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

Vasco M. Carvalho [*] Juan R. Garcia [†] Stephen Hansen [‡] Álvaro Ortiz [§]
Tomasa Rodrigo [¶] José V. Rodríguez Mora [‖] Pep Ruiz [**]

July 22, 2021

Abstract

Financial and payments systems throughout the world generate a vast amount of naturally occurring, and digitally recorded, transaction data. This paper considers billions of transactions from credit- and debit-card data from BBVA, one of the largest banks in the world, as an alternative source of information for measuring consumption, a key component of GDP. We show, through validation exercises against official consumption measures, that transaction data can usefully complement slow-moving national accounts and consumption surveys. We show that this holds (i) over time, as a high frequency consumption proxy both at national and subnational levels; (ii) over consumption categories, rendering it a naturally occurring consumption survey and (iii) over space, as a covariate-rich mobility dataset. We use these features to analyze the impact of the arrival of COVID-19 in Spain and the first national lockdown, and document three results of broad policy relevance for managing lockdowns: (1) strong consumption responses to shop closing and opening, but more muted effects for capacity restrictions; (2) a decline in spending for residents of high-income neighborhoods; (3) higher mobility during the workweek for residents of lower-income neighborhoods, which correlates with increased disease incidence.

Keywords: National Accounts | Transaction Data | Consumption | COVID-19

[*] University of Cambridge, Alan Turing Institute and CEPR. `vmcarvalho.web@gmail.com`
[†] BBVA Research. `juanramon.gl@bbva.com`
[‡] Corresponding author. Imperial College Business School and CEPR. `stephen.hansen@imperial.ac.uk`
[§] BBVA Research. `alvaro.ortiz@bbva.com`
[¶] BBVA Research. `tomasa.rodrigo@bbva.com`
[‖] University of Edinburgh, Alan Turing Institute and CEPR. `sevimora@gmail.com`
[**] BBVA Research. `ruiz.aguirre@bbva.com`

1 Introduction

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

While there is broad agreement among national statistical agencies that unstructured transaction data will play an increasingly prominent role in 21st century national accounting (see Bean (2016), Abraham, Citro, White, & Kirkendall (2018) and Jarmin (2019)), national statistical agencies, academics and policy-makers still largely rely on more traditional structured survey data and slow-moving national accounts updates.[1] Partly, this reluctance reflects concerns regarding the accuracy and representativeness of transaction data. Indeed, 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? While there is a reasonable expectation that economists and government agencies will have increased access to large-scale transaction datasets in the near future, extensive validation against available official statistics is needed in order for transaction data to fulfill its promising role in national accounting.

The first contribution of this paper is to analyse these issues in the context of the universe of credit and debit card transactions mediated by a large global bank, Banco Bilbao Vizcaya Argentaria, S.A (BBVA). Our data consists of the universe of transactions collected from BBVA cardholders and BBVA-operated point-of-sale in Spain, accounting for 2.1 billion transactions.[2] We explore the properties of the data along three different dimensions: as a high-frequency coincident indicator for aggregate and subnational consumption; as a detailed household consumption survey; and as a mobility index.

In each case, we show that 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 official data. This then allows one to make further cuts into the spending data to obtain insight unavailable using external series alone.

Our second contribution is to show how this transaction data, once validated, offers several policy-relevant lessons from the first Spanish lockdown—one of the world's harshest—that are relevant for the numerous countries currently re-entering lockdown. We use the data along each of the three dimensions above to obtain valuable, but otherwise largely hidden, lessons related to the effects of the pandemic and lockdown polices.

First, we exploit subnational high-frequency expenditure data in tandem with systematic changes in lockdown policies across spatial units to evaluate the differential effects of those policies. We show that restrictions of activity that work through limiting capacity and customer density have only a mild effect on expenditure, in particular, when compared with the effect of forcing the closure of large retail establishments.

Second, we exploit the transaction data as a detailed consumption survey, which allows us to track changes in the composition of consumption and the structure of consumption across income classes. We document that residents of the richest zip codes had the largest declines in expenditure during lockdown.

[1] Important exceptions include Gelman, Kariv, Shapiro, Silverman, & Tadelis (2014), Baker (2018) and Olafsson & Pagel (2018), which use data from financial apps to test consumption smoothing theories.

[2] Since BBVA is a global bank, it generates several billion more transactions across other countries in which it has a large market share, for example Turkey, Mexico, and the Southern US. An earlier version of this manuscript included discussion of the global time series, which can be downloaded here https://www.bbvaresearch.com/en/special-section/charts/ which we omit for space constraints.

Further, we show that this is explained because lockdown restrictions, by their very nature, affect more predominantly the pattern of conspicuous consumption prevalent in wealthier individuals.

Third, we show that expenditure in transportation correlates exceedingly well with external mobility measures, and that during the lockdown the mobility of the rich was substantially smaller than that of the poor. Moreover, we also show that differential mobility patterns predict heterogeneity in the incidence of the pandemic across income groups.

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 arising from card spending and point-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, Ganong, Noel, Vavra, Wong, Farrell, & Greig (2020) and Chetty, Friedman, Hendren, & Stepner (2020) in US studies, and Surico, Kanzig, & Hacioglu (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, Camara, & Galbraith (2020) for France; Carvalho, Peralta, & dos Santos (2020) for Portugal; Chronopoulos, Lukas, & Wilson (2020) for the UK; and Andersen, Hansen, Johannesen, & Sheridan (2020) and Alexander & Karger (2020) for the US. Further, our analysis of the effects of lockdown and its easing complements that in Asger Lau Andersen & Sheridan (2020). The latter argue for the importance of behavioral adjustments in expenditure patterns, responding to local disease dynamics even in the absence of lockdown policies. Consistent with this, we find local disease incidence to be a driver of expenditure changes, even when controlling for different levels of lockdown restrictions across space. Unlike Asger Lau Andersen & Sheridan (2020), 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) and Glaeser, Gorback, & Redding (2020), we explore the relation between mobility and disease incidence. Relative to that contribution, we show that in the absence 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.

2 Results

We organize the results by first validating proxy measures derived from Spanish card data against external data in Spain, then applying the proxy to understand an important aspect of the COVID-19 crisis.

2.1 Transaction Data as a High-Frequency Consumption Proxy

Validation

We compare total spending via BBVA cards and PoS terminals with the national account household consumption series ("Non-Durable Household Domestic Final Consumption") for every quarter since 2016. We also compare time series of spending at monthly frequency, and on specific components present in BBVA and national accounts. As detailed in the Supplementary Information, we find that: (i) BBVA card expenditure series correlate highly both with aggregate national-accounts consumption (correlation of 0.874), and within narrowly defined consumption categories series where official data is available, in particular expenditures at gas stations (correlation of 0.784); (ii) that nevertheless, both at the aggregate and sector level the BBVA series is more volatile than the official series. The likely cause of the latter is that card data does not cover stable household expenses—such as rents, school fees, some utilities and subscription services—and that over long spans of time there are likely extensive margin movements, reflecting entry and exit of clients, cards, and PoS in the BBVA sample.

We next validate spending data in the cross section of geographic units. There are no official subnational consumption measures in Spain, so instead we compare BBVA spending to official data on income in Spanish provinces (52 in total) and Madrid postal codes.The correlations are extremely high: 0.975 across provinces, and 0.923 across postal codes. (Further details in Supplementary Information).

The conclusion is that card spending captures important patterns across space and time in national accounts data, albeit with more noise in the latter than the former.

Effects of Lockdown and its Easing

With many major European economies again facing extended lockdowns due to a resurgence in COVID-19 cases, the optimal balance between economic activity and public health is again of paramount importance. We next use the imposition of the first lockdown in Spain in March 2020 and its subsequent, progressive easing to draw lessons for managing restrictions going forward.

The Supplemental Information contains background information on the development of COVID-19 and Spanish government policy responses during March-June 2020, and here we provide a brief summary. A national lockdown was first imposed beginning on March 15th in response to rapidly growing infections. The measures were among the harshest in the world, and resulted in the suspension of all but essential economic activity. After a subsequent fall in cases, the government began *Phase 0* easing on May 4 which permitted small retail stores to operate under strict social distancing guidelines. This first easing stage applied uniformly to all regions in Spain, but further easing was staggered across provinces.

On May 11 some provinces entered *Phase 1* which allowed for larger retail spaces (but not superstores and malls) to reopen at restricted capacity and for outdoor commercial activity (including restaurants) to resume. *Phase 2* then began on May 25 for some provinces and lifted all size restrictions on commercial activity (including malls) and some indoor commercial activities, while still keeping capacity caps in place. *Phase 3* began on June 8 and relaxed further these capacity limits.[3]

Figure 1 plots aggregate expenditure growth in Spain over this period, normalized by average year-on-year (Y-o-Y) growth prior to March 8th. Expenditure growth fell abruptly on the day of lockdown, by about 60 p.p. and remained depressed at that level until early May, when easing of lockdown ensued. The aggregate data is also suggestive of a recovery starting with the nationwide enactment of *Phase 0*. By the 21st of June, when our data end, expenditure growth in Spain is only a few percentage points off its pre-COVID-19 average, denoting a near complete recovery in expenditure.

The staggered adoption of easing phases across provinces, combined with spending data at the day and province level, provides a unique opportunity to study consumption reactions to different kinds of economic restrictions. Figure 2a 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 remained in the more restrictive Phase 0 (in blue). Figures 2b and 2c 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. The easings into Phases 1 and 2 appear to be on average associated with higher spending for switchers vs. stayers. On the other hand, the easing into Phase 3 has a much less marked impact. This provides evidence that shop openings generate more economic impact than the lowering of capacity restrictions. To the extent that capacity restrictions provide public health benefits, this provides a strong

[3] A further *Phase 4* began on June 21 and represented a return to essentially normal economic activity, but we exclude this from our sample below due to too few days entering this period.

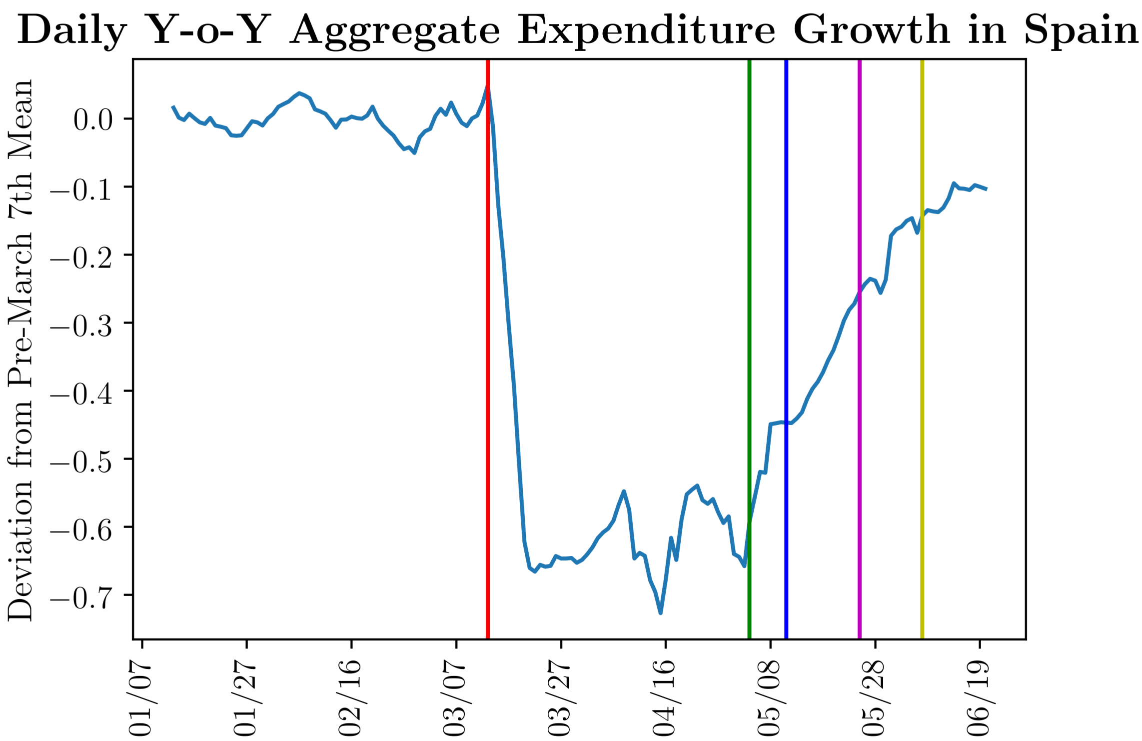

Fig. 1: Moving average (7 day, uncentered) of Y-o-Y growth of expenditure from BBVA series for Spain (aggregate). The vertical lines indicate the timing of events. The first (red) vertical line is drawn on March 13th, the day prior to the announcement of lockdown. The second one is May 4th (start of Phase 0), when easing started nationwide. The third 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 7th.

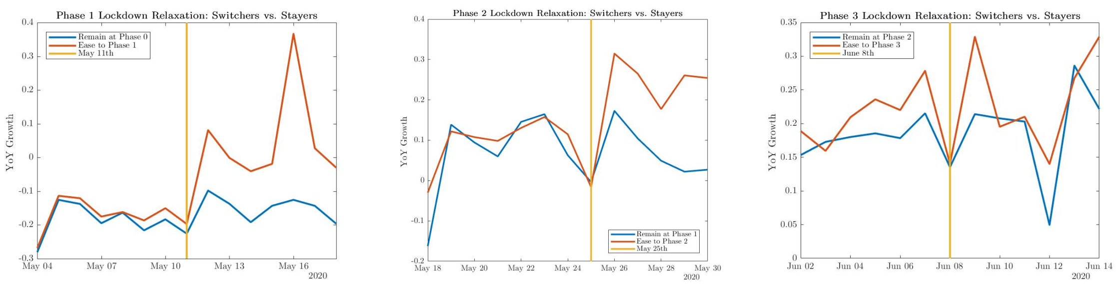

(a) Phase 1 Event-study (b) Phase 2 Event-study (c) Phase 3 Event-study

Fig. 2: Event Study Graphs. 2a: Average Y-o-Y expenditure growth for the provinces which eased into Phase 1 on May 11th (in orange) and average growth for provinces that stayed in the more restrictive Phase 0 (in blue). 2b: Id. but centered around May 25th when some provinces eased into Phase 2 while others remained in Phases 0 and 1. 2c: Id. but centered around June 8th, when some provinces eased into Phase 3 while others remained in previous Phases. All figures use 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.

Table 1: 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 province 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.

	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

rationale for maintaining them in place whenever even moderate infection risk is present.

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 at least partly address this issue we now turn to regression analysis. In Table 1 we present panel regressions of the daily provincial Y-o-Y growth of expenditure on lockdown-phase and easing dummies i.e. binary variables for each province and period, which take a value of one if that particular province is classified in a particular phase of the lockdown - or lockdown easing - a given calendar day and zero otherwise. Note further that, as discussed above, for the week immediately preceding the lockdown, the lockdown itself and Phase 0 of lockdown easing, all provinces move in lockstep, so these categorical variables display the same time pattern for all provinces. Instead, for Phases 1,2 and 3, the time pattern is province-specific, depending on when a particular province advanced to the later lockdown easing phases. 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 Y-o-Y expenditure growth on a series of time dummy variables, where the omitted category is the period before March 8th, one week before any official discussion of 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 of it. The period of strict lockdown, with its associated restrictions on commercial activity, led to a large fall of about 60 p.p. in Y-o-Y growth of expenditure. These patterns are consistent with Figure 1 where one observes that this expenditure contraction coincides with the beginning of lockdown and lasts as long as restrictions remain at their strictest level, up until May 4.

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 pre-lockdown value by the time a provinces reaches Phase 3. Overall, based on these simple means, the Phase 1 and Phase 2 easings which opened 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 1 we additionally control, respectively, for differential disease dynamics across provinces, province fixed effects, and both together. 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 endogeneity 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 1 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 yielding a standard difference-in-differences setup with (*i*) variation in treatment timing across units and (*ii*) multiple treatments. Note also that, relative to the previous specifications, the omitted category is now "Phase 0", the last common policy baseline across all provinces and therefore the interpretation of the coefficients changes. For example, estimates pertaining to Phase 1 now give the percentage *growth* in expenditures for provinces that proceeded to this lockdown easing stage – at whatever calendar date they may have done so – relative remaining at Phase 0 for longer (for further discussion on interpretation and references on this estimator, see our Methods Section 4).

The estimates we obtain are nevertheless similar 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 establishments 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 and, as we show in the Methods section below, are also robust to a further checks related to the possible endogeneity of the timing of lockdown easing.

Finally, these robust correlations notwithstanding, we end this section with a word of caution when interpreting these estimates as the true "causal effect" of lockdown policies. This is because, as is well known, identification of causal effects in our context would require province-level lockdown policies and

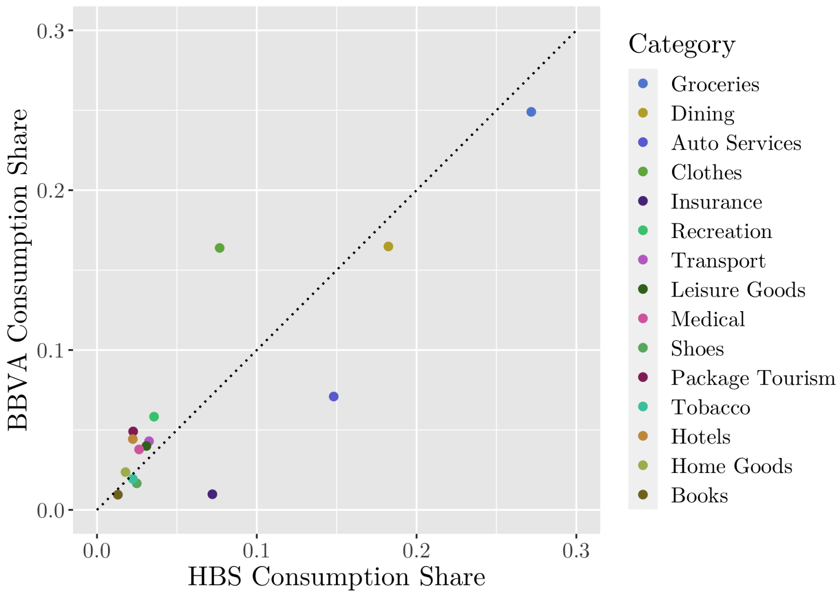

Fig. 3: Consumption Share Comparison in Matched ECOICOP Product Categories

their timings to be "as good as random", at least conditional on time and province fixed effects and, possibly, other relevant time-varying province-level covariates. In particular, while our most demanding specification above attempts to account for all of these, we cannot rule out the presence of other unobserved, time-varying, province-level conditions which (*i*) may have influenced selection into treatment – beyond the province specific prevalence of COVID; for example evolving socio-economic considerations by the Spanish government – and/or (*ii*) have had a bearing on expenditure decisions of household such as the evolution of province-specific expectations of disease prevalence which, in turn, may lead to behavioral expenditure responses that go beyond the particular *de jure* lockdown regime and may not be accounted by province-level disesase prevalence. Thus, to the extent that these time-varying province-level unobservables were operational during lockdown easing in Spain, our estimates may be biased – in either direction – relative to the true causal effect.

2.2 Transaction Data as a Granular Consumption Survey

Validation

National statistics organizations traditionally measure household consumption baskets with representative spending surveys. On the other hand, transaction data derived from card transactions typically contains associated metadata which allows a breakdown of expenditure across goods and services categories. Can these two sources of data be bridged? Can metadata on card transactions stand in for nationally representative consumption surveys?

In the Supplementary Information, we compare in detail household spending across categories as measured by the official Spanish Household Budget Survey (HBS) and by the BBVA dataset, which breaks purchases into one of 77 distinct categories. The two data sources have distinct categorizations, which requires a manual match; in total we find matching categories for 65% of BBVA spending. Figure 3 plots consumption shares in the matched categories in both datasets, which have a correlation coefficient of 0.865.

In a second validation exercise, we consider the subsample of BBVA transactions that involve a BBVA debit or credit card, in which case we have information on the consumer's demographic characteristics.

As we detail in the Supplementary Information, the share of consumption per age and education groups aligns remarkably well between BBVA data and HBS.

Finally, we tabulate total BBVA debit and credit card spending by Madrid postal code, and use postal code income as a proxy for household income. The allocation of consumption across categories according to income derived from BBVA data also aligns exceedingly well with the one observed in HBS.

These three validation exercise demonstrate that information derived from BBVA purchase categories aligns relatively well with information from the HBS along comparable cuts of data, a fact we can use to document the allocation of spending in real time during the onset of the COVID-19 crisis.

Composition of Consumption in the Lockdown

Our first application of using card spending as a consumption survey is to study the spending reallocation induced by the Spanish lockdown (March 15 through May 4). The Supplementary Information lists the 77 BBVA spending categories, and identifies the categories that were directly subject to lockdown restrictions which include a broad set of non-essential shopping categories as defined by the Spanish government.

Table 2: Best and Worst performing categories of expenditure by market share post-lockdown growth. In red, categories restricted during the lockdown.

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 2 lists the top ten and bottom ten spending categories according to the evolution in market share before and after lockdown; categories directly affected by lockdown measures are in red, which (perhaps unsurprisingly) constitute all of the bottom categories and none of the top categories. More notable are the enormous shifts in spending in this period, which some categories collapsing nearly entirely while others increase by 100% or more their market share. The goods and services with market share growth in lockdown relate to basic necessities (such as food), or have very low demand elasticity (such as tobacco). All of them were deemed critical sectors, and remained open for business during the lockdown, albeit with restrictions on capacity and customer density at any given point in time.

Figure 4 provides visual evidence of these spending shifts by plotting the market share across 18 broad spending categories that combine the 77 disaggregated categories. These shares are quite stable up until the week preceding the national lockdown, when a clear reallocation pattern emerges: spending on food and in "hypermarket" (i.e. large superstores) grows considerably, and these two categories 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 quickly once the easing of restrictions began, the composition of consumption returns steadily to pre-lockdown allocations following the entry in the "phase 0" of the easing period, on May 4. We provide further time-series figures in the Supplementary Information to study the evolution of the disaggregated categories.

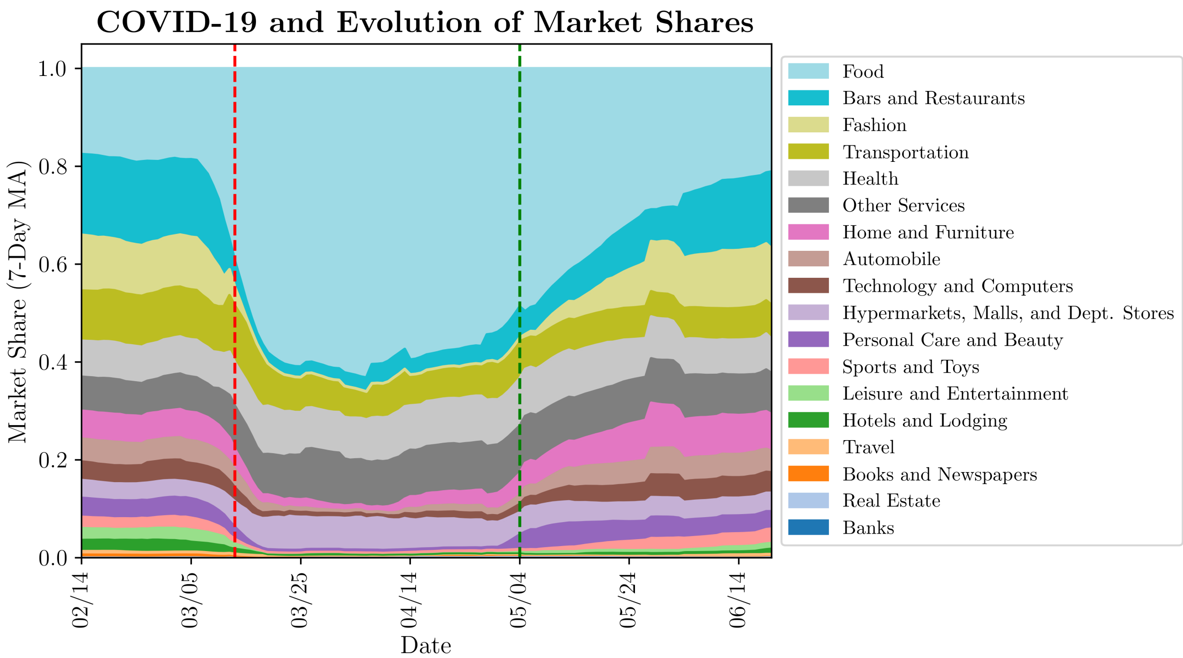

Fig. 4: Evolution of the share of offline spending across categories. Red dashed line=national lockdown begins; green dashed line=national lockdown begins to ease.

Dynamics of Aggregate Consumption across Income Groups during the Lockdown

The shift in consumption during lockdown masks important underlying underlying heterogeneity with respect to income, which our card data allows us to explore in detail using expenditure patterns by different Madrid postal codes. Here we measure spending of BBVA cardholders who have a registered address within a given postal code, and exclude PoS spending since we do not observe the home address of non-BBVA cardholders.

In Table 3 we present the categories that during 2019 were most positively and negatively correlated with postal-code income per capita. One observes a pattern whereby higher-income groups consume goods associated with leisure and market production, while lower-income groups purchase more necessities and engage in home production. Marked in red are those categories whose consumption was restricted during lockdown. Goods associated with the higher-income groups are relatively more affected by lockdown restrictions, which suggests that the consumption basket of higher-income groups became more like that of the poor during lockdown.

The implications of the alternative consumption baskets consumed by different income groups can be seen in figure 5, which plots a moving average of expenditure growth for Madrid postal codes binned by quintile according to income per capita. The sharpest declines in spending during lockdown concentrate in the richest postal codes, which is consistent with the rich being unable to consume their normal goods basket due to restrictions.

In the Supplemental Information we perform more formal statistical analysis in order to quantify these effects more rigorously, and control for disease dynamics that might also drive neighborhood-level spending. These regressions not only confirm the that wealthier neighborhoods were the ones experiencing the largest fall in expenditure. They additionally suggest that areas more affected by the pandemic had larger declines in expenditure.

Table 3: Categories more positively and negatively correlated with average income across Madrid postal codes. In red, categories restricted during the lockdown.

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

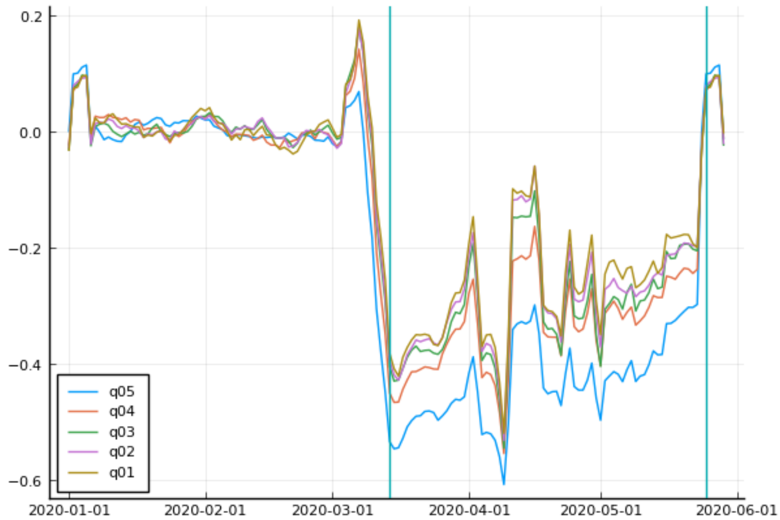

Fig. 5: Y-o-Y growth rate of expenditure in Madrid's postal codes during 2020 by postal code average income (in quintiles). Normalized by the average Y-o-Y growth before 08/03/2020. The two vertical lines indicate (i) the lockdown day (March 15th) and (ii) the day the whole of Madrid went into phase 1 of the easing process (May 25th).

2.3 Transaction Data as a Real-Time Mobility Proxy

Validation

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 disaggregated 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.

Besides shopping for essential goods, the main source of mobility during Spain's lockdown was commuting for work. We use card data to measure this by considering BBVA spending categories that relate directly to transportation: 'bus trips'; 'gas stations'; 'parking'; 'tolls'; 'taxi'; 'urban transport'; 'trains'. To validate this as a travel-to-work measure, we compare transportation spending growth against the 'work places' and 'transit' stations categories from Google's Mobility Report for Spain, which expresses

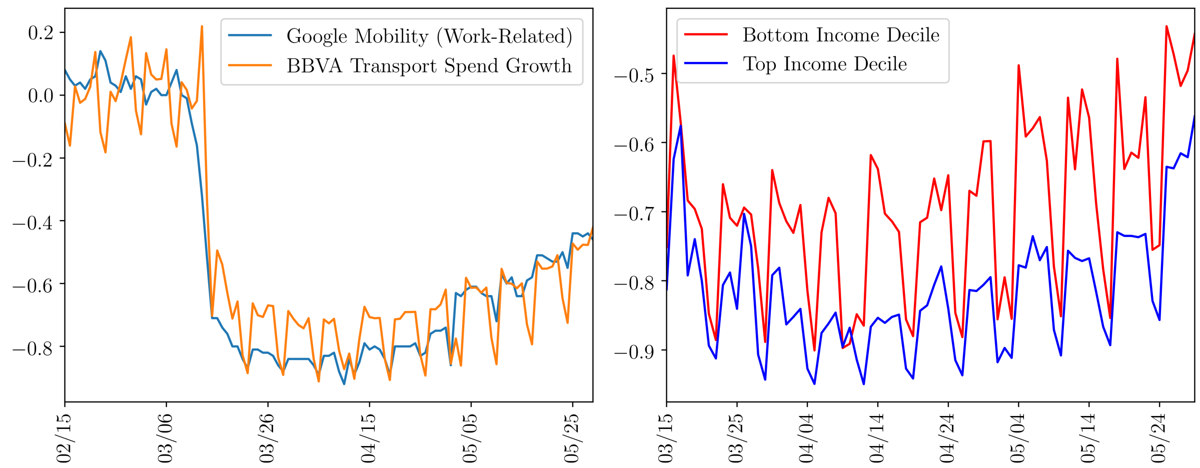

(a) Transport Spending vs Mobility (b) Transport Spending and Income

Fig. 6: Left panel: comparison of Google Mobility Report for Spain for work-related categories 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 among top-income and bottom-income Madrid postal codes during lockdown period.

time spent in these locations in percentage change terms using mobile phone location data. Figure 6a plots the two series, which track each other closely throughout the sample, albeit with more weekly seasonality in the card spending data. In the overall sample of days reported in Figure 6a the correlation is 0.94.

Income and mobility

Previous literature has highlighted that lower-income workers are more likely to have jobs for which working from home is not possible (Dingel & Neiman, 2020), but whether such workers continue to work, or suspend their labor market activity and remain at home, is not clear. Figure 6b plots the change in transportation spending during lockdown among cardholders residing in the lowest- and highest-decile Madrid postal codes (by income per capita). The average spending reduction relative to pre-COVID baseline for the former is 66% and for the latter is 85%, which is the maximum average reduction for any postal code decile (see Coven & Gupta, 2020 for evidence on mobility by postal code in New York City that comes from mobile phones). Strikingly, these differences emerge primarily during the workweek: transport spending falls across postal codes appear much more similar during weekends than 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. It also suggests that a substantial number of workers unable to work from home continue to work in lockdown, even if in theory only essential workers were supposed to leave home.

To further explore the relationship between income and mobility during lockdown, we look at spending patterns by Madrid postal code during the peak of the lockdown in April 2020. Figure 7a plots the share of online spending in total spending during this period against postal code income per capita.[4] The raw correlation between the variables is 0.43 (p-value $<$ 1e-13), although the plot makes clear there is substantial variation in online shopping behavior across all income groups. This nevertheless provides evidence that residents of higher-income postal codes are more able to shop online and avoid leaving

[4] In the Supplementary Information we show that richer neighborhoods also had a substantial *increase* in online spending of food (a necessity good) during the pandemic.

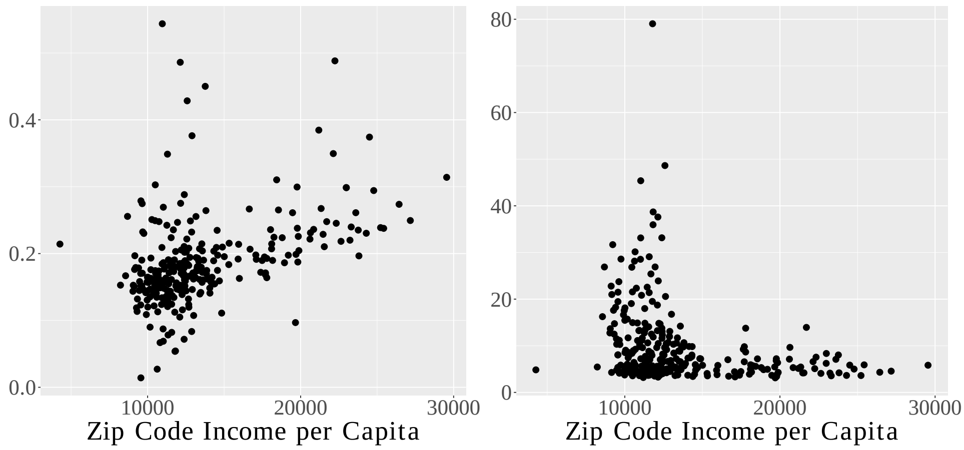

(a) Online Spend Share (b) Distance Travelled for Offline Spending

Fig. 7: This figure compares shopping behavior related to mobility across Madrid postal codes during April 2020. Left panel: Share of total spending in April 2020 purchased online. Right panel: for each postal code, we estimate the distance travelled in kilometers for making offline purchases in April 2020.

their homes during lockdown.

We next examine the distance traveled across postal codes conditional on making offline purchases, which by definition requires leaving one's home. To do this, we first create a dyadic dataset in which we tabulate the offline purchases made by residents of each postal code in all other postal codes in Madrid (including one's own postal code). share_{ij} is the share of offline spending[5] of postal code i purchased in postal code j. We then compute d_{ij}, the distance in kilometers between geographic centroids of postal codes i and j. Finally, our estimate of the distance traveled for offline shopping of postcode i residents is $\sum_{j \neq i} d_{ij}\text{share}_{ij}$. That is, we weight the physical distance between postcodes by spending shares, and impute a zero distance to purchases made in own postal code. Figure 7b plots this estimate against postal code income per capita.[6] Among postal codes with income per capita above 15,000, the average distance traveled for offline purchases is 5.7km and the interquartile range is (4.2km, 6.6km). Among postal codes with income per capita below 15,000, the corresponding statistics are 11.2km and (5.2km, 13.3km). The implication is that not only are residents of poorer postal codes less likely to make purchases online, but also more likely to travel greater distances when they leave home to make offline purchases. Both facts combine to provide further support to the idea that substantial mobility inequality existed across income groups during Spain's lockdown. Moreover this demonstrates that card purchase data can be informative about physical movements across narrow geographic units.

The infection cost of mobility

A natural next question is whether mobility has health consequences. To the extent that travel outside the home makes it more likely to interact with others, it may increase the risk of contracting coronavirus. Our results above motivate us to use card spending on transportation as an input into a disease model to explore this connection. Furthermore, we explore how different modes of transportation affect disease incidence. From 1 February 2020 through 30 April 2020, two modes of transport make up 75% of total

[5] In the Supplementary Information we further extend the analysis counting the number of transactions (instead of the share of purchases). We show that during the lockdown the number of offline transactions performed outside their zipcode by residents of richer neighborhoods felt much more than in poorer ones, while there are no substantial changes within the neirborhood of residence.

[6] There exist a literature aiming to understand distance to the consumption point. See for instance Miyauchi, Nakajima, & Redding (2021)

spending in Madrid postal codes on transportation: gasoline (63% of total spending) and urban transport (12%). We take the former as a proxy for car transportation, while the latter represents spending on Madrid's public transportation system. A reasonable expectation is that public transportation brings travelers into closer contact with others, so might represent a particularly high-risk form of mobility during the pandemic. This represents another application of card data as a consumption survey, as the detail provided by the spending categories allows us to dig into impacts of different types of travel.

Cumulative COVID-19 Incidence within Postal Code				
	(1)	(2)	(3)	(4)
Total Transport Spending	0.472**			
	(0.027)			
Car Transport Spending		0.055		
		(0.039)		
Urban Transport Spending			0.760***	
			(0.201)	
Urban + Car Spending				0.070*
				(0.036)
Income per Capita	-0.008	0.027	-0.039	0.018
	(0.047)	(0.043)	(0.046)	(0.044)
Senior Share	30.872***	32.372***	23.588***	31.729***
	(3.576)	(3.591)	(4.124)	(3.564)
R^2	0.272	0.261	0.297	0.267
N	248	248	248	248

Table 4: Estimated coefficients of ordinary least squares model for total cases per 1,000 residents at Madrid post code level. Standard errors in parentheses. *** p<0.01; **p<0.05; *p<0.1.

To begin the analysis, we regress total COVID-19 cases per 1,000 residents in each postal code on income per capita (measured in units of EUR 1,000), the share of residents above 65, and total spending per capita on transportation of different forms during February, March, and April 2020. The estimated coefficients are in table 4. As expected, the share of older residents is a strong predictor of total cases but we find no effect of income per capita. More pertinent for our purposes, we find a moderately strong impact of total transportation spending on cases. The estimated coefficient implies that one standard deviation change in total transport spending generates a 0.143 standard deviation change in COVID-19 cases. This is consistent with transport spending correlating with social contact and disease exposure, which thereby increase disease incidence.

We also find strong heterogeneity in the association between types of transport spending and disease. Spending on car transportation has no significant effects on COVID-19 incidence, but spending on urban transport has very strong effects. The estimated coefficient implies that a one standard deviation change in urban transport spending generates a 0.267 standard deviation change in COVID-19 cases, nearly twice the effect of generic spending. The final column pools urban and car spending. As expected, given car spending makes up most of this combined category, the effects are quite similar to car spending alone.

This highlights that the mode of transportation may be as big a component of health risk as mobility *per se*. Prior to lockdown, higher-income neighborhoods have a slightly higher share of urban transport spending in total transport spending than lower-income ones. During lockdown, urban transport shares are uncorrelated with income at the postal code level.

There are many factors that the cross-sectional regressions do not control for. Distance from the center of Madrid, occupational structure, quality of housing stock, and population density are all factors that might potentially drive the relationship between disease and mobility. To address these sources of confounding, we next adopt a panel regression framework that allows us to study the impact of transportation spending at daily frequency on disease outcomes *within* post codes while controlling for postal code fixed effects. The methods section below formally describes the Poisson regression model we adopt.

Daily COVID-19 Incidence within Postal Code				
	(1)	(2)	(3)	(4)
Lagged Total Transport Spending	0.036***			
	(0.0004)			
Lagged Total Car Spending		0.070***		
		(0.0007)		
Lagged Urban Transport Spending			0.125***	
			(0.0021)	
Lagged Urban Transport + Car Spending				0.051***
				(0.0006)
Lockdown Indicator	1.637***	1.658***	1.461***	1.634***
	(0.0178)	(0.0177)	(0.0177)	(0.0177)
Lagged Daily Incidence	0.023***	0.023***	0.026***	0.028***
	(0.0002)	(0.0002)	(0.0002)	(0.0002)
Postal Code F.E.	Y	Y	Y	Y
N	26784	26784	26784	26784

Table 5: Estimated coefficients of Poisson regression model for postcode-level COVID-19 incidence. Standard errors in parentheses. The p-values associated with each coefficient are less than 1e-15.

Table 5 reports the estimated coefficients the panel model. 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). All lagged transport spending indicators are positive and highly significant, including car transportation. The interpretation is that there is a robust relationship between current disease incidence in postcodes, and transportation spending across all categories several weeks prior. Again, though, the effect of urban transport spending is particularly high. In the Poisson model the average treatment effect is the estimated coefficient value multiplied by the mean of the dependent variable, in this case 3.92. In these terms, a unit increase in lagged urban transport spending increases daily incidence by 0.49, while the corresponding numbers for overall and car spending are 0.14 and 0.27, respectively.[7]

Overall, then, we observe that transport spending is a good proxy of mobility, as well as a predictor of disease. Since we also observe that residents of poorer postal codes travel more during the workweek in lockdown, the overall implication is that they are also more subject to disease risk than residents of richer neighborhoods. This is another sense in which card data helps uncover the distributional impact of COVID-19, in this case on expected health outcomes instead of consumption behavior.

3 Discussion

The increasing abundance of detailed and granular financial transactions stored by banks and payment systems is potentially transformative for economic measurement. National statistics agencies are at the earliest stages of engaging with nontraditional data, and our results suggest the value in complementing traditional surveys with naturally occurring transaction data. These efforts are particularly important in low- and middle-income countries, where more standard high-quality and high-frequency indicators of consumption maybe too costly to produce.

Transaction data also provides timely signals to policymakers about the impact of economic shocks and policy interventions, which is especially important at times of high uncertainty and rapid change

[7] We do not observe whether a unit of urban transport spending generates more or less movement through space than a unit of gasoline spending, which would also be an important input into a model of disease risk and spending.

as during the current COVID-19 crisis. We draw three lessons from the first Spanish lockdown in early 2020 from BBVA card data that are more broadly relevant as many countries in the world again limit economic activity to control disease spread.

First, the closing and opening of establishments had a dramatic effect on spending, which reacts abruptly to both measures. On the other hand, social distancing policies and restrictions of capacity have a much more limited effect. This highlights that lockdown policies are not an either/or policy. When countries ease out of lockdown, shop openings are important for stimulating economic activity, but capacity restrictions can be maintained for longer periods at relatively low economic cost while protecting health.

Second, 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 spending decline by more. The resulting increased savings suggests that private households in high-income neighborhoods accumulated assets during the crisis that could help finance the large government deficits resulting from employment support and other measures.

Third, detailed transaction data on transportation and commuting expenditures reveals that residents of poorer neighborhoods are more likely to travel during the workweek during lockdowns, and that this correlates with higher disease incidence. Importantly, though, the mode of transportation appears to affect disease. Investment in additional safety measures for users of public transportation, and in transportation infrastructure that promotes social distancing without increasing pollution (e.g. cycle lanes), could mitigate these impacts.

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.

4 Methods

Transaction Data

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

In the supplementary information we present some summary statistics of this large dataset. In total, 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). 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.

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. In the Supplementary Information we 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.

When analyzing these data we calculate Y-o-Y growth 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 Wednesday with first Wednesday, etc.). We then measure the 2019-2020 Y-o-Y growth for the same day of the week. This controls for weekly seasonality to some extent, but to further control for weekly variation in some of the graphs we use the 7 day moving average. In figure 2 is particularly important to control for day of the week variation, so we show the residuals on day of the week dummies.

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.

Postal-code level Data

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").

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

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 inhabitants of 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.

Further analysis of lockdown easing

The standard difference-in-differences analysis for the effect of lockdown easing 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 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. 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 period 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.

A second concern that arises, as articulated in Goodman-Bacon (2018), is that the treatment effect may not be stable over time. In our context, this means that the expenditure effects of lockdown easing may be 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, when all Spanish provinces remained in either Phase 0 or 1. Thus, within this subsample, we have three groups of provinces: early-switchers, easing into Phase 1 on May 11th, late-switchers on May 18th and never-switchers (till May 25th). Based on this classification we can use the Goodman-Bacon (2018) 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.

Poisson regression model for disease outcomes as function of spending

Let $y_{i,t}$ be the number of new COVID cases in postal code i on day t, and $x_{i,t}$ be the level of transport spending of postal code i resident on day t, measured in EUR 1,000 units. The within-postcode disease predictor we use is $x_{i,(t-28):(t-14)} = \frac{1}{14}\sum_{\tau=t-28}^{t-14} x_{i,\tau}$, which accounts for two aspects of transport spending. First, it is potentially noisy, so averaging over multiple days helps dampen the impact of idiosyncratic, day-level spending variation. Second, it accounts for the incubation time of coronavirus before the onset of COVID-19, as well as delays in testing and the recording of cases in official statistics. Our construction focuses on the health impact of transport spending on a given day on disease outcomes two-to-four weeks later. Averaging also helps control for the uncertainty in the exact timing of the health effects.

We model $y_{i,t}$ using a Poisson regression model[8] with mean

$$\mu_{i,t} = \beta_1 y_{i,t-1} + \beta_2 x_{i,(t-28):(t-14)} + \beta_3 \text{Lockdown}_t + \gamma_i.$$

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 any time-invariant postcode characteristics that might affect disease outcomes or transport spending.

Data Availability

- In the Supplementary Materials we make available codes and data - both expenditure series and necessary covariates - pertaining to the national-, province- and broad category- level data, allowing researchers to fully replicate key COVID-19 results in the paper (Figure 1, Figure 2, Figure 4 and Table 1) and to conduct their own national, subnational or expenditure-category analysis in this context. Note, however, that this study builds from proprietary card transaction data from BBVA, a Spanish commercial bank. Both the individual-level card source data and aggregations at the zip code level or highly disaggregated category level involve highly sensitive personal information about customers and/or may disclose proprietary commercial information on local bank activities. Therefore, we are unable to render data fully publicly accessible beyond what is deposited in the Supplementary Materials. In particular, we are unable to publicly share replication materials involving: historical time series for Spain, cross-country expenditure data, detailed category of expenditure information or zip code level data. Individual researchers interested in these more detailed data sets should direct their query to BBVA Research.

- The Spanish Household Budget Survey is publicly accessible data, and can be obtained from the web-page of the "Instituto Nacional de Estadística"

- The data on income at census tract level (CUSEC) from where the income at postal code level is calculated is also public, and can also be obtained from the web-page of the "Instituto Nacional de Estadística"

- The data on incidence of the pandemic at Madrid Health District level can be obtained from: https://www.comunidad.madrid/servicios/salud/2019-nuevo-coronavirus.

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

Acknowledgements

The authors would like to thank four anonymous reviewers and editors for their valuable suggestions. We also thank Raman Singh Chhina for superb research assistance, and Gary Koop for inspiring comments on a previous draft. Vasco M. Carvalho gratefully acknowledges funding from the Leverhulme Trust and the European Research Council, grant #101001221, MICRO2MACRO. Stephen Hansen gratefully acknowledges funding from ERC consolidator grant 864863. José V. Rodríguez Mora gratefully acknowledges the support of the UK Economic and Social Research Council (ESRC), award reference ES/L009633/1.

References

Abraham, K., Citro, C. F., White, G. D., & Kirkendall, N. K. (Eds.) (2018). *Reengineering the Census Bureau's Annual Economic Surveys*. Washington, DC: National Academies Press.

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. (2020). Consumer Responses to the COVID-19 Crisis: Evidence from Bank Account Transaction Data. CEPR Discussion Paper 14809.

Asger Lau Andersen, N. J., Emil Toft Hansen, & Sheridan, A. (2020). Pandemic, shutdown and consumer spending: Lessons from scandinavian policy responses to covid-19. Tech. rep.

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.

Bean, C. (2016). Independent review of uk economic statistics: Final report. Tech. rep., GOV.UK.

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.

Carvalho, B., Peralta, S., & dos Santos, J. P. (2020). What and how did people buy during the Great Lockdown? Evidence from electronic payments. Working Papers ECARES 2020-20, ULB – Universite Libre de Bruxelles.
URL https://ideas.repec.org/p/eca/wpaper/2013-307531.html

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.

Dingel, J. I., & Neiman, B. (2020). How many jobs can be done at home? Working Paper 26948, National Bureau of Economic Research.
URL http://www.nber.org/papers/w26948

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.

Glaeser, E. L., Gorback, C., & Redding, S. J. (2020). Jue insight: How much does covid-19 increase with mobility? evidence from new york and four other u.s. cities. *Journal of Urban Economics*, (p. 103292).
URL https://www.sciencedirect.com/science/article/pii/S0094119020300632

Goodman-Bacon, A. (2018). Difference-in-differences with variation in treatment timing. Tech. rep., National Bureau of Economic Research.

Jarmin, R. (2019). Evolving measurement for an evolving economy: Thoughts on 21st century us economic statistics. *Journal of Economic Perspectives*, *33*(1), 165–84.
URL https://www.aeaweb.org/articles?id=10.1257/jep.33.1.165

Miyauchi, Y., Nakajima, K., & Redding, S. J. (2021). Consumption access and the spatial concentration of economic activity: Evidence from smartphone data. NBER Working Paper 28497.

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.

Author contributions

Carvalho, Hansen and Rodríguez Mora analysed the data and wrote the manuscript. García, Ortiz, Rodrigo and Ruiz analysed the data.

Supplementary Information to: Tracking the COVID-19 Crisis with High-Resolution Transaction Data

Vasco M. Carvalho Juan R. Garcia Stephen Hansen* Álvaro Ortiz
Tomasa Rodrigo José V. Rodríguez Mora Pep Ruiz

July 22, 2021

*Corresponding author:stephen.hansen@imperial.ac.uk

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(%)
18-25	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.

1 Summary Statistics

In table 1 we present basic statistics of our main dataset along with comparisons of the age and education structure of BBVA clients with the whole of the Spanish population.

2 Time Series Validation and Subnational Correlations

2.1 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.[1] 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.

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.

Figure 1a 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. Figure 1b plots the INE Gas Retail Sales series against that the corresponding BBVA series, rescaled by the estimated elasticity.

[1] 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.

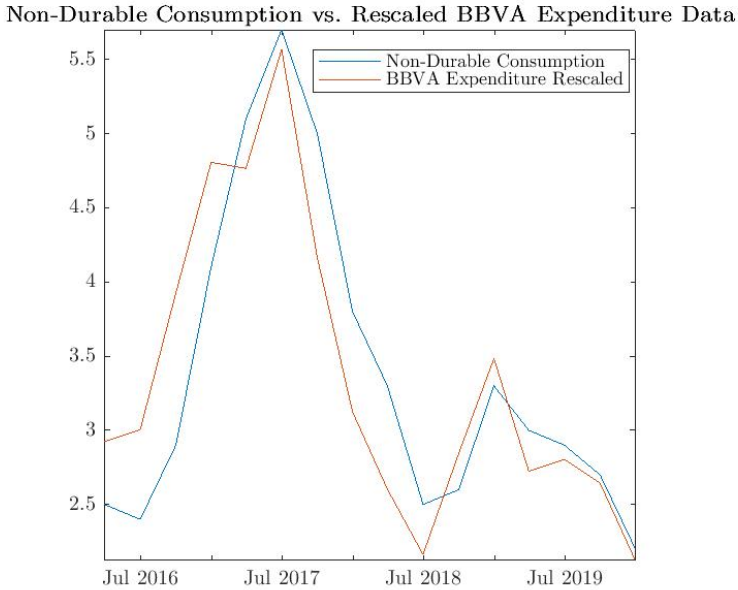

(a) Quarterly year-on-year growth rate of BBVA aggregate offline expenditures series vs. Quarterly year-on-year growth rate of Household Non-Durable Domestic Consumption in Spain. BBVA offline expenditure series are rescaled by 0.401, the elasticity of national accounts non-durable consumption growth to BBVA offline expenditure growth. The quarterly consumption series is sourced from the Spanish national accounts.

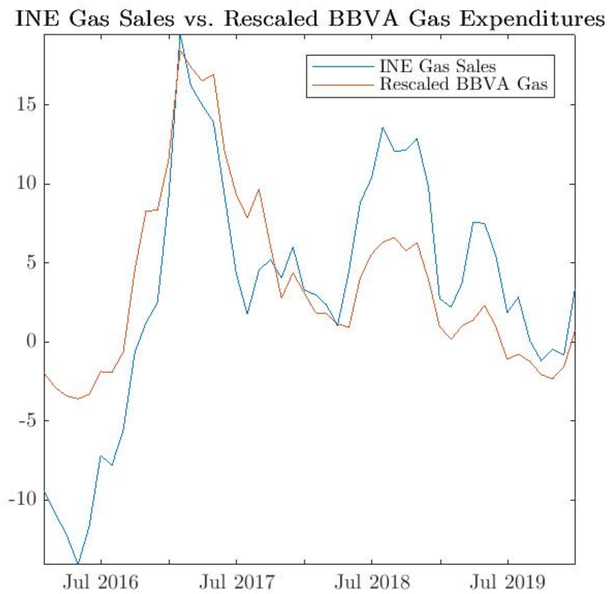

(b) Monthly year-on-year growth rate of BBVA expenditures at gas stations vs. monthly year-on-year growth rate of Spain's National Statistics Institute (INE) Retail Trade Index for gas stations. BBVA's gas expenditure series is rescaled by 0.346 All source data is nominal and not deseasonalized.

Fig. 1: Comparisons of time series of official statistics and BBVA consumption data.

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.

2.2 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 policymakers. 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 paper we 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.

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.

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[2]

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.

[2]Notice that the slope is larger than one. This is compatible with the poor spending a larger share of their consumption in utilities, housing and durable goods rarely paid by card. It is also compatible with BBVA having higher market share in richer areas.

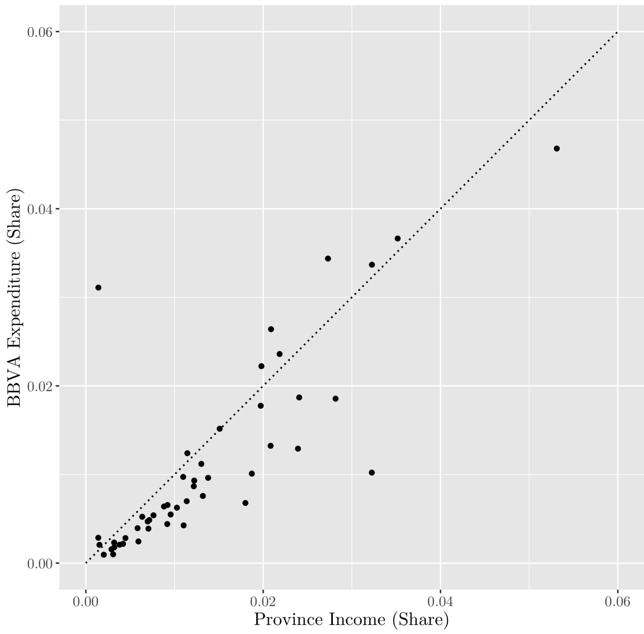

(a) Share of income of each Province in Spanish GDP (INE) vs. the Share of Expenditure in each Province, 2019 (BBVA data). It does not include, two outlying provinces, Barcelona and Madrid (the correlation is higher when we include them).

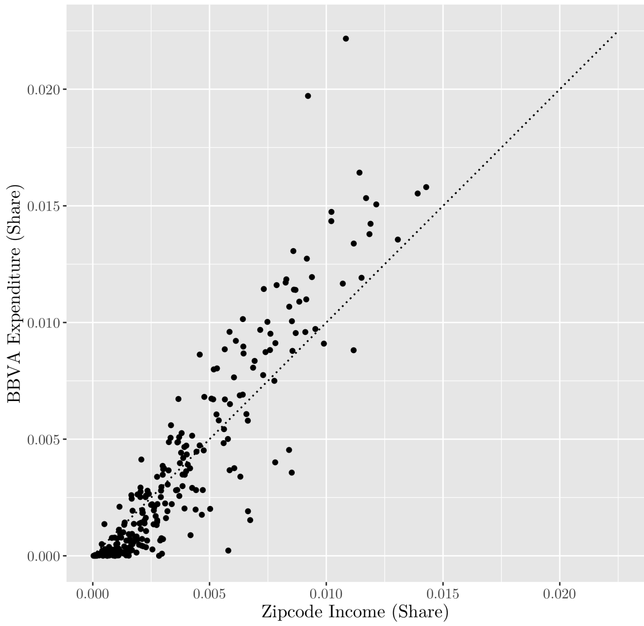

(b) Share of income of each Madrid postal code in Madrid's aggregate income (INE) vs. the share of expenditure by BBVA cardholders residing in each Madrid postal code over total expenditures by BBVA cardholders residing in the Madrid province (BBVA data), 2019. For postal code income we aggregate up from "secciones censales" level-data from INE.

Fig. 2: Cross sectional correlations of BBVA expenditure with income

3 Summary of Development of COVID-19 in Spain and Easing Phases

Spain has been one of the hardest hit countries by the COVID-19 pandemic. The first confirmed COVID-19 infection in Spain dates from the 31st of January 2020. 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. Throughout March the pandemic intensified throughout Spain, with 94,417 confirmed cases and 8,189 confirmed deaths by March 31st.

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 strict national lockdown policy which greatly restricted mobility and commercial activity. In particular, all citizens had to remain in their residences at all times, with exceptions only made for shopping for basic staples and medicine and for dealing 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. Note that, in spite of the large differences in disease incidence across provinces, this policy was implemented uniformly across the whole country.

Starting on May 4th, Spain initiated a phased lockdown easing process, aimed at a gradual re-normalization of activity. Below 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.

4 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 accounting aggregates. In this section of the SI we 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 Validation of expenditure categories

In table 2 we show the categories in which BBVA divides expenditure, marking in red those that were restricted by government decree during the lockdown. 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

Table 2: Description of Categories of Expenditure. We mark in red those categories that were restricted during the lockdown.

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		

housing), others are in practice not present in card spending (utility bill payments), while others have an ambiguous relationship with the categories in table 2 (non-alcoholic beverages). We are able to match 15 categories of ECOICOP with categories from table 2 with some confidence, which in total make up 48.2% of HBS spending in 2018 and 65% of BBVA card spending in 2019. Absent from card payment data there are, for instance, imputed (housing) rental values, actual (housing) rental payments, car purchasing, and utility bills which together make up for another 34% of national spending.

4.2 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). 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. 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 3 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

Table 3: Consumption Share Comparison by Demographics

		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

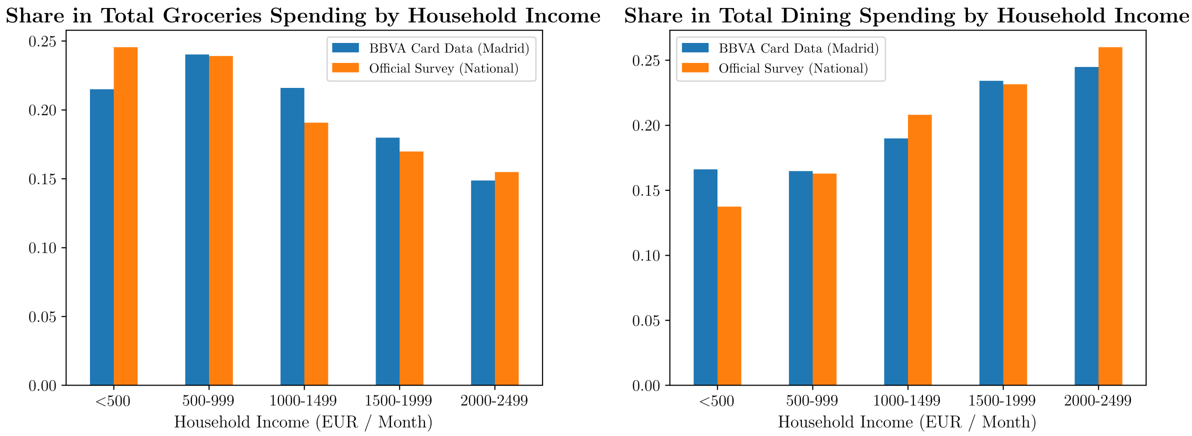

(a) Groceries (b) Dining

Fig. 3: Comparison of Consumption Shares by Income in Food Categories

more on housing services, utilities, etc.

4.3 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. 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. Our main income proxy is the income of postal codes within Madrid.

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 4 in the main text.[3] In general, the within-income-group correlation in consumption shares remains very high: it ranges from 0.83 to 0.95.

Figure 3 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:

[3]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.

1. It only applies to Madrid rather than Spain as a whole. 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.

2. it comes 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. We illustrate next how income affects consumption across Madrid. To do so we compute consumption shares in each of the categories listed in table 2 across Madrid postal codes, and then correlate the share of individual categories with income per capita. In the main text we list 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.

4.4 Consumption shares during lockdown

In the main text, we identify the top ten and bottom ten sectors according to market share growth during March 15 - May 4 compared to the pre-lockdown period. Figure 4 plots the daily evolution of these shares.

The aggregate evolution of these two sets of expenditure categories is illustrative of the dynamics of the crisis. 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 grew to an average value of 60%, while the worst performing categories made up on average just 1.6% of consumption. Again, these patterns quickly reverse after lockdown easing.

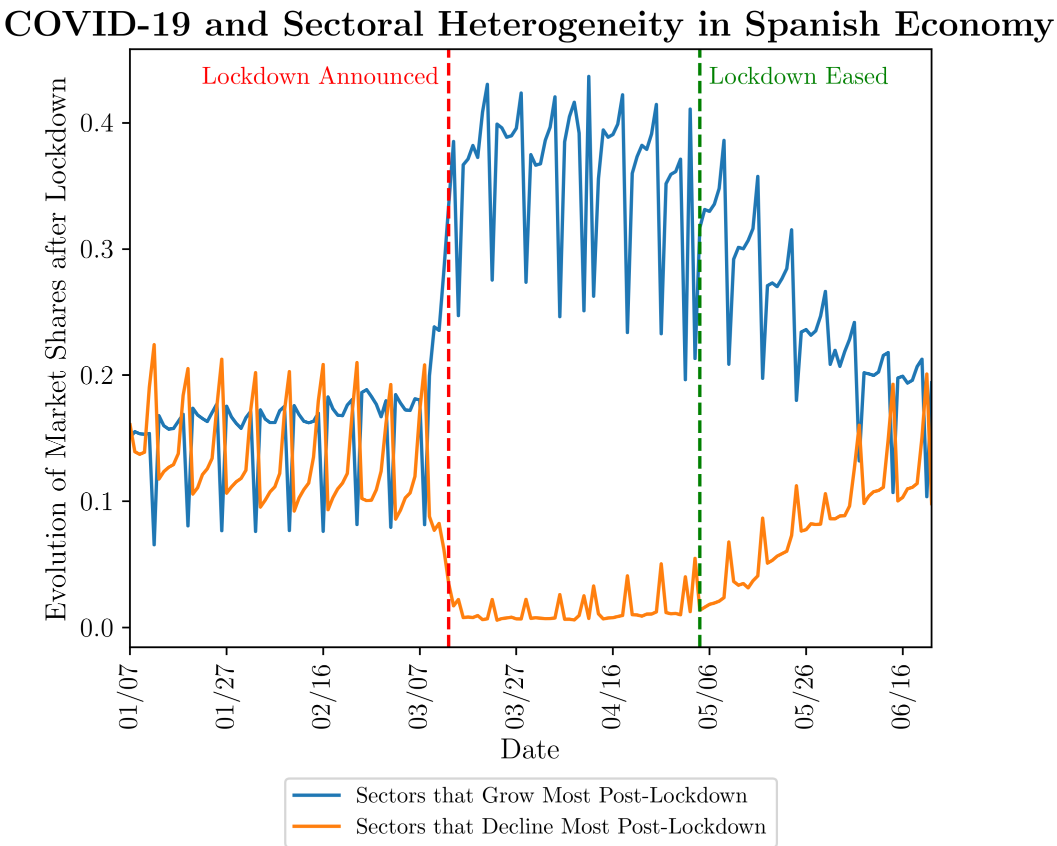

Fig. 4: Evolution of the market shares of the sectors that increased their share the most during the lockdown, and those that decreased their share the most.

5 Evolution of Consumption Across Income Groups and level of incidence of the Pandemic.

To quantify these effects more rigorously, and control for disease dynamics that might also drive neighborhood-level spending,[4] we use more formal statistical analysis. In Table 4 we present a set of panel regressions of Y-o-Y daily expenditure growth across postal districts on measures of income and incidence of COVID-19 within the district (in the Supplementary Information we detail how we calculate daily incidence at postal code level). These regressions not only confirm the that wealthier neighborhoods were the ones experiencing the largest fall in expenditure. They additionally suggest that areas more affected by the pandemic suffer larger declines in expenditure. All regressions include day fixed effects, thus controlling for the effects of any common factor in spending across Madrid (such as lockdown policies).

	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 4: 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.

In column (1) of table 4 we correlate contemporaneous new confirmed cases with the change in aggregate expenditure. 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, locations most affected by COVID-19 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 across richer and poorer neighborhoods. In column (2) we include only day fixed effects, while in column (5) we include also postal code fixed effects, controlling for time-invariant unobservable differences across postal codes (and voiding the use of quintile dummies when not interacted with the

[4]Since all of Madrid was in the same lockdown phase during our period of analysis, different government policies cannot explain the results

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.

The tendency of richer postal codes to contract spending more is robust to unobserved, time-invariant effects across days and postal codes as well as to disease incidence. At the same time, disease incidence remains a significant predictor of spending in all regressions in which it is included even after *conditioning on the level of income and on the lockdown policy implemented.*

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

6 Mobility across income groups.

6.1 Food transactions and income during the pandemic.

An interesting issue related to mobility is the extent to which higher income households (or zip codes in our analysis) resort to online purchasing by more than poorer people in order to shelter from the disease. One illustrative exercise we have explored is to focus on one of the key consumption categories during the lockdown: food consumption, which in our data corresponds to categories 3-6 from Table 2 in this Supplemental Appendix. We have tabulated the total number of transactions for online food transactions during April 2019 and during April 2020 by zipcode. The figure below plots income per capita against food transactions per capita by zipcode.[5] The gradient of the regression line in 2020 is much higher than in 2019, which suggests that higher-income neighborhoods were shifting more of their food purchasing online. Thus the overall online activity of higher-income groups appears related to the purchasing of necessities.

7 Mobility measured by number of transactions

We explore here the raw number of transactions as an alternative way to characterize the relationship between mobility and income. One issue is that higher-income groups on average have higher consumption, which translates into more transactions. This is one of the reasons we conducted the analysis in the main text in shares since this measures relative spending in different locations. Instead, when we work with raw transactions we compare total amounts in April 2019 and in April 2020 to help isolate the effect of the pandemic from that of being higher-income.

The particular exercise we conduct is to tabulate the set of "offline transactions" (i.e. those in which the card swiped a point-of-sale located in a physical shop) by postal code of cardholder residence. We further divide these into transactions that take place in the same postal code as the resident lives, and into those that take place in outside zip codes. In Figure 6 we plot these tabulations, where both transactions and income are in per-capita terms.[6]

For total transactions inside home postal code, one observes almost no difference between April 2019 and April 2020. This suggests that the frequency of local shop visits did not change markedly during the pandemic relative to normal times (although the composition of spending presumably does). In contrast, there is a large difference in transactions outside home postal code. The income/transaction volume gradient is much less steep during the pandemic than in 2019. In combination with the other evidence in the paper, one interpretation is that residents of higher-income postal codes were more likely

[5]The number of transactions is based on an index value BBVA provided to us and is not the actual count; instead the relative values have meaning so that a zipcode with 2.0 has twice as many food transactions per capita as one with 1.0

[6]As above, the number of transactions is based on an index value BBVA provided to us and is not the actual count.

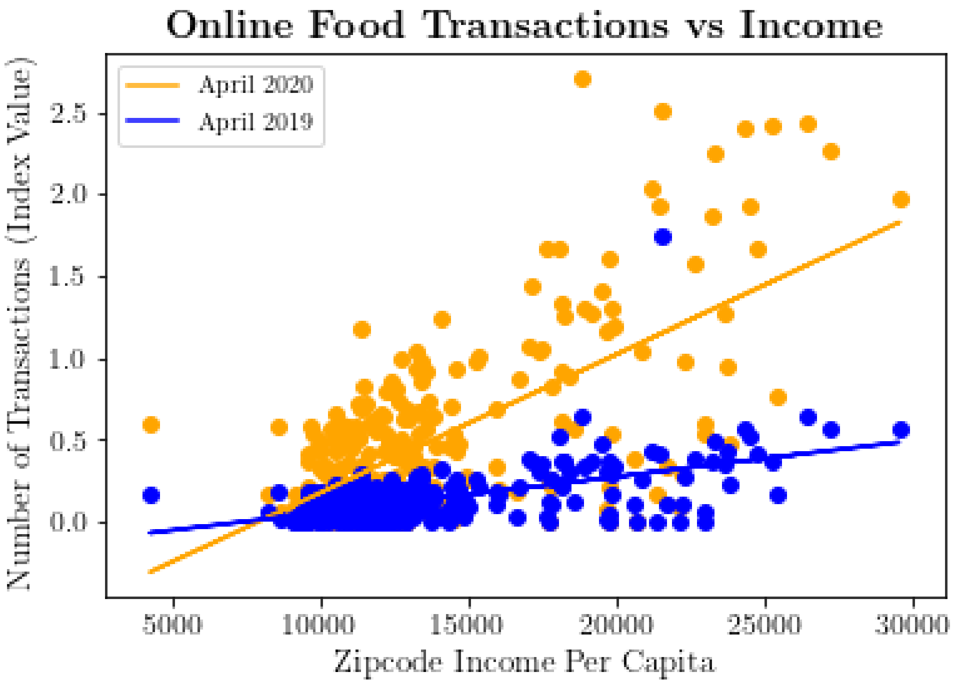

Fig. 5: Online purchases of food (vertical axis) by mean income of the zipcode (horizontal axis) in April 2019 and April 2020.

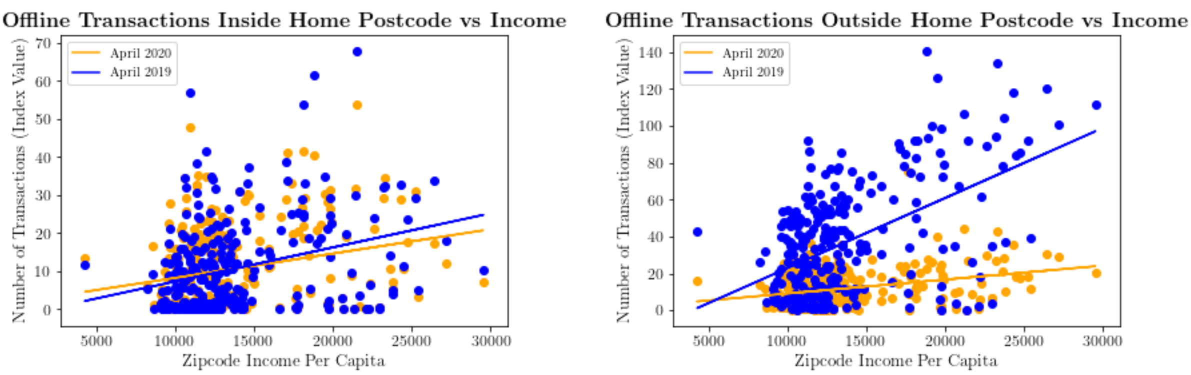

(a) Groceries

(b) Dining

Fig. 6: Offline number of transactions within and outside the zipcode of residence against the average income of the zipcode.

to stop commuting during lockdown. This would eliminate outside transactions that happen during the workweek at a faster rate for higher-income people. To the extent that outside-own-zip transactions are riskier in terms of disease than inside transactions, this evidence also suggests that higher-income residents were able to more effectively shield.

References

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.