run a set of panel regressions of Y-o-Y expenditure growth across the postal districts on measures of income and incidence of COVID-19 within the district. These regressions not only confirm the impression that wealthier neighborhoods were the ones experiencing the largest fall in expenditure. They also suggest that, in addition, areas more affected by the pandemic suffer larger declines in expenditure.

All regressions include day fixed effects, thus controlling for the effects of the common policies: differently than in section 3.2.2, the lockdown and easing policies both took place at the same time in all postal codes, our unit of analysis here. In column (1) we correlate contemporaneous new confirmed cases with the change in aggregate expenditure.[29] In column (4) we add fixed effects for each postal code. We obtain a statistically significant negative correlation, even when controlling by postal district fixed effects. Thus, more affected locations have suffered a more substantial drop in expenditures independently of the policy in place (lockdown/easing).

In columns (2) and (5) we return to considering the differential effects of the lockdown on richer and poorer neighborhoods. In column (2) we include only day fixed effects, while in column (5) we include also postal code fixed effects (thus voiding the use of quintile dummies when not interacted with the lockdown). Finally, in columns (3) and (6) we include both the income of the district and the degree in which it is affected by the pandemic, controlling only for day fixed effects in the former, and adding postal code fix effects in the later.

Thus, the visual finding that in wealthier districts the decline in expenditure during the pandemic has been steeper seems robust, and this is the case independently of the incidence of the pandemic. This should not be surprising given the result of the previous

[28]The data can be obtained from: https://www.comunidad.madrid/servicios/salud/2019-nuevo-coronavirus. There are some technical caveats. We have information on disease incidence for health districts, while we have information on expenses from BBVA by postal code, and we have socioeconomic information at "sección censal" level. Unfortunately the three levels do not have a perfect match, but we have detailed geo-location information of the three levels, so we can place them in the map exactly. To merge the three sources of data we have used the following procedure: (i) The smallest in size of the three units is by far the "seccion censal", which consists of very homogeneous divisions of around 1500 individuals. Postal codes and health districts are larger, and of comparable sizes. (ii) We calculate the socioeconomic status of each postal code by merging the information of all the "secciones censales" that are completely included within the postal code. (iii) In order to attribute COVID-19 Incidence to each postal code, we assume that incidence is uniformly distributed across the specific health district, and impute to each "seccion censal" within the health district its proportional share. We then sum the imputed COVID-19 incidence of the "secciones censales" that are within each postal code to determine the degree of incidence within it. An additional issue is that the reported number is not the daily incidence, but the accumulated one the previous 14 days (or aggregated) and there seem to be revisions of the data when cases are diagnosed incorrectly, etc. We calculate daily incidence as the difference between the reported accumulated incidence one day and the one reported the previous day.

[29]In table 11, in appendix C, we present the same table using as a measure of the incidence in of COVID-19 in the district not the daily cases but the aggregate total at the end of the period. It produces equivalent results.

section. The goods and services more restricted during the pandemic are those more closely associated to the consumption of richer individuals: one needs to spend in luxuries in order to stop consuming them.

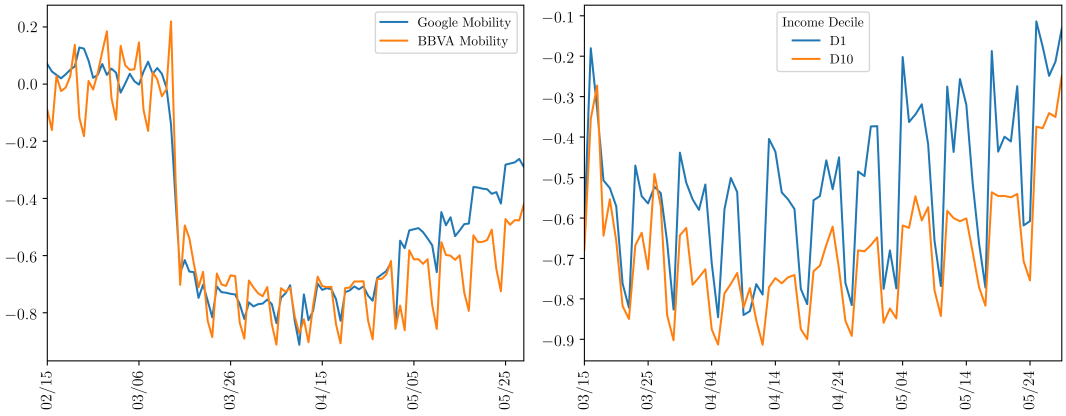
Likewise, incidence of the pandemic by itself has a detrimental effect on expenditure even after *conditioning on the level of income an on the lockdown policy implemented.* It seems intuitive that larger incidence in a neighborhood would result in lower expenditure, as direct knowledge of people suffering the illness is likely to increase one's precautions and to increase the degree of self-isolation, thus decreasing spending. And these effects would be *above and beyond* the ones induced by the province-level lockdown policies. Nevertheless, this is not proof of causality, as we cannot do away with the possibility that incidence is acting as a covariate for unobserved characteristics than drive consumption adjustment; albeit it would have to be orthogonal to our income proxies.

Overall, we conclude that both disease incidence and socio-economic status were important drivers of expenditure adjustment during lockdown.

5 Transaction Data as a Real-Time Mobility Proxy

The final aspect of information that we focus on from card spending is mobility patterns. Mobility and its determinants have become major issues during the COVID-19 pandemic due to the control of movement being a key goal of social distancing policies (see for example Allcott, Boxell, Conway, Gentzkow, Thaler, & Yang, 2020 and Simonov, Sacher, Dubé, & Biswas, 2020), but mobility studies typically rely on data captured from users' mobile phones. In countries like the USA, this data is available at fairly dis-aggregated spatial units and also contains information on user characteristics. In other countries, such data is much rarer and so alternative mobility proxies are important to find.

We define a card-expenditure-based mobility proxy by considering BBVA spending categories from table 3 that relate directly to transportation: categories 52, 53, 54, 55, 56, 58, 59. The idea is that when individuals spend money on these categories, they do so intending to move in space which should be correlated with movement captured by mobile-phone-based measures. We cannot control for whether this spending is for transport on the day the purchase is made (if for example households purchase tickets in advance), and so validating this series against external sources of mobility data is important. We do so using aggregate Spanish mobility data from the Google Mobility Reports, which as explained above draw on mobile phone location data. Figure 3 introduces the Report's measure of time spent at home, but in addition there are five other categories that measure time spent outside the home (time at work, time at transit stations, etc.). We create a single mobility number from the Reports by averaging at daily frequency the growth in time spent in these outside-the-home activities. Against this we plot the change in



(a) Transport Spending vs Mobility (b) Transport Spending and Income

Fig. 11: Left panel: comparison of Google Mobiliy Report for Spain against BBVA card data spending on transportation subcategories. The baseline for computing growth for the BBVA series is the spending average from 1 January 2020 through 14 February 2020. Right panel: change in transport spending by postal code decile in Madrid.

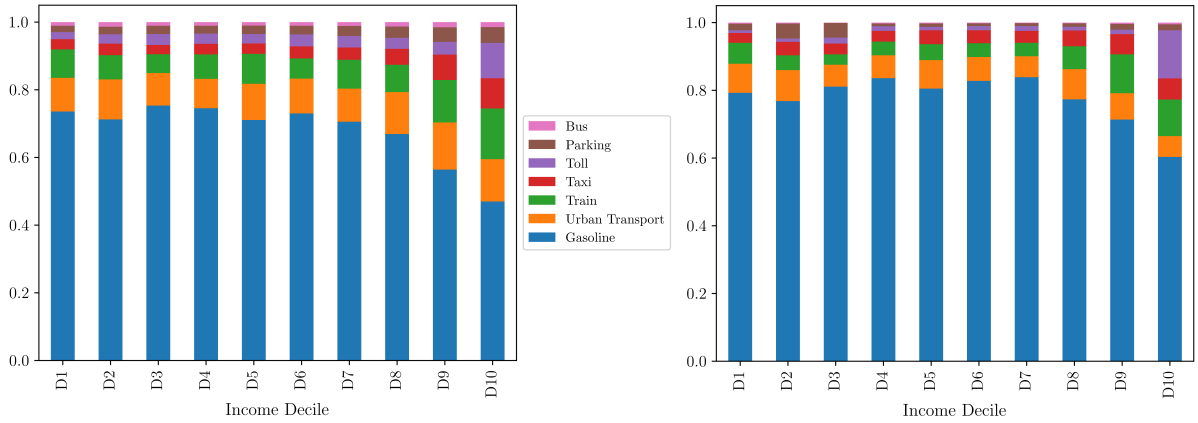
spending on the BBVA transportation categories in figure 11a. One can observe a very tight relationship between the two series during lockdown, which will be our main period of interest, although the two series diverge somewhat during lockdown easing. In the overall sample of days reported in figure 11a the correlation is 0.94.

Since card spending on transportation tracks mobility during lockdown very closely in the aggregate, we use it to explore heterogeneity across Madrid postal codes in mobility during lockdown. In principle the only people traveling outside the home during lockdown were key workers unable to telecommute. To the extent these occupations are skewed towards lower-income workers, one might expect to see differential mobility patterns across poorer and richer neighborhoods.[30] Figure 11b shows exactly this. Total transportation spending reductions during lockdown are significantly lower for customers living in the lowest decile postal codes by income per capita compared to those living in the highest decile. The average spending reduction for the former is 66% while for the latter the reduction is 85%, which is the maximum average reduction for any postal code decile.

Strikingly, these differences emerge primarily during the workweek: transport spending falls across postal codes appear much more similar during weekends that during working days. This strongly suggests that mobility differences across income groups arise because of different work patterns, not because of an innate preference for travel by lower-income households.

At the same time as mobility patterns shift in lockdown, we can also use spending on

[30]See Coven & Gupta, 2020 for evidence on mobility by postal code in New York City that comes from mobile phones.



(a) Transport Spend Shares: April 2019

(b) Transport Spend Shares: April 2020

Fig. 12: Spending Shares on Transportation Categories by Income per Capita Deciles of Madrid Postal Codes

Income Decile	D1	D2	D3	D4	D5	D6	D7	D8	D9	D10
2019	0.10	0.12	0.10	0.09	0.11	0.10	0.10	0.12	0.14	0.12
Lockdown	0.09	0.09	0.06	0.07	0.08	0.07	0.06	0.09	0.08	0.06

Table 8: Share of Urban Transport in Total Transport Spending by Income

the individual transportation categories to see how the *composition* of spending changes during lockdown. Figure 12a plots the share in overall transportation spending of different components for postal codes according to income deciles during April 2019 and during April 2020. One can observe that in normal times higher-income households devote more spending to categories that are more time efficient (taxis, toll payments). During lockdown, all income groups shift more spending towards gasoline which indicates greater reliance on car transportation conditional on traveling at all. As table 8 shows, though, this adjustment is not uniform across income groups. While higher-income households adjust substantially downwards the share of total transport spending on Madrid's public transportation system (urban transport for short in the plots), the lowest-income decile shifts by very little their public transport spending. Thus not only are lower-income households traveling more during lockdown during the workweek, they do so on more epidemiologically risky modes of transport.

5.1 The infection cost of mobility

The fact that households in lower-income postal codes travel relatively more than those in higher-income postal codes during lockdown, and do so at least in part on Madrid's public transportation system, suggests they may be relatively more exposed to infection.

Daily COVID-19 incidence within Postal Code	
	(1)
Lagged spending on urban transport	0.5729***
	(0.008120)
lockdown	1.590***
	(0.01792)
Lagged daily incidence	0.02644***
	(0.0001981)
Postal Code F.E.	Y
N	26784

Table 9: Estimated coefficients of disease model. Standard errors in parentheses.

We now seek to quantify this cost, which can be interpreted as the extra disease burden faced due to working in sectors that demand commuting during the lockdown.

We first introduce some notation to describe a simple regression model for how mobility patterns impact disease. We consider a panel data environment where i denotes a postal code and t denotes a day. We let $y_{i,t}$ denote the number of new COVID cases confirmed by official health statistics (data described above). To model mobility patterns, first let $x_{i,t}$ denote total spending on urban transportation of BBVA clients who live in postal code i on day t. We then compute the growth rate $g_{i,t} = \frac{x_{i,t}}{\overline{x}_i} - 1$ where $\overline{x}_i$ is average urban transport spending from 1 January 2020 through 14 February 2020. We express spending using this convention to keep in line with the mobility index as above. Finally, we obtain a lagged measure of mobility as $lg_{i,t} = \frac{1}{14}\sum_{\tau=t-28}^{t-14} g_{i,\tau}$. We build this lag into the model to account for the incubation time of coronavirus, delays in testing, and delays in confirmed cases appearing in official statistics.

We model $y_{i,t}$ using a Poisson regression model[31] with mean

$$\mu_{i,t} = \beta_1 y_{i,t-1} + \beta_2 lg_{i,t} + \beta_3 \text{Lockdown}_t + \gamma_i.$$

Lockdown_t is an indicator variable for whether a day falls in the post-lockdown period (recall that our case data begins in late February prior to lockdown) and γ_i is a postal code fixed effect which controls for income, age structure, distance from center of Madrid, etc.

Table 9 reports the estimated coefficients. As expected, we find significant and positive effects of the lockdown on case growth (since COVID-19 cases peaked during this time) as well as of lagged new cases (since infection dynamics are persistent). Less obvi-

[31]The Poisson model accounts for the discrete, non-negative count nature of the daily case data. If one instead uses an OLS model, the qualitative results we discuss below continue to hold.

ously, we also find a positive and statistically significant effect of lagged urban transport spending on current cases, which supports the idea that traveling in Madrid during the epidemic increases the probability of infection. The average marginal effect implied by the estimated coefficient is a quantitatively important 2.25 (obtained by multiplying the coefficient by 3.92, the sample average of new daily cases).

The estimated infection model allows us to conduct a counterfactual exercise to quantify by how much COVID-19 cases in lower-income neighborhoods would fall if households could reduce their urban transport spending to the level of the richest decile. The estimated Poisson model produces a fitted mean value

$$\hat{\mu}_{i,t} = \hat{\beta}_1 y_{i,t-1} + \hat{\beta}_2 lg_{i,t} + \hat{\beta}_3 \text{Lockdown}_t + \hat{\gamma}_i$$

To generate a counterfactual mean $\hat{\mu}_{i,t}^{\text{counter}}$, we replace $lg_{i,t}$—whenever t falls in the lockdown period—with a counterfactual mobility growth rate $lg_{i,t}^{\text{counter}}$ equal to the average value of $lg_{i,t}$ computed over all *top-decile* postal codes. From this we obtain

$$\hat{\mu}_{i,t}^{\text{counter}} = \hat{\beta}_1 y_{i,t-1} + \hat{\beta}_2 lg_{i,t}^{\text{counter}} + \hat{\beta}_3 \text{Lockdown}_t + \hat{\gamma}_i$$

This way of formulating the counterfactual ignores the fact that the entire disease path might potentially shift given alternative mobility patterns, but here we hold fixed all covariates except lagged transport spending. This likely underestimates the true impact of reducing mobility because reducing cases on a given day likely leads to case reductions in future days as well, due to lowering the stock of infected residents.

Income Decile	Total Cases	Total Reduction in Counterfactual	Percent Reduction in Counterfactual
1	6800	381	5.60
2	9285	443	4.77
3	9669	535	5.53
4	9418	366	3.88
5	11714	474	4.04
6	9522	344	3.62
7	8679	293	3.37
8	13705	450	3.28
9	15254	331	2.17

Table 10: This table reports the outcome of a counterfactual exercise in which we impose the mobility patterns of the highest-income decile of Madrid postal codes on all postal codes during lockdown. Here we report the outcome of the counterfactual by income decile. Total Cases comes from Madrid health data. Total reduction is the difference in predicted daily cases in the original regression model and with counterfactual mobility. The final column reports this number is a percentage of the total cases.

Our main counterfactual object of interest is the difference in cases under $\hat{\mu}_{i,t}$ and under $\hat{\mu}_{i,t}^{\text{counter}}$. We report these in table 10. In total we estimate the reduction in cases associated with the counterfactual as 3,613 which is 3.44% of the total cases in the data. These reductions are over-represented among lower-income postal codes especially in percentage terms. This is due to such areas engaging in the highest travel intensity during lockdown relative to their pre-lockdown level and doing so on urban transportation.

While more sophisticated modeling is clearly needed to jointly model disease progression and mobility proxies obtained from spending data, the basic application we present here demonstrates the novel insights one can obtain about inequalities in disease burden from transaction data.

6 Conclusion

The ability to track economic conditions at high frequency is important for making effective and timely policy choices. This is especially the case when conditions are changing rapidly and are subject to high levels of uncertainty, as is currently the case throughout the world due to the COVID-19 pandemic.

The current crisis comes at a time when the world is as rich in digital data as it has ever been, including detailed and granular information about transactions as stored by banks and payment systems. A pressing challenge is to use this data to provide signals to policymakers about the impact of COVID-19 and the policy interventions made to limit its spread.

More generally, the richness of naturally occurring transaction data, holds the promise of (at least) complementing existing (i) lower frequency national accounts and consumption series; (ii) cross-sectional consumption surveys and (iii) mobility data.

The results in this paper can be summarized along these dimensions. First, by validating our data, whenever possible, with respect to publicly available national accounts and surveys, we conclude that transaction data can be usefully deployed to track the economy and exploit time series, cross-sectional and geographical dimensions of the data. At the same time, we document that care should be taken due to the higher volatility and biases in coverage inherent to card data. This not withstanding, our results suggest that transaction data maybe useful to policy-makers and researchers alike, particularly in a middle-income and developing country context, where more standard high quality and high frequency indicators of consumption maybe too costly to produce.

Second, our proof-of-concept application - tracking the COVID-19 crisis - suggests a coherent, granular and high frequency narrative for expenditure adjustments during the pandemic. The decline in expenditure was abrupt and responded strongly to lockdown, social distancing policies and their progressive easing. Likewise, the recovery in expen-

diture seems fast, and more V than L-like. Underlying this decline in expenditure is a large reallocation across expenditure categories, away from social goods and luxuries. As a result, higher income groups - those who consume such goods relatively more in normal times - saw their incomes decline by more. Further, detailed transaction data on transportation and commuting expenditures reveals that the dynamics of transaction data can be used to explain disease incidence at the local level. Finally, the differential patterns of transportation expenditure - as a proxy for mobility during a pandemic - across rich and poor are associated to differential disease incidence across these groups.

Overall, our paper demonstrates how transaction data can be used to assess economic conditions. We show that such data is able to capture many relevant patterns in spending and that, importantly, it does so in near-real time. Moreover, its unprecedented granularity offers the possibility of using it as a high-resolution *"microscope"*; not only for deciding how best to weather future shocks - pandemic-related or otherwise - but also to provide the tools for an ever more granular and covariate-rich analysis of both economic events and economic models.

References

Alexander, D., & Karger, E. (2020). Do Stay-at-Home Orders Cause People to Stay at Home? Effects of Stay-at-Home Orders on Consumer Behavior. SSRN Scholarly Paper ID 3583625, Social Science Research Network, Rochester, NY.

Allcott, H., Boxell, L., Conway, J. C., Gentzkow, M., Thaler, M., & Yang, D. Y. (2020). Polarization and Public Health: Partisan Differences in Social Distancing during the Coronavirus Pandemic. Working Paper 26946, National Bureau of Economic Research.

Andersen, A., Hansen, E. T., Johannesen, N., & Sheridan, A. (2020a). Consumer Responses to the COVID-19 Crisis: Evidence from Bank Account Transaction Data. CEPR Discussion Paper 14809.

Andersen, A. L., Hansen, E. T., Johannesen, N., & Sheridan, A. (2020b). Pandemic, Shutdown and Consumer Spending: Lessons from Scandinavian Policy Responses to COVID-19. *arXiv:2005.04630 [econ, q-fin]*.

Baker, S. R. (2018). Debt and the Response to Household Income Shocks: Validation and Application of Linked Financial Account Data. *Journal of Political Economy*, *126*(4), 1504–1557.

Baker, S. R., Farrokhnia, R., Meyer, S., Pagel, M., & Yannelis, C. (2020a). How Does Household Spending Respond to an Epidemic? Consumption During the 2020 COVID-19 Pandemic. Working Paper 26949, National Bureau of Economic Research.

Baker, S. R., Farrokhnia, R. A., Meyer, S., Pagel, M., & Yannelis, C. (2020b). Income, Liquidity, and the Consumption Response to the 2020 Economic Stimulus Payments. Working Paper 27097, National Bureau of Economic Research.

Bodas, D., López, J. R. G., López, T. R., de Aguirre, P. R., Ulloa, C. A., Arias, J. M., de Dios Romero Palop, J., Lapaz, H. V., & Pacce, M. J. (2019). Measuring retail trade using card transactional data. Working Papers 1921, Banco de España;Working Papers Homepage.

Bounie, D., Camara, Y., & Galbraith, J. W. (2020). Consumers' Mobility, Expenditure and Online-Offline Substitution Response to COVID-19: Evidence from French Transaction Data. SSRN Scholarly Paper ID 3588373, Social Science Research Network, Rochester, NY.

Chetty, R., Friedman, J. N., Hendren, N., & Stepner, M. (2020). How Did COVID-19 and Stabilization Policies Affect Spending and Employment? A New Real-Time Economic Tracker Based on Private Sector Data. Working Paper.

Chronopoulos, D. K., Lukas, M., & Wilson, J. O. S. (2020). Consumer Spending Responses to the COVID-19 Pandemic: An Assessment of Great Britain. Working Paper 20-012.

Coven, J., & Gupta, A. (2020). Disparities in Mobility Responses to COVID-19. Working Paper.

Cox, N., Ganong, P., Noel, P., Vavra, J., Wong, A., Farrell, D., & Greig, F. (2020). Initial Impacts of the Pandemic on Consumer Behavior: Evidence from Linked Income, Spending, and Saving Data. Becker Friedman Institute Working Paper 2020-82.

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

Olafsson, A., & Pagel, M. (2018). The Liquid Hand-to-Mouth: Evidence from Personal Finance Management Software. *The Review of Financial Studies*, *31*(11), 4398–4446. URL https://doi.org/10.1093/rfs/hhy055

Simonov, A., Sacher, S. K., Dubé, J.-P. H., & Biswas, S. (2020). The Persuasive Effect of Fox News: Non-Compliance with Social Distancing During the Covid-19 Pandemic. Working Paper 27237, National Bureau of Economic Research.

Surico, P., Kanzig, D., & Hacioglu, S. (2020). Consumption in the Time of COVID-19: Evidence from UK Transaction Data. CEPR Discussion Paper 14733.

Watanabe, T., Omori, Y., et al. (2020). Online consumption during the covid-19 crisis: Evidence from japan. Tech. rep., University of Tokyo, Graduate School of Economics.

A Commercial Restrictions in the different Phases of the Easing Process

This is a summary of the restrictions imposed during the different stages of the easing process.

"Phase 0"

- From May 4th to May 17th small retail spaces can provide goods (or services, such as hairdressers) to individual customers (one by one) and only by previous appointment. These establishments need to separate seller from costumer by a screen.
  The facilities can not be within malls or within any bigger retail spaces.
  From May 18th, these shops can serve without the need of previous appointment (but still to only with one customer in the premises).
- Bars, restaurants and cafeterias can provide food and beverages "to go", but they can not be consumed in the premises.

"Phase 1"

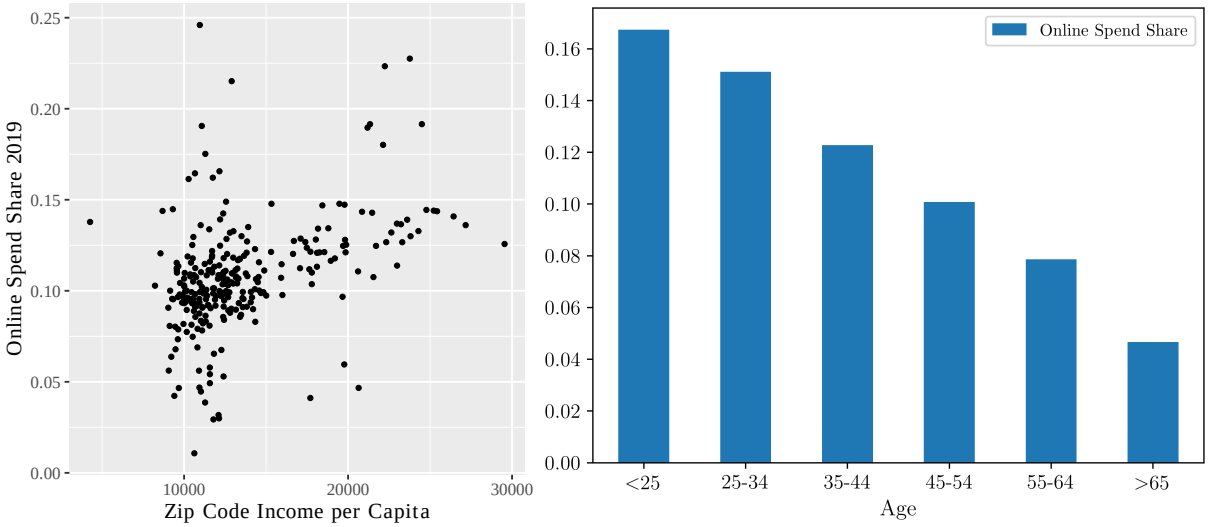
- Shops and retail spaces of less than 400 m^2 can open but restricting the number of customers to less than 30% of the capacity of the space.
- Car inspection facilities and garden centers can open, but restricted to individual customers and by appointment.
- Bars, restaurants and Cafeterias can open terraces (limited to 50% of capacity).
- Open Air markets can open with a 33% capacity and with a maximum of 1/4 of the shops opened at any given time.
- Hotels can be used, but not their common areas (cafeterias, restaurants, etc.)

"Phase 2"

- Shops can open independently of their size, but limiting customers to 40% of the capacity of the establishment.
- Malls can open at 30% of capacity.
- Bars, restaurants and Cafeterias can serve in the interior at a maximum of 40% capacity and with physical separation between costumers.
- Sporting events can take place, but with no public on the premises.
- Weddings with less than 50 persons inside premises or 100 if it is open air.
- Art Exhibitions and Cultural equipments may open at 30% capacity
- Sports equipments and swimming pools at 30% capacity
- Theaters, Cinemas and Life performances with audience limited to less than 50 in interiors and 400 in exteriors.
- Congresses and conferences with up to 50 attendees.
- Hotels can use common areas

"Phase 3"

- Shops, Retail Spaces, bars and restaurants with capacity restricted to 50%
- Malls can open with no capacity restriction
- Open air markets limited to 50% of stalls open at any time.
- Terraces of Bars and restaurants limited to 75% capacity.
- Weddings with limited of 75 persons (or 75% of capacity) inside premises or 150 open air.
- Casinos and Betting houses, limited to 50% capacity and a maximum of 50 customers.
- Summer Camps, limited to 1/3 capacity and with activities limited to 80 people inside and 200 outside.



(a) Income per Capita

(b) Age

Fig. 13: Online Spending Shares in 2019 and Household Characteristics

B Online vs Offline Spending

Card data provides insights into the distribution and extent of online shopping across households, an issue that has become prominent in discussions of consumption during lockdown. In this section we use BBVA card data to show (1) that online spending is correlated with both income and age, (2) that in Spain online spending share during the lockdown was negatively correlated with previous usage of online spending , and (3) that the lockdown implied a large substitution towards online spending but that this pattern corrected itself extremely fast once the easing policies took place.

(1) Figure 13 plots online spend shares in 2019 against per capita income at Madrid postal code level as well as online spend shares against age. Overall there is a strong correlation between higher income per capita and online spending (0.42, p-val 3.9e-13), although as the figure shows there is significant dispersion in online spending in the center of the income distribution. There is also a clear relationship between age and online spending, with a monotonic decrease in online share in age. These initial patterns suggest a high degree of heterogeneity in shopping behavior across households, an issue we revisit below in the analysis of lockdown consumption.

(2) In figure 14 we plot, at Madrid postal code level, the 2019 online spend share against the percentage change in online spending between 2019 and lockdown (we remove postal codes with very small 2019 online spend shares to avoid distortions to the percentage change metric). One observes a negative relationship between these series: postal codes with a lower online spend share increase proportionally more their online spend during lockdown. This is consistent with households that previously had not made a shift to online spending beginning to do so in order to maintain consumption level during lockdown. This pattern is different from the one that has been observe in Japan, where (Watanabe et al., 2020) show that the consumers that have increased online spending the most are those with previous online experience.

(3) Nevertheless this substitution does not seem to have generated a persistent consumption pattern. Once the restrictions of the lockdown have been relaxed, the share of online spending has returned fast to its previous levels. In figure 15a it is patent that during the lockdown the share of online spending grew very fast initially, not because online spending grew, but because

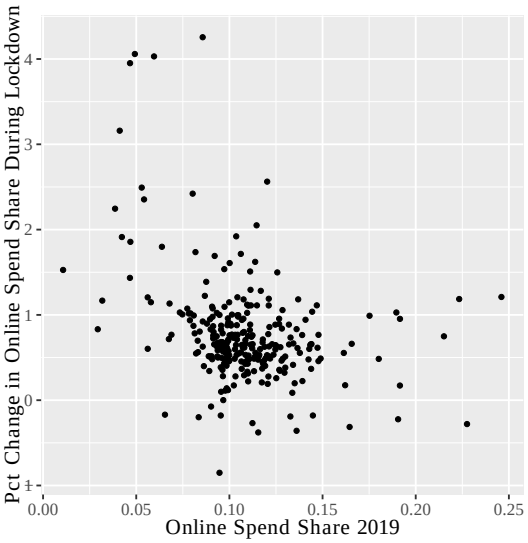
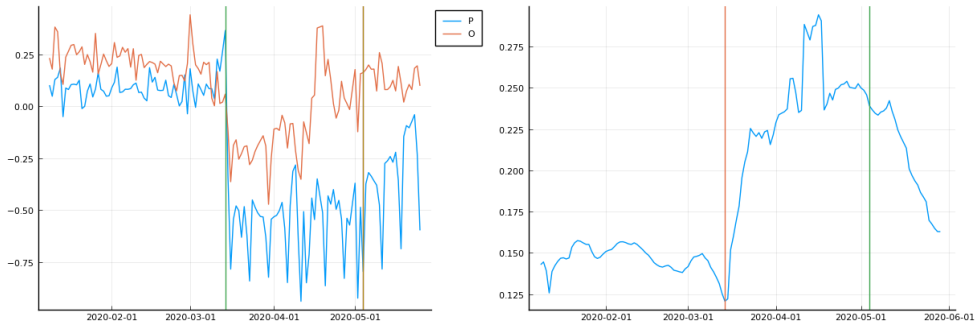


Fig. 14: income and online



(a) Y-o-Y expenditure growth online and offline.

(b) Share of online spending.

online spending decreased less than offline spending. The recovery of offline sales once the easing process is started implies that the share of online sales returns very fast to its previous values.

It has to be noted that on this issue Spain might be an exception, as in *all* other BBVA markets where we have data online spending increased during the pandemic. Still, our data in those markets is not detailed enough to make clear inference.

C Regression on Madrid Postal Codes

Table 11 is equivalent to table 7 in the text, but instead of using the new cases per day we regress into teh accumulated cases in the postalcode interacted with a lockdown dummy. As it is clear the results are qualitatively identical.

	Daily Y-o-Y Expenditure Growth at each Postal Code					
	(1)	(2)	(3)	(4)	(5)	(6)
Covid Incidence	-3.497***		-3.864**			
	(0.853)		(1.268)			
lockdown & Covid Incidence	-4.346***		-3.916***	-5.044***		-4.666***
	(0.766)		(0.852)	(0.824)		(0.934)
quintile 02		0.017	0.039			
		(0.058)	(0.056)			
quintile 03		0.007	0.046			
		(0.055)	(0.055)			
quintile 04		-0.132**	-0.087			
		(0.043)	(0.047)			
quintile 05		0.094	0.190			
		(0.184)	(0.209)			
lockdown & quintile 02		-0.139	-0.113		-0.109	-0.081
		(0.098)	(0.093)		(0.102)	(0.096)
lockdown & quintile 03		-0.234*	-0.185		-0.190	-0.138
		(0.100)	(0.095)		(0.105)	(0.099)
lockdown & quintile 04		-0.253**	-0.210**		-0.257**	-0.204**
		(0.080)	(0.074)		(0.084)	(0.078)
lockdown & quintile 05		-0.329**	-0.232*		-0.309*	-0.192
		(0.112)	(0.116)		(0.120)	(0.125)
day F.E.	Yes	Yes	Yes	Yes	Yes	Yes
Postal Code F.E.				Yes	Yes	Yes
N	37,213	37,213	37,213	37,212	37,212	37,212
R^2	0.027	0.019	0.029	0.126	0.125	0.126

Table 11: Regression of Madrid postal code daily Y-o-Y growth rates on lockdown dummy variable, daily COVID-19 incidence per capita in each postal code, Income quintile of postal code and interactions with lockdown period. Standard errors clustered at the Madrid postal code level.