

The Distributional Impact of the Pandemic

Sinem Hacıoğlu Hoke
Diego R. Känzig
Paolo Surico

November 2020



WID.WORLD
THE SOURCE FOR
GLOBAL INEQUALITY DATA

The Distributional Impact of the Pandemic*

Sinem Hacıoğlu Hoke[†] Diego R. Känzig[‡] Paolo Surico[§]

November 2020

Abstract

The top quartile of the income distribution accounts for almost half of the pandemic-related decline in aggregate consumption, with expenditure for this group falling much more than income. In contrast, the bottom quartile of the income distribution has seen the smallest spending cuts and the largest earnings drop but their total incomes have fallen by much less because of the increase in government benefits. The decline in consumers' spending preceded the introduction of the lockdown, whose partial lifting has triggered a stronger recovery in sectors with a lower contact rate. The largest spending contractions are concentrated in the most affluent regions. These conclusions are based on detailed high-frequency transaction data on spending, earnings and income from a large fintech company in the United Kingdom.

JEL classification: D12, E21, G51 and H31.

Keywords: spending, earnings, income, benefits, heterogeneity, pandemic.

*This paper updates and supersedes the CEPR Discussion Paper No. 14733 "Consumption in the Time of Covid-19: Evidence from UK Transactions Data". We are grateful to Andrea Galeotti, Atif Mian, Elias Papaioannou, Amir Sufi, Matthew Waldron and Tom Waters for useful comments and suggestions. We thank Francesco Amodeo and Sebastian Hohmann for valuable research assistance. The views expressed are those of the authors and do not reflect those of the Bank of England or any of its Committees. Authors do not have any personal conflict of interest with the company that have provided the data. Surico gratefully acknowledges financial support from the European Research Council (Grant 771976). All data have been anonymized at the source by the data provider.

[†]Bank of England. King's College, London (Data Analytics for Finance and Macro). E-mail: sinem.hacioglu@bankofengland.co.uk, Web: www.sinemhaciogluhoke.com.

[‡]London Business School. E-mail: dkaenzig@london.edu. Web: www.diegokaenzig.com

[§]London Business School and Center for Economic and Policy Research. E-mail: psurico@london.edu. Web: <https://sites.google.com/site/paolosurico/>.

1 Introduction

The global recession associated with the Covid-19 pandemic has proved to be unprecedented in both breadth and depth. The health crisis has very rapidly spread over goods and labour markets, education and finance, eventually reaching every sector of the economy. Virtually regardless of what metric is chosen, whether about infection, employment, businesses or any other socio-economic measure, the size of the shock has already been an order of magnitude larger than anything seen in the post-WWII era.

Not only the aggregate effect has been large, affecting virtually all nations, businesses and households, but also the impact has been so heterogeneous across society that some groups have clearly lost out far more than others. This has led to both the emergence of new inequalities and the acceleration of existing ones: the health of older patients have become more vulnerable, the education of children from disadvantaged families has been more adversely affected than that of children from affluent households, and job security has become an even more significant concern among less skilled workers.

When it comes to economic outcomes such as consumption, earnings, income and government policies, there are many forms of heterogeneity, inequality and redistribution at play. Four prominent examples through which the pandemic may influence household behaviour and their standards of living can be detected: (i) along the income distribution, between poorer and richer families, (ii) over time, before and after the lockdown measures, (iii) across sectors, depending on health fears and contact rates, (iv) between geographical areas, alongside a number of local characteristics.

In this paper, we look at each of these four dimensions of heterogeneity, in an effort to provide a coherent framework to track the distributional impact of the pandemic in *real-time*. The motivation of our approach is twofold. On the academic side, household heterogeneity and distributional dynamics can help to uncover the channels through which macroeconomic shocks transmit to the real economy. On the policy side, immediate knowledge of what interventions are more effective and what groups of society respond more can support the design of targeted policies in real-time.

We use detailed transaction data from one of the U.K.'s largest personal financial manager, Money Dashboard (MDB). The app is a real-time account aggregator that collates the financial transactions of a user's current, credit and savings accounts, regardless of provider, within a single platform. The granularity of the data allows us to construct measures of spending, earnings and income at the

user-level as well as looking at sub-categories of consumption across sectors and regions.

Main findings. Our analysis highlights a few patterns. Along the income distribution, more affluent users have cut their expenditure by far more than any other group and well in excess of the fall in their income. Importantly, their contribution of 45% to the decline in aggregate spending during the pandemic is about 10 percentage points higher than their contribution to the level of aggregate spending in normal times. At the other end of the spectrum, low-income users have experienced the largest drop in earnings but the smallest decline in spending. In part, this reflects that they have benefited more than any other group from government benefits to boost their income. While our findings for the second quarter of 2020 are consistent with an increase in personal saving rate around 20% ([Office for Budget Responsibility, 2020](#)), they also suggest that higher income households may have contributed disproportionately to this increase in the saving rate.

Spending began to fall before lockdown measures and social distancing were introduced in March, suggesting that health fears and income uncertainty may have also contributed to the fall in aggregate consumption, over and above social distancing measures. Some of these measures were lifted in June, triggering a significant, but partial recovery. Across spending categories, services, and in particular sectors with a higher contact rate such as restaurant, travel, holidays and entertainment, have seen both the largest decline and the slowest recovery. On the other hand, on-line shopping and food-delivery spending have increased significantly during the pandemic.

As for regional heterogeneity, we project the year on year growth rate in local spending for 2020Q2 to area-specific measures of the health crisis (number of coronavirus-related deaths per 1,000 inhabitants), government support (share of furloughed workers) and income inequality (share of MDB users in the top quartile of the income distribution). While all measures bear some correlation with the pattern of local spending across geographical areas, only the higher-income indicator proves to be a robust, stable and significant predictor of local demand, implying that a 2% increase in the share of higher-income households in a geographical area is associated with a further 1% decline in local consumers' spending. This is consistent with the observation that expenditure is cut more at the top than at the bottom of the income distribution.

Related literature. Our analysis is related to a recent, international and uncoordinated effort to track the economy in real time using high-frequency financial

transaction data from bank accounts and survey data to shed light on spending patterns across time and space. [Baker et al. \(2020\)](#), [Cox et al. \(2020\)](#) and [Chetty et al. \(2020\)](#) document in great details the initial response of U.S. household to the pandemic, giving particular emphasis to the differences across political affiliation, income groups and sectors of the economy respectively. [Coibion, Gorodnichenko, and Weber \(2020\)](#) exploit the different timings of lockdown across U.S. states to identify pandemic-related cuts in spending as well as losses in income and wealth for half of their survey participants.

[Carvalho et al. \(2020\)](#) look at a large sample of households in Spain and focus on both the pervasive regional heterogeneity in the spending response to the lockdown and the distinction between on- and off-line sales. [Andersen et al. \(2020a\)](#) explore transaction level data from a large bank in Denmark and document a significant reduction in aggregate card spending. Subsequently, [Andersen et al. \(2020b\)](#) compare the implications of Denmark's strict measures against the virus versus Sweden where restrictions were mild, and find that the significant drop in activity has been caused by the virus itself regardless of the specific restrictions against it. [Bounie, Camara, and Galbraith \(2020\)](#) track consumer spending and mobility in France using individual card transactions and show that, although household expenditure declined significantly, online shopping partially offset the fall in aggregate consumption.

With these important studies, we share both the use of granular transaction-level data and the focus – for the United Kingdom – on the decline and recovery in household expenditure. The similarities of the consumption patterns across countries should reduce external validity concerns. Unlike most of these papers, but similar to [Andersen et al. \(2020a\)](#) for Denmark and Sweden as well as [Cox et al. \(2020\)](#) and [Chetty et al. \(2020\)](#) for the United States, we look also at earnings, income and government benefits which, together with detailed information on spending, allows us to track the evolution of the balance sheet positions of British households during the Covid-19 crisis. Furthermore, we highlight the large distributional changes and the significant heterogeneity in spending along the income distribution that has emerged as a result of the pandemic.

As for U.K. studies, [Bourquin et al. \(2020\)](#) focus on the income side of the MDB data but do not look at spending like we have done in [Hacioglu-Hoke, Känzig, and Surico \(2020\)](#). In independent work, [Chronopoulos, Lukas, and Wilson \(2020\)](#) use the MDB data to analyse the pattern of expenditure on groceries, dining and drinking, alcohol and gambling. Relative to [Bourquin et al. \(2020\)](#), we look *jointly* at spending, earnings and income, and therefore are able to document the pervasive heterogeneity in the consumption responses and balance sheet changes across

the income distribution of British households. Relative to [Chronopoulos, Lukas, and Wilson \(2020\)](#), we focus on a broader and more detailed set of expenditure categories and exploit the information on spending, earnings and income to uncover the distributional impact of the Covid-19 crisis. In subsequent work, [Delestre et al. \(2020\)](#) use the MDB data to study how different elements of the income protection offered during the crisis have impacted British households.

Structure of the paper. The paper is organized as follows. In Section 2, we describe the data source and report descriptive statistics for some main variables of interest. The findings on spending, earnings and income over all users and across income groups are reported in Section 3, which also discusses the role of government benefits. The focus of Section 4 is on the effects of the lockdown measures imposed by the government and its partial lifting in the subsequent months. This section also presents evidence across different sectors. In Section 5, we look at spending heterogeneity across regions and relates that to indicators of the health crisis, government support and income inequality. In the Appendix, we discuss further features of the MDB data, its representativeness and timeliness relative to other surveys and national accounts and provide additional analyses and results.

2 Data description

The data we use in this paper are provided by Money Dashboard (MDB), a free online personal financial management company operating in the U.K. Their main product is an app that gives its users the flexibility to link multiple accounts (current, savings or credit card accounts) and provides them with a set of tools for categorizing and keeping track of their income and spending. The number of users has increased significantly over time, from around 10,000 in 2012 to close to 100,000 by the end of 2019.

The raw dataset is composed of transactions. All incoming and outgoing transactions in the accounts linked to the app are collected with the transaction date, the transacted amount, whether it was a credit or debit, a transaction description, a user identifier and the account that has been used.¹ The users in the dataset have been anonymised. In terms of user characteristics, we observe a partial postcode, the gender and the year of birth.

¹When a user registers with the app and links her accounts, the app downloads up to three years of past transactions from all linked accounts. After the initial registration, the app continuously receives updated transactions. However, the users must re-authorise MDB to continue to download new transactions every 90 days.

Using machine learning techniques, the data provider categorizes the transactions into almost 200 categories and thousands of merchants. The categories are as detailed as cinema, taxi, insurance, parking, dining out, mortgages etc, while merchants cover the most known businesses in the U.K., e.g. Tesco, Sainsbury's or Waitrose for groceries, Pret a Manger, Eat or McDonalds for take-away or snacks, and Amazon and eBay for on-line shopping. Similar to outflows, we observe incoming payments such as salary and interest income.²

Based on these classifications, we construct monthly measures of consumption expenditure, earnings and income. On the expenditure side, we follow the classification of individual consumption by purpose (COICOP) by the United Nations Statistics Division as closely as possible to create broader categories of expenditure, such as total, non-durable, durables, and services expenditure. To be able to analyze the effects of government policies such as the lockdown, we also construct measures of expenditure at the weekly frequency. We express all consumption and income measures in real terms by deflating the series by the U.K. Consumer Price Index (CPI). For more information, see Appendix A.2.

For our main analysis, we focus on a balanced panel of users for the period between January 2019 and June 2020. We restrict our sample to users that have consistently used the app throughout this period and have a fully updated set of transactions as of 5 July 2020. Furthermore, we exclude users with implausibly low/high expenditure and income, as well as users that use the app for business-related purposes. The final dataset is composed of 8,365 users and is made up of 15,261,988 transactions in the period between 1 January 2019 and 5 July 2020. For more information on the selection of users, see Appendix A.1.³ To capture changes in both the intensive and extensive margin, we include all these users in our analysis, even if they have no earnings or expenditure in a specific sub-category in

²The MDB algorithm is able to tag around 75% of all transactions. Users also have the option to manually add or change tags. However, in our analysis we rely exclusively on the automatic tagging provided by MDB. The main reason behind this choice is that, unlike the user tagging, it minimizes measurement errors, avoids recollection biases and ensures a timely and consistent system of categorization across users and time. However, our results are robust to using the user tags instead or if we construct alternative measures of income and expenditure using all transactions, including untagged, net of identifiable transfers. Furthermore, we have also verified that the performance of the tagging algorithm has remained stable over the course of the pandemic, both in the aggregate and by income group (see Appendix D for more detail).

³The number of users in our panel is considerably smaller than the overall number of users in 2019-2020. There are a number of reasons for this. First, we impose quite strong restrictions on the users to ensure that we do not include users with incomplete accounts. Second, we lose some users that have not yet re-authorized MDB to download the most recent data. Third, over the past 6 months, MDB has been migrating users from using Yodlee to their own direct integration with the banks' Open Banking API. While this change is undoubtedly positive in the longer term, MDB has lost some of their users during this process. This illustrates again the importance of focusing on a set of users that we observe consistently throughout the sample of interest.

a given period.

It should be noted that while our dataset contains information on users' ATM cash withdrawals (see Section 4), it only covers electronic transactions and payments. However, cash use among U.K. consumers has been falling significantly in the last two decades: while in 2008 cash payments accounted for two thirds of all payments, in 2018 the share of cash payments has declined to 28%. As of 2019, 98% of adults holds a debit card while around 65% have a credit card. Due to the widespread use of contactless payments, the share of cash is predicted to fall further to 9% by 2028.⁴ Hence, our dataset appears representative of the current transaction environment.⁵

Table 1 presents some descriptive statistics on the users in our sample. On average, users link four bank accounts from two different banks to the app and have around 100 transactions each month. The median net salary among the users is £2,327 and their median age is 36. For an average user, approximately a third of the total monthly expenses is on non-durable goods, and more than half of it is on services. Groceries account for about half of the non-durable expenses while spending on restaurants is on average £110 per month with a considerable dispersion within different percentiles. The average mortgagor faces over £900 of monthly repayments and the average renters pays over £700 each month.

In Appendix A, we present further information on the demographic and geographical features of the data. As shown in Figures B.1 and B.2, on average, Money Dashboard users are relatively young and more likely to be based in the Greater London area. Once we normalize the number of users by the population of the region they live in, however, the geographical distribution is more even (Figure B.3).

To evaluate the extent to which the MDB data may be representative of the British population, we explore both cross-sectional and time-series dimensions of the MDB data in Appendix B.2. Along the cross-section of users, we show that both (i) the distribution of 2019 after-tax income and total expenditure and (ii) the share of 2019 spending for the main COICOP consumption categories from the MDB data align well with the same statistics based on the Living Costs and Food Survey (LCFS), which has been constructed to be representative of British households. As for the time-series dimension, we compare the quarterly

⁴U.K. Finance, Payment Markets Summary 2019, <https://www.ukfinance.org.uk/sites/default/files/uploads/pdf/UK-Finance-UK-Payment-Markets-Report-2019-SUMMARY.pdf>.

⁵Reassuringly, some of our findings in Section 4 also line up with the Office for National Statistics' (ONS) March release of Retail Sales Inquiry on the changes in sales. Source: Office for National Statistics - Monthly Business Survey - Retail Sales Inquiry <https://www.ons.gov.uk/businessindustryandtrade/retailindustry/bulletins/retailsales/march2020>.

household final consumption expenditure from the national accounts as published by the ONS to an equivalent time series obtained from aggregating the financial transactions by a rotating panel of MDB users. The two time series exhibit a similar trend in level and share most turning points in the year-on-year growth rates version, with the latter displaying a remarkable correlation of 0.79 over the 2016-2020 period.

Finally, in Appendix [A.3](#), we provide details on the differences and similarities between our method and two more traditional sources: the LCFS and the Understanding Society Survey (USS). In short, two main advantages of these are representativeness of the British population and recording spending both in electronic and cash forms. The two main advantages of our method are timeliness as opposed to the six to twelve months of delay typically associated with surveys, and the fully automated and electronic data collection in real-time which minimize serious non-classical measurement errors and recollection biases associated data collection in surveys (see for instance [Attanasio, Hurst, and Pistaferri \(2014\)](#) and [Pistaferri \(2015\)](#)). In summary, analysis based on financial transaction data may present a useful complement to the well-established approaches of measuring aggregate consumption in national statistics based on survey data.

3 Spending, income and benefits

In this section, we explore the evolution of spending, earnings and savings in an effort to track how the balance sheet of British households has changed in response to the pandemic. We start from looking at the median values across all users in our sample and then present the analysis on heterogeneity across income groups. A main take away is that the spending of top earners has declined the most and far in excess of both their earnings and income. But for households at the bottom of the income distribution, the relatively smaller drop in spending has been so far smaller than the fall in earnings but larger than the decline in income. In the final part of this section, we explore the role of government benefits to account for the dynamics of income and spending of the less affluent households.

3.1 The overall impact of the pandemic

A simple way to capture the general trends on the household balance sheets is to summarize the evolution of inflows versus outflows for all users in our sample. This is the goal of [Figure 1](#), which reports the year on year growth rate in median monthly spending (dashed orange line) and median monthly income (solid blue

line) for the first half of 2020. To net out regular growth and age effects, we express the series as indices, normalized to 100 in January before computing the year on year growth rates.

The evidence in Figure 1 reveals that the start of the pandemic and the consequent lockdown in March 2020 are associated with a drop of over 10% in the median expenditure relative to March 2019. The fall has deepened in the subsequent months, reaching its trough of almost -40% in April before beginning the recovery in May. By the end of the second quarter of 2020, the level of expenditure is still significantly subdued.

The pattern of expenditure stands in contrast to the evolution of income. The blue solid line in Figure 1 reveals a more gentle fall throughout the second quarter of 2020, possibly influenced by government support policies such as the Coronavirus Job Retention Scheme. But there is little sign of recovery in June, when the median income records an almost 10% fall relative to its 2019 counterpart.

Taken together, the time profiles of expenditure and income suggest that overall the saving rate of British households has increased significantly over the start of the pandemic, with an average gap between median income and median expenditure in excess of 20%. However, as we explore in the next section, the summary statistics reported in Figure 1 hide pervasive heterogeneity across income groups.

3.2 The heterogeneous responses across income groups

The impact of the pandemic shock on health, education, labour market access and many other socio-demographic indicators has been very heterogeneous across society creating new and reinforcing old inequalities (Blundell et al., 2020). As the access to personal and professional opportunities are often correlated with the level of income, in this part of our analysis we decompose the findings of the previous section by income groups.

In Figure 2, we present the median monthly spending (top panel), earnings (middle panel) and income (bottom panel) for four income groups that roughly coincide with the quartiles of the ex-ante income distribution in 2019 among the users in our sample. These are people with after-tax income below £20,000 (in red), between £20,000 and £30,000 (in green), within the £30,000 to £40,000 range (in orange) and above £40,000 (in blue) respectively. Earnings refer to their main salary whereas income includes also financial income, rental income, pensions and benefits, among other sources.

The top panel of Figure 2 makes clear that the fall in total spending is much more pronounced among higher-income users, who have witnessed a trough

of roughly 45% in May before a partial recovery in June to 35% less than the pre-pandemic level. At the other end of the income distribution, less affluent households have cut their expenditure by almost 30% in April, 25% in May and are in June over 15% below their 2019 levels. The spending of the two groups in the middle of the income distribution falls in between the extremes of top and bottom groups, with the ranking of their consumption decline mirroring the ranking of their 2019 income levels.⁶

To appreciate the importance of each group spending for aggregate consumption and to place the expenditure patterns during the pandemic in historical context, in Table 2 we record – for each set of users along the ex-ante income distribution – their contribution to (i) the overall spending *level* in 2019Q2 and (ii) the overall spending *decline* in 2020Q2. A main take away from this exercise is that the contribution of higher-income users is disproportionate: while in normal times (as measured by the same quarter in the previous year) users with after-tax income above £40,000 contribute to about 35% of the level of aggregate spending, during the pandemic they have driven about 45% of the overall decline. The increase of 10 percentage points in the contribution of high-income users to the aggregate spending pattern seems mirrored by a similar decline in the contribution of users at the bottom of the income distribution, whose spending share declines from about 18% in normal times to 9% during the pandemic.⁷

The middle panel of Figure 2 explores the possibility that the heterogeneity in spending reported above may reflect heterogeneity in earnings. In fact, users at the bottom of the income distribution are experiencing the largest fall in salary, which are in May and June 2020 more than 20% below of what they were in 2019. This is substantially larger than the 10% drop in earnings faced by the middle group or the 15% witnessed by higher-income users. The bottom panel, however, reveals that the ranking of groups change if one looks at income instead. While for all other groups the income decline is quantitatively similar to the earnings decline in the middle panel, this is not the case for low earners. Users at the bottom of the income distribution have experienced an income fall of around 12% in June against the backdrop of an almost double earnings decline in percentage terms. In Section 3.5, we discuss the crucial role of government benefits to account for this

⁶This is consistent with the evidence for United States in Cox et al. (2020), who also find that higher-income households have cut expenditure much more than lower-income families.

⁷As discussed by Cox et al. (2020) for the U.S., a possible explanation for the finding of a larger spending decline among top earners is that their basket features a larger shares of non-essential goods and services, which are typically either easier to postpone or whose provision has been disrupted more during the pandemic. Crawford et al. (2020) show evidence of a larger prevalence of non-essential items also in the consumption basket of British top earners. In Section 3.4, we study this explanation in more detail.

pattern.⁸

3.3 Saving patterns

Having documented the evolution of spending, earnings and income across the income distribution, we are now in the position to go back to the saving patterns, as implied by the difference between the overall decline in spending and income reported in Figure 1 for all users. Which group is driving the average effect? The answer is provided in Figure 3, whose Panel A depicts for each income group the percentage decline in spending (green bar), earnings (blue bar) and income (orange bar) in June 2020 relative to the previous year. The main inference to draw from this chart is that the decline in earnings exceeds the decline in expenditure only for users at the very bottom of the income distribution. For all other groups, the pandemic shock is associated with a significant positive gap between the decline in earnings and the decline in spending, which is increasing in the level of income. On the other hand, looking at the orange bars reveals that also users with after-tax income below £20,000 have faced a fall in income that is smaller than the fall in their spending. It should be noted, however, that the change in personal savings implied by the gap between the drop in income and the drop in spending for this group is smaller than for users with income in the range £20,000-30,000 range and is only a fraction than for the groups above the median income.

Another way of inferring the pattern of saving *flows* across British households during the pandemic is provided in Panel B of Figure 3, which reports the median of the individual user's saving rates for each income group. The individual saving rate is computed based on all incoming and outgoing transactions (including both tagged and untagged transactions net of identifiable transfers, see Appendix D for more information). As such it is likely to be a noisier measure of personal saving rates than the inferred measure from Panel A. Notwithstanding the different measurement strategies, the bottom chart of Figure 3 paints a picture similar to the top chart. More affluent households have enjoyed significantly higher saving rates during the pandemic, in the ballpark of 20%. This is in stark contrast with users at the bottom of the income distribution, whose individual saving rates have averaged around 5% and, as we will show, have been able to avoid significant dissavings mostly thanks to the government support. The pattern of savings

⁸In Appendix C, we also explore spending heterogeneity along other dimensions such as age, housing tenure and gender on the one hand, and receiving any type of government benefit, being a pensioner and having children on the other hand. The largest variations are across housing tenure status and benefit receipts but appear quantitatively smaller than when we split the sample along the income distribution.

across income groups presented in this section chimes with independent evidence provided by Haldane (2020) based on an on-line survey conducted by Ipsos Mori on behalf of the Bank of England in 2020Q2.

3.4 The impact of government restrictions

A possible reason for the relatively larger spending decline among higher income users is that the consumption basket of this group of households may be more tilted towards non-essential goods and services, whose purchases have been physically more restricted (than essentials) by the lockdown measures imposed by the government. To investigate this hypothesis, we group consumption into three categories depending on whether the associated physical transactions were unrestricted, partially restricted or fully restricted during the Spring 2020 lockdown (see Table A.2 in Appendix A for a full classification of transaction types). Before exploring the evolution of spending patterns across income groups along the restricted/unrestricted categorization, it is useful to note that, in a typical year such as 2019, high-income, middle-income and low-income MDB users tend to spend about 42%, 39% and 37% respectively of their total consumption on goods and services whose physical purchase has been restricted during the lockdown.

In Figure 4, we document how the monthly spending in sectors that were either restricted during the pandemic (left panel) or unrestricted (right panel) has evolved over the course of 2020, for each income group considered previously. A comparison of the two charts suggests three considerations. First, the spending decline in unrestricted sectors was much lower than in fully restricted sectors (about half of the decline on average). Second, while the ranking of the spending drops among the four income groups is the same under either categorization (the drop is increasing in the pre-pandemic income level), higher income households tend to cut their expenditure relatively more in unrestricted sectors than poorer households. The relatively stronger response in combination with the higher pre-pandemic share in restricted expenditure can likely account for part of the heterogeneity in the total expenditure responses, however, it appears that the channel is not strong enough quantitatively to account for the bulk of the heterogeneity by income group. Third, to the extent that the spending drop in unrestricted sectors cannot be attributed (by construction) to the physical restrictions imposed by the government, the results in Figure 4 suggest that at least half of the consumption decline in restricted sectors (or two thirds of the drop in overall consumption) may be due to fears and uncertainty about the evolution of the pandemic rather than to the government restrictions per se.

3.5 The role of government benefits

In a previous section, we have seen that users at the bottom of the income distribution experienced a significantly larger drop in earnings than in income. In this section, we explore what category of income may account for this finding. Unlike high-income users, people making less than £20,000 a year are unlikely to have any financial or rental income. This implies that the gap between earning and income among less affluent households is likely to reflect pensions, allowance and benefits, with the latter being related to either family, work, low-income or health. In our sample, there has not been any significant increase in neither the intensive nor the extensive margins of family benefits during the second quarter of 2020.⁹ This leaves all other benefits, and in particular those related to work and having low-income, as the most likely candidate to explain the gap between income and earnings at the bottom of the income distribution. For a lack of a better term, we will loosely refer as 'government benefits' to this broader group that excludes family-related benefits.

In the United Kingdom, there are three main programmes through which adults can access working-related benefits: the long-standing (i) Job Seeker Allowance (JSA), (ii) Employment Support and Allowance (ESA), and (iii) the more recently introduced Universal Credit (UC). The latter is phasing out most other forms of benefits, including JSA and ESA, and is the largest of the three.¹⁰ It should be noted that many users may have been already UC recipients before the pandemic because of personal disability, other health conditions or employment and family circumstances. This implies that when looking at the role of benefits is important to consider both the extensive margin of new applications as well as the intensive margin of larger disbursements to existing recipients of (non work- or income-related) universal credit.

We start investigating the role of government benefits in the top panel of Figure 5, which presents for each income group the share of users that witnessed an increase in their government benefits disbursements. This includes also new applicants, who went from zero pounds to receiving the amount they were eligible for. The panel makes it clear that the bottom group (and to a lesser extent users with income between 20k and 30k) experienced a much sharper rise in the share of users receiving a higher government benefit disbursement, going from only 2.5% in February (relative to January) to 11.5% in June (also relative to January). The second largest increase is recorded for the second bottom group (green line)

⁹On April 6, child benefits were raised by 1.7% to £21.05 (£13.95) for the eldest (other) child.

¹⁰To be eligible for UC, the applicant and partner must have £16,000 or less in savings between the two. On-the-job applicants are allowed as long as their earning is below a fixed threshold.

at around 4% whereas only 2% of users with income above £30,000 have seen an increase in their government benefits.

To isolate the intensive margin of the benefits contribution to income in each group, the middle panel of Figure 5 reports the median value (in pounds) of the monthly disbursement associated with government benefits conditioning on users who have received some form of these benefits during 2019-2020. While there seems to be only a modest change at the top of the income distribution, the two bottom groups have seen an increase of more than 50% in the value of their median monthly disbursement. This represents a significant income change for these groups and therefore it may rationalize part of the finding related to savings in Figure 3.

A simple way to corroborate the important role played by government benefits in shaping the patterns of both income and spending at the bottom of the income distribution, is to focus on two otherwise similar groups of users: (i) those in the two bottom income groups who faced an increase in the disbursement of government benefits during 2020Q2 and (ii) those in the same two bottom income groups but who faced no increase in benefits over the same quarter (or received no benefits at all). To the extent that (a) benefits are an important driver of income and (b) income is an important driver of spending for less affluent households (i.e. they are more likely to be hand-to-mouth), we would expect to see a smaller decline and a faster recovery in spending for the group who benefitted from government support.

Prima facie evidence in support of this working hypothesis is provided in the last panel of Figure 5. Users at the bottom of the income distribution facing an increase in government benefits have experienced a far smaller trough in spending, around -20% , than users in the same income group that have not experienced an increase in benefits and by June their consumption went back to pre-pandemic levels. In sharp contrast, lower-income users with no increase in government benefits are, on average, still 20% below their spending level in the same month of last year. Taken together, the results in this section suggest that government benefits seem to be filling an important gap in the balance sheet of more disadvantaged households, allowing them to support their spending significantly in the face of a very material fall in earnings.

4 The effect of the lockdown on spending patterns

In the previous section, we have compared the decline in median spending to the fall in median earnings and income. As the latter are best measured at monthly

frequency, we have aggregated for each user all purchases within a month. In this section, however, we are interested in assessing the effect of the lockdown and its partial lifting, over time and across consumption categories. For this purpose, we focus on weekly average spending.¹¹ We start by looking at the main components of total expenditure, namely non-durable goods, services and durables. Then, we move to a detailed analysis of various sub-categories, especially within the service sector. As before, we express all series in real terms by deflating them by the U.K. CPI in the corresponding month.

4.1 The main consumption categories

On March 23, the U.K. government announced nationwide lockdown measures. However, softer measures had been in place earlier. On March 15, elderly and people in vulnerable groups were asked to self isolate. On March 16, the government issued an advice against all non-essential travels and going to pubs, restaurants and cinemas, and closed schools until further notice. Starting from March 20, pubs, cafes, restaurants, bars and gyms were officially closed. In the meantime, most of the employers in the U.K., especially in Greater London, asked their employees to work from home whenever possible.¹² On June 15, the government started to implement a series of further relaxations: high street retailers, department stores and indoor markets opened after almost three months of lockdown. The remaining businesses, restaurants, pubs and hairdressers were given green light to open after July 4 which therefore sits outside our sample.

Based on the narrative above, the main focus of this section will be on two dates: March 23 and June 15, which are denoted by two vertical lines in Figures 6 and 7. Both charts compare the average weekly spending in 2020 with the corresponding week of 2019. More specifically, the left column of Figure 6 report the average weekly expenditure on all goods and services (top row), non-durable

¹¹Whenever distributional aspects may be important, using the median (rather than the average) helps to hedge against outliers. So in the rest of the paper and especially in Section 3 where we compare earnings, income and spending at monthly frequency, we have reported median values also for spending. But for a sectoral analysis at weekly frequency, as we do in this section, we face the issue that many users make no purchase in several specific sub-categories over a significant fraction of weeks, implying that for those sectors the median value of weekly spending is often zero. Accordingly, for this section and this section only, we report average value of weekly spending (rather than median values). We have verified, however, that for the main categories of consumption in Figure 6, where we observe purchases from most users in most weeks, using average rather than median values makes only a small quantitative difference.

¹²Dingel and Neiman (2020) report that over 40% of the jobs in the U.K. can be performed at home while the Office of National Statistics (ONS) reports it as less than 30% of the workforce. Source: ONS Coronavirus and home-working in the U.K. labour market: 2019. <https://www.ons.gov.uk/employmentandlabourmarket/peopleinwork/employmentandemployeetypes/articles/coronavirusandhomeworkingintheuklabourmarket/2019>.

goods (second row), services (third row) and durable goods (bottom row), normalised to 100 in January. The red line refers to 2019 and the blue line to 2020. The right column shows the year on year growth rate in weekly spending indicators. To smooth out occasional spikes due to the slightly different timing of each week across the two years, we display the growth rate as a two week moving average. All charts exclude recurring bills such as utilities, council taxes and phone bills and, for non-durables goods, we further exclude groceries. Unlike its national statistics counterpart, our measure of total spending does not include actual or imputed rents. In the Appendix, we present a version of Figure 6 that also includes recurring bills.

A few interesting patterns emerge from Figure 6. First, the large decline in spending on non-durable goods and services started a couple of weeks before the vertical line on March 23 when the lockdown was introduced. Interestingly, the decline preceded also the introduction of social distancing policies on March 16, suggesting that at least partially, the initial drop in consumption has not been related to lockdown measures per se but more likely to fears and uncertainty on the evolution of the health and economic crisis. Second, the partial lifting of the restrictive measures to people's movements on June 15 has been associated with a significant recovery in average weekly spending on non-durable goods and services in the left column. It should be noted, however, the level of spending in the second half of June and beginning of July has been subdued relative to the pre-pandemic levels. Furthermore, when compared to the same weeks of 2019, the pick up over the last weeks of June and the first of July appears seasonal. Indeed, the year on year growth rate in the second column depicts a more modest recovery.¹³ Third, the pattern on durable goods sharply contrasts with the other categories: spending starts falling on the week of the lockdown announcement, reaches a trough in April before going back to pre-pandemic levels by early July.

4.2 Sectoral heterogeneity

In the previous section, we have documented that the spending decline has been more pronounced in the service and non-durable goods sectors whereas the recovery has been much faster for durable goods. In this section, we zoom in on some of the most salient subcomponents of the main categories of the previous section. The main goal here is to explore the hypothesis of some substitutability across goods and services.

¹³The MDB data display a bunching in the amounts transacted at the very beginning and end of each month. The first and last two days of June 2019 fell on a weekend, making the point of comparison lower and thus the year on year decline at weekly frequency less pronounced.

For the sake of exposition, in Figure 7, we report only the year on year growth rate of weekly expenditure, which by construction controls for seasonal factors. The top left panel report retail spending in blue and on-line shopping at Amazon and Ebay as orange line. The drop of the blue line is immediate, large and persistent whereas the rise of the orange line is delayed, large and short-lived, suggesting some form of substitution between high-street and on-line shopping. The top right panel shows that spending on alcohol and tobacco has received a sizeable boost during the lockdown and has made further gains after its partial relaxation, even though the time series is quite erratic.

Moving to the second row, the left column tells another story of possible substitution between restaurants (in blue) and food delivery (in orange). While spending on the former has collapsed even before the lockdown and then only mildly recovered after its partial lifting, food delivery has enjoyed persistent and increasing gains, feeding the notion that some of this substitution effects may extend beyond some further relaxations of the socially restrictive measures. In contrast, the increase in groceries around the start of the lockdown was short-lived, as visible in the right column.

In the left panel of the third row, we compare spending on recreation at home (orange broken line) and away/outside (blue solid line). Interestingly, despite the collapse in admissions to cinemas, theatres, concerts and museums, which is already evident before March 23, recreation at home has not picked up to compensate for that, implying limited substitution between the two. In contrast, DIY/Home repair expenditure on the right has experienced a volatile pattern with a significant rise around the start of the lockdown, a short-lived contraction in April and finally a robust pick up in may and June, being back now to the same level of last year.

The final two rows paint a picture that is very consistent with the rest of the sectors. Travel & Holidays, which is also characterized by a high contact rate like restaurant and recreation, witnesses a similar drop in demand, both in terms of size and persistence. In line with our findings for the U.K., [Abay, Tafere, and Woldemichael \(2020\)](#) use Google search data for a large panel of countries to document the temporal substitution from physical to online services and significant contractions of demand for hotels, restaurants and retail trade. We also observe a pronounced fall in spending on Transportation & Fuel but it shows a potentially quicker recovery. Finally, the last row suggests that there has been little substitution between cash holding (as measured by ATM withdrawals) and credit card spending as both fall by a similar magnitude and are still relatively subdued.

5 Regional variation

It is often argued that London and the South East are different from the rest of England and the U.K. along a number of socio-economic dimensions. In this section, we therefore explore the geography of the Covid-19 economic crisis across geographical areas of the United Kingdom. In particular, we focus on total expenditure and, for each statistical area, we compute the spending growth rate in 2020Q2 relative to the same period in the previous year. Two main results emerge from this exercise. First, the spending decline in Panel A of Figure 8 varies significantly across the U.K. statistical regions. Second, the most affected areas are Greater London and the South East with a -45% and a -37% decline respectively. In contrast, the least affected areas are Northern Ireland and Wales with drops of -10% and a -23% respectively.

Further insights into this spatial heterogeneity are gained from Panel B of Figure 8. This uses the first two digits of each user's home postcode to construct a heat map of the median percentage decline in total expenditure for more than one hundred smaller geographical areas. Darker (lighter) shades reflect larger (smaller) declines whereas the light grey denotes areas for which we do not observe a sufficient number of users to estimate the spending declines reliably. The heat map reveals pervasive heterogeneity in the economic costs across the country. The percentage declines range from between -6% and -16% in a few locations around Belfast, Wales and the Scottish Highlands to between -48% and -58% in a handful of areas across Greater London, Hampshire, Berkshire, Surrey and Aberdeenshire.

An important question amid this regional variation refers to what observables correlates well with the geographical patterns of spending. While it is impossible to draw any causal inference without a proper identification strategy, a simple correlation analysis helps to shed some light on the empirical regularities that may deserve further investigation. For the purpose of this descriptive exercise, we construct three indicators that relate to the health crisis, the extent of government support to retain workers and income inequality respectively. As for the health crisis, we compute the number of coronavirus deaths per 1,000 inhabitants using the data publicly available through the ONS website. In terms of government support across the country, we look at the share of furloughed workers published by HMRC. Finally, we use the MDB data to compute the share of high-income users (i.e. with after-tax annual income in excess of £40,000 during 2019) in each geographical area.

The top panel of Figure 9 reports the scatter plot between the percentage

decline in median spending over 2020Q2 relative to 2019Q2 for each area and the share of Covid-19-related deaths for every 1,000 inhabitants. A larger number of causalities (with respect to population) seems related to a larger spending decline, which is suggestive of a possible association between fears of contracting the virus and consumption patterns. On the other hand, the middle panel reveals that areas with a larger share of furloughed workers tend to experience a smaller drop in household expenditure, consistent with the notion that government support may have avoided an even deeper consumption crisis. Finally, the bottom panel of Figure 9 uncovers a stronger and significantly negative correlation between the share of higher-income users and the percentage change in median consumers' spending across U.K. geographical areas.

The analysis in Figure 9 is based on unconditional correlations and all three indicators bear some degree of comovements with each other. To gather some further insights into the most promising empirical regularities, in Table 3, we move one step further and project the percentage decline in spending for each geographical area onto the three measures of the health crisis, the government support and the income inequality. To elicit both the individual and marginal contribution of each dimension, we run a battery of regressions with all possible combinations of regressors, ranging for each of the three indicators in isolation, to regressions in pair and finally all three together.

A main result from this exercise is that only the share of higher-income users is a significant and robust predictor of a larger decline in spending, consistent with the analysis in Section 3. For every 2% increase in the share of users with income above £40K, local spending falls by a further 1%. In contrast, while the relationships between regional spending change and the share of coronavirus deaths on the one hand and the share of furloughed workers on the other hand have the expected sign, these variables clearly do not have the predictive power displayed by the share of higher-income users.

It is worth emphasizing that the correlations in this section are only suggestive and do not imply any causality. Still, they may help to bridge the findings of empirical studies that observe income at the individual level, like ours, with the results from very detailed and granular spending analyses, like for instance [Carvalho et al. \(2020\)](#), that only report income at the municipal or regional level. Using user-level data to construct a regional measure of income, we have shown that a simple regression analysis at the regional level – such as the one in Table 3 – is able to capture the statistical association between spending and income at the individual level that we have documented in Section 3.

6 Conclusions

Covid-19 is the greatest health and economic crisis of our time. Despite the extraordinary global efforts among policy makers, health researchers, data scientists and economists, many fundamental questions about its evolving and likely persistent economic effects remain unanswered. At the heart of this uncertainty lies the difficulty in acquiring real-time information that could help to identify which groups of society have been most affected and why. In this paper, we join a blossoming empirical literature that aims to track the economy in real-time by constructing indicators of economic activity using high-frequency transaction data from linked bank accounts.

We exploit the richness of household balance sheet data from a large fintech app provider in the U.K. to construct measures of spending, earnings and income at the individual level and look at the impact of the pandemic across the income distribution. A main finding is that the most affluent households have cut their spending the most during 2020Q2, with the top quartile accounting for about 45% of the decline in aggregate consumption. This represents an increase of 10 percentage points relative to the contribution of this group to total household expenditure during normal times. In contrast, fintech app users with less than £20,000 annual income in 2019 have contracted their expenditure the least, possibly reflecting a larger spending on essential goods and services.

Low-income users have, however, experienced the largest drop in earnings, around 15%, but their incomes have declined by much less. We provide evidence suggesting that an increase in government benefits since the start of the pandemic likely accounts for the significantly smaller fall in income (relative to earnings) for this group and, therefore, may also explain part of the less pronounced decline in their spending relative to the other groups. Consistent with this hypothesis, we show that when we restrict our analysis to users (i) at the bottom of the ex-ante income distribution as of 2019 and (ii) facing an increase in government benefits during 2020Q2, the level of their spending in June 2020 has been similar to the pre-pandemic levels. Nevertheless, our results suggest that users at the top of the income distribution may have seen a disproportionately larger increase in their personal savings.

As for the large swings in consumption during the first half of 2020, our evidence reveals that a significant part of the spending drop actually occurred before the introduction of lockdown measures and social distancing policies in the second half of March, suggesting that fears and uncertainty about contracting the virus and/or about future income have likely contributed to shape the dynamics of

household expenditure. In line with this view, we document that the spending increase in June, after the partial lifting of some of the restrictive measures, has been more pronounced in sectors characterized by a lower contact rate such as retail, food delivery and durable goods. In contrast, the spending on categories such as restaurants, travel & holidays and recreation outside home is still significantly below pre-pandemic levels.

Finally, we document pervasive heterogeneity in the breadth and depth of the spending decline across regions of the United Kingdom. London and the South East exhibit the largest contractions in the second quarter of 2020 whereas Northern Ireland and Wales have experienced a much milder spending reduction. We explore whether the number of coronavirus deaths per 1,000 inhabitants, the share of furloughed workers and the share of users within the top income bracket have any statistical power to predict a larger change in household expenditure. Our findings suggest that a 2% increase in the share of high-income users is associated with a very significant 1% of further decline in consumption. In contrast, there is little evidence of any systematic relationship between local spending changes and either the share of Covid-19 casualties or the share of furloughed workers across regions.

The analysis presented in this paper is not meant to substitute the fundamental role played by well-designed, representative surveys operating in many advanced and emerging economies, such as 'Understanding Society' and 'Living Costs and Food' in the United Kingdom. Rather, by exploiting the joint presence of spending, earnings and income at the individual level as well as the timeliness of their release, transaction data from linked bank accounts (such as for instance those provided by a fintech app) can complement more traditional survey methods. This paper has focused on tracking the heterogeneous effects of macroeconomic shocks on the distribution of household means and ends in real-time. This is of paramount importance for academic research seeking to identify the most salient channels of the transmission mechanism and for policy analysis trying to design policies targeted to specific groups of society.

References

- Abay, Kibrom A., Kibrom Tafere, and Andinet Woldemichael.** 2020. "Winners and Losers from COVID-19 : Global Evidence from Google Search." World Bank Policy Research Working Papers 9268.
- Andersen, Asger Lau, Emil Toft Hansen, Niels Johannesen, and Adam Sheridan.** 2020a. "Consumer responses to the COVID-19 crisis: Evidence from bank account transaction data." *Covid Economics*, 7: 88–114.
- Andersen, Asger Lau, Emil Toft Hansen, Niels Johannesen, and Adam Sheridan.** 2020b. "Pandemic, Shutdown and Consumer Spending: Lessons from Scandinavian Policy Responses to COVID-19."
- Attanasio, Orazio, Erik Hurst, and Luigi Pistaferri.** 2014. "The Evolution of Income, Consumption, and Leisure Inequality in the United States, 1980–2010." in *Improving the Measurement of Consumer Expenditures*, NBER(C. Carroll, T. Crossley, and J. Sabelhaus (Eds.)): 100–140.
- Baker, Scott R, R.A. Farrokhnia, Steffen Meyer, Michaela Pagel, and Constantine Yannelis.** 2020. "How Does Household Spending Respond to an Epidemic? Consumption During the 2020 COVID-19 Pandemic." National Bureau of Economic Research Working Paper 26949.
- Blundell, Richard, Monica Costa Dias, Robert Joyce, and Xiaowei Xu.** 2020. "COVID-19 and Inequalities*." *Fiscal Studies*, 1–29.
- Bounie, David, Youssouf Camara, and John W. Galbraith.** 2020. "Consumers' Mobility, Expenditure and Online-Offline Substitution Response to COVID-19: Evidence from French Transaction Data." CIRANO CIRANO Working Papers 2020s-28.
- Bourquin, Pascale, Isaac Delestre, Robert Joyce, Imran Rasul, and Tom Waters.** 2020. "The effects of coronavirus on household finances and financial distress." Institute for Fiscal Studies briefing note 14908.
- Carvalho, Vasco M., Juan R. Garcia, Stephen Hansen, Álvaro Ortiz, Tomasa Rodrigo, José V. Rodríguez Mora, and José Ruiz.** 2020. "Tracking the COVID-19 Crisis with High-Resolution Transaction Data." Cambridge University CWPE2030.
- Chetty, Raj, John Friedman, Nathaniel Hendren, Michael Stepner, and The Opportunity Insights Team.** 2020. "How Did COVID-19 and Stabilization Policies Affect Spending and Employment? A New Real-Time Economic Tracker Based on Private Sector Data." Harvard University mimeo.
- Chronopoulos, Dimitris K., Marcel Lukas, and John O.S. Wilson.** 2020. "Consumer Spending Responses to the COVID-19 Pandemic: An Assessment of Great Britain." St. Andrews University 20-012.
- Coibion, Olivier, Yuriy Gorodnichenko, and Michael Weber.** 2020. "The Cost of the Covid-19 Crisis: Lockdowns, Macroeconomic Expectations, and Consumer Spending." National Bureau of Economic Research Working Paper 27141.
- Cox, Natalie, Peter Ganong, Pascal Noel, Joseph Vavra, Arlene Wong, Diana Farrell,**

- and Fiona Greig.** 2020. "Initial Impacts of the Pandemic on Consumer Behavior: Evidence from Linked Income, Spending, and Savings Data." *Brookings Papers of Economic Activity*.
- Crawford, Rowena, Alex Davenport, Robert Joyce, and Peter Levell.** 2020. "Household spending and coronavirus." Institute for Fiscal Studies briefing note 14795.
- Delestre, Isaac, Robert Joyce, Imran Rasul, and Tom Waters.** 2020. "The effects of coronavirus on household finances and financial distress." Institute for Fiscal Studies briefing note BN 303.
- Dingel, Jonathan and Brent Neiman.** 2020. "How many jobs can be done at home?" *Covid Economics*, 1: 16–24.
- Hacioglu-Hoke, Sinem, Diego Känzig, and Paolo Surico.** 2020. "Consumption in the time of Covid-19: Evidence from UK transaction data." CEPR discussion paper 14733.
- Office for Budget Responsibility.** 2020. "Fiscal sustainability report: July." Office for Budget Responsibility.
- Pistaferri, Luigi.** 2015. "Household consumption: research questions, measurement issues, and data collection strategies." *Journal of Economic and Social Measurement*, 1-4: 123–149.

Table 1: Summary statistics

	Mean	<i>p</i> 10	<i>p</i> 25	<i>p</i> 50	<i>p</i> 75	<i>p</i> 90
<i>User Characteristics</i>						
#Banks	2.30	1.00	1.00	2.00	3.00	4.00
#Accounts	4.38	1.00	2.00	4.00	6.00	8.00
Age	37.90	26.00	30.00	36.00	44.00	53.00
Salary	2669.05	1080.92	1636.51	2327.32	3297.48	4374.32
<i>Monthly Transactions</i>						
#Transactions	100.20	44.00	64.00	92.00	126.00	164.00
Total expenditure	1475.87	458.21	731.96	1190.62	1880.62	2768.53
Nondurables	554.85	122.04	240.21	439.43	764.68	1139.19
Durables	140.19	6.57	16.97	46.95	124.96	317.02
Services	836.99	218.18	367.57	622.38	1029.21	1609.66
Groceries	311.08	40.83	107.31	236.14	443.05	683.18
Restaurant	110.63	13.19	32.75	73.20	141.20	234.84
<i>Monthly Mortgage and Rent Payments</i>						
Mortgage payments	935.52	295.07	483.85	749.11	1102.15	1592.76
Rents	762.60	90.95	300.00	576.00	945.00	1472.66

Note: Summary statistics on user characteristics (top panel), monthly transactions, salary and selected expense types (middle panel), mortgage and rent payments (bottom panel) for our main sample of users. The table reports the mean as well as 10, 25, 50, 75 and 90 percentiles. Salary, expenditure, and mortgage/rent payments refer to the period from January to March 2020.

Table 2: Percentage spending decline by groups of the ex-ante income distribution

Income group	Share of users	Share in expenditure	Contribution to decline
< 20K	23.8	17.5	9.8
20 – 30K	28.6	20.9	16.3
30 – 40K	23.0	26.9	28.6
> 40K	24.6	34.7	45.3

Note: The table reports the share of users by after-tax income groups of (i) below 20K, (ii) between 20 to 40K, (iii) from 30k to 40K, and (iv) above 40K, based on their 2019 income distribution. It also reports these income groups' share in aggregate total expenditure in the second quarter of 2019 (third column) and how much each group contributes to the decline in spending in the second quarter of 2020 relative to the same period of 2019 (fourth column). Total expenditure has been deflated by the U.K. CPI.

Table 3: Regional regressions


	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Covid-19 deaths	-0.0728* (0.0366)			-0.0718 (0.0372)	-0.0508 (0.0341)		-0.0498 (0.0347)
furloughed workers		0.535 (0.443)		0.506 (0.437)		0.530 (0.438)	0.510 (0.438)
higher-income users			-0.546*** (0.135)		-0.505*** (0.134)	-0.546*** (0.136)	-0.505*** (0.134)
Constant	-0.258*** (0.0315)	-0.472*** (0.136)	-0.186*** (0.0338)	-0.410** (0.139)	-0.157*** (0.0429)	-0.344* (0.138)	-0.310* (0.141)
Observations	108	108	108	108	108	108	108
Adjusted R^2	0.028	0.001	0.104	0.028	0.113	0.105	0.114

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$


Note: Dependent variable: percentage decline in regional spending over 2020Q2 relative to 2019Q2. Sources: the variable 'Covid-19 deaths' refers to the number of Covid-19 deaths per 1,000 inhabitants and is available from the ONS; 'furloughed workers' stands for the share of furloughed workers as reported by HMRC; 'higher-income users' is the share of users with after tax income above £40,000 in 2019 from the MDB sample. The 107 geographical areas are determined on the basis of the first two digit of MDB users' home post code. We exclude areas for which there are only 15 users or less.

Figure 1: Change in total expenditure and income



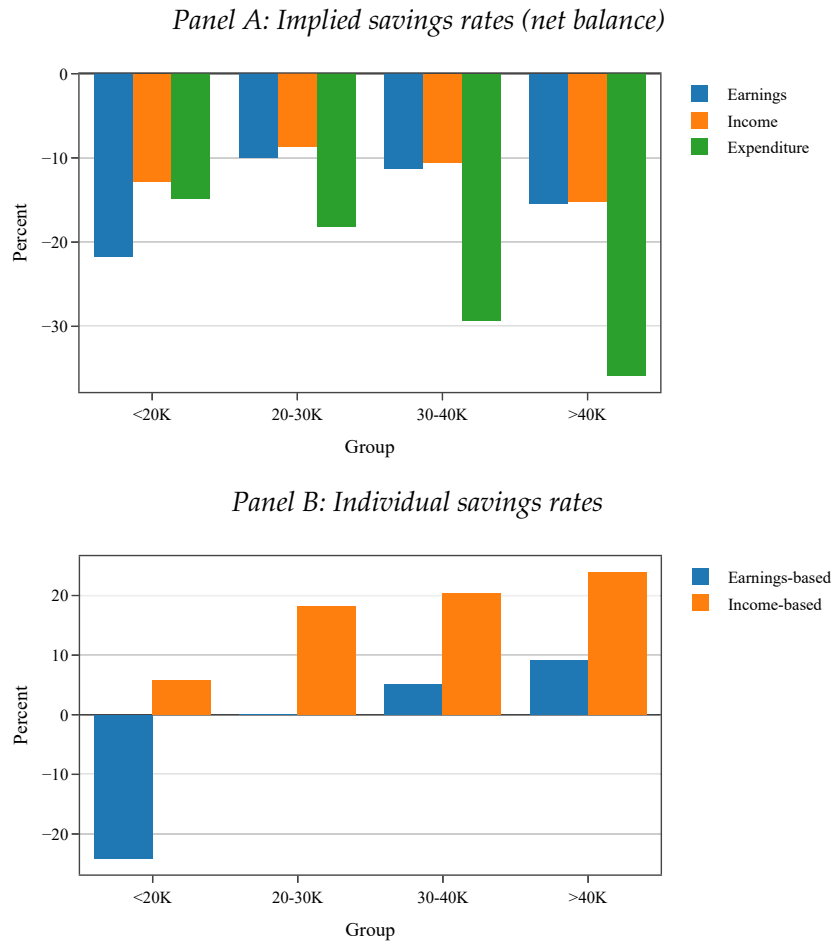
Notes: The figure records the year on year change in median income and total expenditure. To net out regular growth and age effects, we normalise all series to 100 in January in both years. Looking at year on year changes further takes care of the seasonality in the data.

Figure 2: Average spending, earnings and income by income groups




Notes: This figure shows the heterogeneity in spending behavior, earnings and income by four different income groups. Depicted is the year on year change in median expenditure, earnings and income series, normalized to 100 in January in both years. The users are classified into four income groups based on their annual after-tax income in 2019.

Figure 3: Personal saving rates along the ex-ante income distribution




Notes: This figure shows how personal savings were impacted by the pandemic. In the top panel, we show the percentage decline in earnings, income and expenditure in June 2020 relative to June 2019 for users with different after-tax income in 2019 that have been grouped according to the level displayed on the horizontal axis. This sheds light on the implied saving rates across these groups. In the bottom panel, we also provide more direct evidence, showing the savings rates as of June 2020 constructed from individual account balances ($s = 1 - C/Y$). For the income measure, we use both earnings as well as total income as measured by all incoming flows net of money transfers and refunds.

Figure 4: Monthly expenditure in restricted and unrestricted sectors



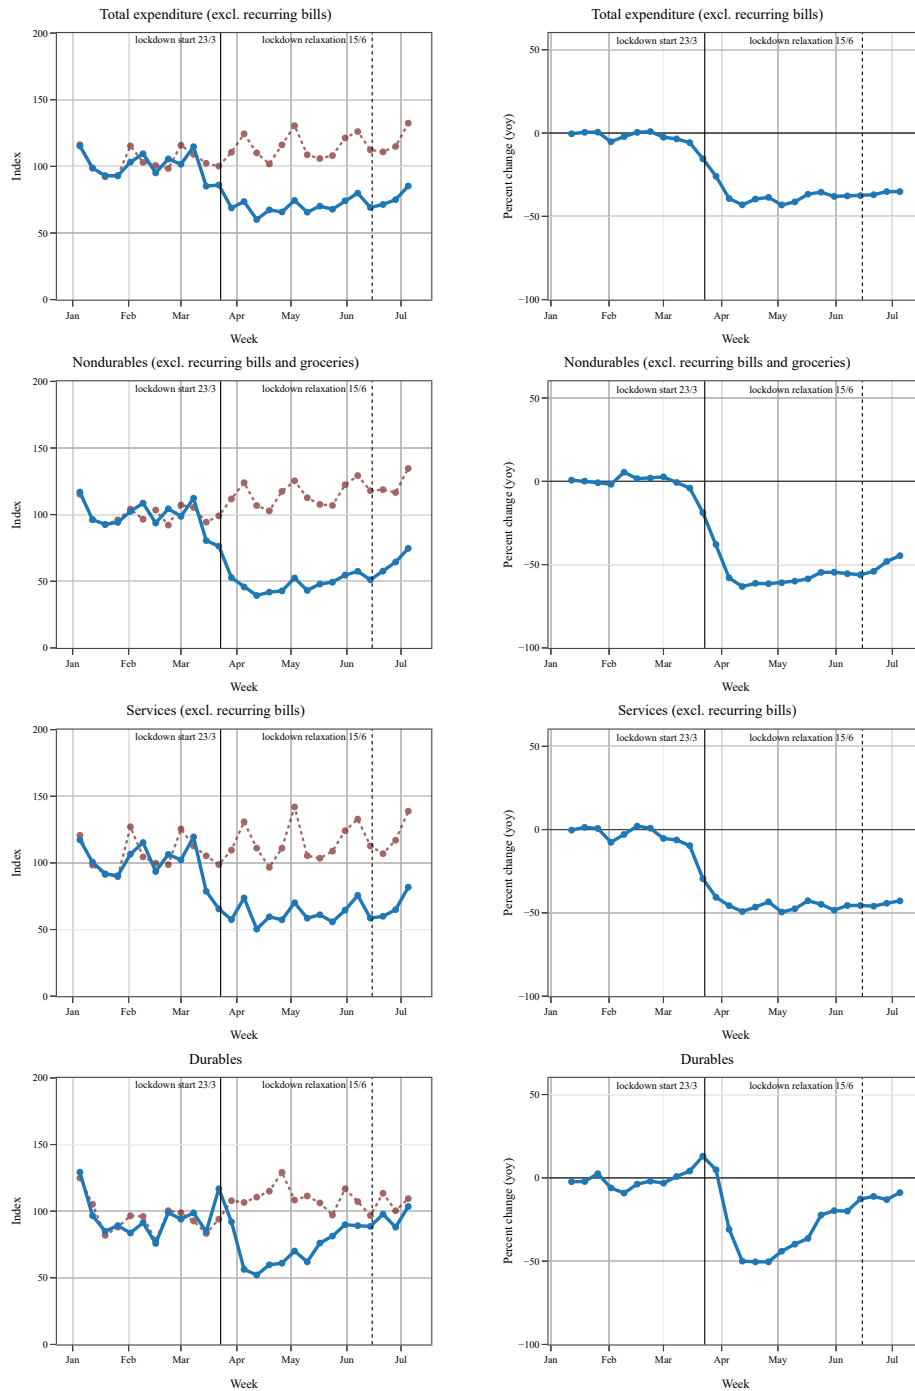
Notes: The figure shows the monthly median expenditure spent in restricted and unrestricted sectors over the course of the pandemic by the four income groups considered.

Figure 5: The role of government benefits




Notes: This figure shows the role of government benefits for the four different income groups. The top panel shows the share of users that experience an increase in benefits relative to January. The middle panel depicts the median value of the benefits, conditioning on the users receiving benefits. The bottom panel shows the change in spending for the two lower income groups by benefit status.

Figure 6: Average weekly expenditures by main categories



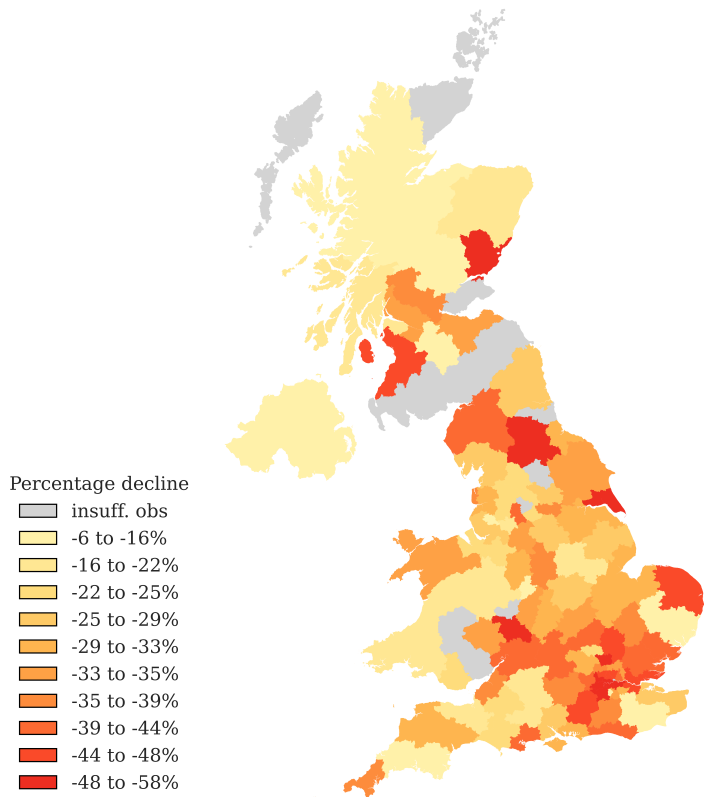

Notes: The figure shows the average weekly total expenditure, as well as spending in non-durable goods, services and durable goods normalised to 100 in January (left); year on year changes (right). All series have been deflated by the U.K. CPI. Normalising both series to 100 in January nets out regular growth and age effects.

Figure 7: Average weekly expenditures by sectors




Notes: Average weekly expenditure for a number of subcategories of interest expressed as year on year changes. Food delivery is composed of transactions in Deliveroo, Uber Eats and Just Eat. All series have been deflated by the U.K. CPI. Normalising the series to 100 in January in both years nets out regular growth and age effects.

Figure 8: The decline in total expenditure in the U.K



Notes: Percentage decline in spending over 2020Q2 relative to 2019Q2 across different regions in the U.K. based on the first two digits or MDB users' postcodes. Top panel: percentage decline in spending in the UK and its main regions. Bottom panel: heatmap of the percentage decline in spending; darker areas indicate larger drops in spending. We exclude postcodes for which we observe only 15 users or less.

Figure 9: The relationship between change in spending and share of deaths, furloughed workers and higher income users



Notes: The vertical axis in each chart refers to the percentage decline in spending over 2020Q2 relative to 2019Q2 across different regions in the U.K. based on the first two digits or MDB users' postcodes. The horizontal axis in each row stands for (i) the number of Covid-19 deaths per 1,000 inhabitants, (ii) the share of furloughed and (iii) the share of users with after tax income above £40,000 in 2019 from the MDB sample, respectively in each geographical area.

Appendix

A Data appendix

A.1 Data cleaning and user selection

A key advantage of data from financial aggregators such as MDB is that users can link all their accounts, thus generating a complete picture of their income and expenditure. However, to get an accurate picture of a user's finances, it is important to exclude users with incomplete accounts. To this end, we restrict our sample to users which

1. have at least one current account
2. transact at least 200 pounds in debits and have a minimum of 5 transactions in each month in our sample (Jan 2019 - June 2020)
3. have refreshed the app, and with that their accounts, in July 2020, ensuring that all relevant transactions have been downloaded by MDB

In the MDB data, there is also the possibility of duplicate accounts, for instance if a user registers with the service multiple times. MDB have developed a procedure to identify such accounts. We exclude accounts that MDB has identified as duplicates from our analysis.

There are also some users that use MDB for business-related purposes. Since we are interested in personal incomes and consumption expenditure, we want to exclude such accounts from our analysis. To this end, we exclude (1) accounts that are explicitly registered as a business account with the bank, (2) accounts that include the words "company", "business" or "limited" in the account name, (3) accounts that include a transaction to HMRC with the word "VAT" in the transaction description. Furthermore, we exclude users with disproportionately large expenditure tagged as business expenses, users with an implausibly large number of accounts (more than 10 active accounts) as well as users with monthly debits in excess of 100K.

Because our analysis depends on the transaction tags supplied by MDB, we impose some further restrictions on monthly expenditure and income calculated based on the tagged transactions. In particular, we require that users have monthly total expenditure of at least 100 pounds (and less than 10K) in all months but one in our sample. Furthermore, we require that they have some transactions tagged as income in at least 8 months in 2019 and 2 months in 2020 and that their annual after-tax income in 2019 is at least £5,000. Because the MDB sample is

tilted towards higher-income individuals and to further minimize the risk to include users who use MDB for business-related purposes, we also exclude users with an annual after-tax income in 2019 above £60,000. However, our results are robust to using higher thresholds. Finally, except when we look specifically at pensioners (in Appendix C), we restrict our sample to users in the working age population.

A.2 Construction of expenditure and income variables

To construct measures on expenditure and income, we rely on the transaction tags as well as the merchant descriptions provided by MDB. As discussed in the main text, MDB uses machine learning techniques to categorise transactions into different groups based on the description of a given transaction ('automatic tags'). The transaction description is of the same format as it would appear on bank statements and typically describes the source or nature of the transaction. The MDB algorithm is able to tag around 75% of all transactions. The remaining transactions are left untagged. Users have the ability to manually add tags or change the tags provided by MDB. A further 5% of transactions are manually tagged by users in this way ('user tags'). As part of the anonymisation of the data, sensitive parts of the transaction descriptions have been redacted. This makes it difficult to categorise transactions ourselves and because of this we are reliant on the tags provided by MDB.

As explained in the main text, we only rely on the automatic tags in our analysis, as they are consistently applied across users and time and not subject to recollection biases or the like. We also exclude all transactions not tagged by the algorithm from our analysis because it is hard to assess whether these transactions are actual expenditure/income or merely reflect transfers between different users or accounts. While this means that we may understate expenditure and income to a certain extent, it ensures that these measures do not erroneously include unrelated transactions such as transfers. Furthermore, constructing finer measures than total expenditure and income is only possible based on the tagged transactions. However, our results turn out to be robust with respect to the tag type as well as the treatment of untagged transactions, see Appendix D.

Expenditure measures. To construct measures of consumption expenditure, we closely follow the COICOP guidelines by the United Nations Statistics Division to group different transaction tags into broader categories of expenditure, such as total, non-durable, durables, and services expenditure, as well as smaller subcat-

egories such as retail, restaurants, or travel & holidays. The transaction tags used for non-durables, durables, and services are presented in Table A.1. Total expenditure consists of all tags in the non-durables, durables and services groups. In addition, Table A.2 reports all the transaction tags available in the data. Some of these spending types have been restricted during the lockdown period, e.g. concert or hotel expenses. Some spending categories have been unrestricted, e.g. food and groceries, or only partially restricted, e.g. when online shopping is an option for DIY. It is not possible to categorise all spending types under these three categories therefore some are left unclassified, e.g. benefits or loans.

For some subcategories of expenditure, there are no matching transaction tags provided by MDB, e.g. for online shopping or food delivery. For these categories, we use information on the merchants. In particular, we proxy online shopping by transactions with *Amazon* and *Ebay* as the merchant. For food delivery, we look at transactions with *Uber Eats*, *Deliveroo*, and *Just Eat* as the merchant.

For each user we then construct measures of expenditure for these categories by summing all debit transactions with the relevant tags over a specific time interval (e.g. a week or a month). Spending at the weekly frequency features some bunching towards the beginning and the end of a given month because of recurring bills such as utilities, broadband, mobile or council taxes. For this reason, we also construct measures of expenditure excluding these recurring bills, which makes the series more easily interpretable.

Income, earnings, and benefits. On the income side, we construct measures of earnings, benefits and total income. For earnings, we rely on the tag *Salary or Wages (Main)*. There are also tags for secondary salaries in the data, *Salary (secondary)*, however, these are poorly populated and including them does not change any of our results.¹

To construct a measure of total income, we construct different subgroups of income. In particular, we create a measure of financial and other income (*Irregular Income or Gifts, Interest income, Rental income (whole property), Dividend*), a measure of pensions (*Pension, Work pension*), and a measure of overall benefits (*Benefits, Family benefits, Other benefits*). Total income is then equal to the sum of earnings as well as all other income.

For our measure of government benefits, we use the *Benefits* tag. This tag includes income-support, work, and some health related benefits, issued by the

¹As discussed in Bourquin et al. (2020), in earlier vintages of the MDB data there was a problem with the algorithm to identify salaries, which resulted in too few transactions being tagged as salaries from March 2020 onwards. In the new vintage of the data, MDB have addressed this issue and the salary tags should be reliable again.

U.K. Department for Work and Pensions. We do not include family and other benefits. However, we have verified that unlike income related benefits, these benefits were very stable over our sample of interest, both along the intensive and extensive margin.

Table A.1: Transaction categories for nondurable, durable and services consumption (automatic tags only)

	Services	Nondurables	Durables
1	Bank charges	Insurance - other	Appliances or Electrical
2	Banking Charges	Life insurance	Art
3	Beauty treatments	Lunch or Snacks	Cycling
4	Breakdown cover	Media bundle	DIY
5	Broadband	Medical treatment	Furniture
6	Child - Everyday or Childcare	Memberships	Furniture, Furnishing, Gardens
7	Childcare Fees	Mobile	Garden
8	Children - other	Mobile phone insurance	Home
9	Cinema	Museum/exhibition	Home DIY or Repairs
10	Concert & Theatre	Music	Jewellery
11	Contents or Other Insurance	Parking	Lighting
12	Council Tax	Parking or Tolls	Musical Equipment
13	Course and Tuition Fees	Payment Protection Insurance	Personal Electronics
14	Dental insurance	Pet Insurance	Software
15	Dental treatment	Phone (landline)	
16	Device rental	Phone or Internet	
17	Dining and drinking	Physiotherapy	
18	Dining or Going Out	Postage / Shipping	
19	Driving Lessons	Public Transport	
20	Dry cleaning and laundry	Road charges	
21	Education - other	School fees	
22	Enjoyment	Service / Parts / Repairs	
23	Entertainment, TV, Media	Social club	
24	Flights	Spa	
25	Gambling	Sports Club Membership	
26	Gym Membership	Sports event	
27	Hairdressing	Take-away	
28	Hairdressing, Health, Other	Taxi	
29	Health insurance	Taxis or Vehicle Hire	
30	Hire Purchase	Tradesmen fees	
31	Hobbies - other	Transport	
32	Hobbies or Activities	TV / Movies Package	
33	Hobby Club Membership	TV Licence	
34	Holiday	Vehicle hire	
35	Holidays	Vehicle insurance	
36	Home appliance insurance	Vehicle Running Costs	
37	Home insurance	Vet	
38	Hotel/B&B	Web hosting	
39	Income insurance	Zoo/theme park	
40	Insurance		

Table A.2: List of transaction types based on automatic tags and their classification based on restricted, partially restricted and unrestricted transactions during lockdown

	Transaction Type	Unrestricted	Restricted	Partially restricted	Transaction Type	Unrestricted	Restricted	Partially restricted
1	Accessories				41	Dental treatment	•	
2	Administration - other	•		•	42	Designer clothes		•
3	Alcohol	•			43	Device rental	•	
4	Appearance			•	44	Dining and drinking	•	
5	Appliances or Electrical			•	45	Dining or Going Out	•	
6	Art			•	46	Dividend		
7	Art Supplies			•	47	Donation to organisation		
8	Bank charges				48	Driving Lessons	•	
9	Banking Charges				49	Dry cleaning and laundry	•	
10	Beauty products			•	50	Education - other		•
11	Beauty treatments				51	Electrical equipment		•
12	Benefits		•		52	Electricity	•	
13	Bills				53	Energy (Gas, Elec, Other)	•	
14	Books / Magazines / Newspapers	•		•	54	Enjoyment	•	
15	Breakdown cover	•			55	Entertainment, TV, Media		
16	Broadband	•			56	Expenses	•	
17	Bursary				57	Eye care		
18	Business Expenses				58	Family benefits		
19	Car fund				59	Financial - other		
20	Cash				60	Fines		
21	Child - Clothes			•	61	Flights		•
22	Child - Everyday or Childcare		•		62	Flowers		•
23	Child - Toys, Clubs or Other			•	63	Food, Groceries, Household	•	
24	Childcare Fees		•		64	Fuel	•	
25	Children - other				65	Furniture		•
26	Cinema		•		66	Furniture, Furnishing, Gardens		•
27	Clothes			•	67	Gambling		•
28	Clothes - Designer or Other			•	68	Games and gaming	•	
29	Clothes - Everyday or Work			•	69	Garden		•
30	Coal/Oil/LPG/other	•			70	Gas	•	
31	Concert & Theatre				71	General savings		
32	Contents or Other Insurance	•	•		72	Gifts or Presents		•
33	Council Tax				73	Gym Membership	•	
34	Course and Tuition Fees			•	74	Hairdressing	•	
35	Credit Card				75	Hairdressing, Health, Other	•	
36	Credit card repayment				76	Health insurance	•	
37	Current account				77	Hire Purchase		•
38	Cycling	•			78	Hobbies - other		•
39	DIY			•	79	Hobbies or Activities		•
40	Dental insurance	•			80	Hobby Club Membership	•	

Table A.2: List of transaction types based on automatic tags and their classification based on restricted, partially restricted and unrestricted transactions during lockdown (continued)

	Transaction Type	Unrestricted	Restricted	Partially restricted	Transaction Type	Unrestricted	Restricted	Partially restricted
81	Hobby Supplies				121	Parking or Tolls		
82	Holiday		•	•	122	Payday Loan		•
83	Holidays		•		123	Payday loan funds		
84	Home			•	124	Payment Protection Insurance	•	
85	Home DIY or Repairs			•	125	Paypal account		
86	Home appliance insurance	•			126	Penalty charges		
87	Home insurance	•			127	Pension		
88	Hotel/B&B				128	Pension or Investments		
89	ISA		•		129	Personal Care - Other		•
90	Income insurance				130	Personal Electronics		•
91	Insurance	•			131	Personal Loan		
92	Insurance - other	•			132	Pet - Everyday or Food	•	
93	Interest charges				133	Pet - Toys, Training, Other		
94	Interest income				134	Pet Insurance	•	
95	Investment - other				135	Phone (landline)	•	
96	Investments or Shares				136	Phone or Internet	•	
97	Irregular Income or Gifts				137	Photography		
98	Jewellery			•	138	Physiotherapy		•
99	Legal	•			139	Postage / Shipping	•	
100	Life insurance	•			140	Printing		
101	Lighting			•	141	Public Transport		•
102	Lunch or Snacks			•	142	Refunded purchase		
103	Media bundle				143	Rent		
104	Medical treatment	•			144	Rental income (room)		
105	Medical, Dental, Eye Care		•		145	Rental income (whole property)		
106	Memberships			•	146	Repayments		
107	Mobile	•			147	Rewards/cashback		
108	Mobile app	•			148	Road charges		
109	Mobile phone insurance	•			149	Salary (secondary)	•	
110	Mortgage or Rent				150	Salary or Wages (Main)		
111	Mortgage payment				151	Saving (general)		
112	Mortgage release				152	Savings		
113	Museum/exhibition				153	Savings (general)		
114	Music				154	School fees		
115	Musical Equipment	•			155	Secured loan repayment	•	
116	No Tag			•	156	Service / Parts / Repairs	•	
117	One-off or Other				157	Sharedealing account		
118	One-off or Other Payment				158	Shoes		•
119	Other benefits				159	Social club	•	
120	Parking		•		160	Software	•	

Table A.2: List of transaction types based on automatic tags and their classification based on restricted, partially restricted and unrestricted transactions during lockdown (continued)

	Transaction Type	Unrestricted	Restricted	Partially restricted	Transaction Type	Unrestricted	Restricted	Partially restricted
161	Spa		•		178	Toys		
162	Sports Club Membership		•		179	Tradesmen fees	•	
163	Sports Equipment		•	•	180	Transfers		
164	Sports event		•		181	Transport		•
165	Stationery			•	182	Unsecured loan funds		
166	Stationery & consumables			•	183	Unsecured loan repayment		
167	Store card repayment				184	Vehicle	•	
168	Student Loan funds				185	Vehicle Running Costs	•	
169	Student loan repayment				186	Vehicle hire		•
170	Supermarket	•			187	Vehicle insurance	•	
171	TV / Movies Package	•			188	Vet	•	
172	TV Licence	•			189	Water	•	
173	Take-away	•			190	Web hosting	•	
174	Tax Payment				191	Winnings		
175	Taxi		•		192	Work pension		
176	Taxis or Vehicle Hire		•		193	Zoo/theme park		•
177	Toiletries	•						

A.3 Comparison of different spending indicators in the U.K.

There are two main information sources which inform both scholars and policy makers about how household expenditure evolves in the U.K.. The first is the Office of National Statistics' Living Costs and Food Survey (LCFS). This is a yearly survey whose results are released twelve months after the previous financial year ends.² The LCFS aims to understand consumption at the household level and reports the average weekly expenditure alongside comparisons with the previous financial years. The survey results feed into the calculation of consumer price indices and national output.

The second data source is the Understanding Society Surveys (USS), formerly conducted as the British Household Panel Survey.³ This survey is conducted in multiple waves. As data collection and processing takes place for one wave, the other waves are taken to the development stage. The latest wave, Wave 9, has started in 2016 with the development phase. The data collection took place until the third quarter of 2019. Final data processing and documentation has been finalised in the second half of 2019 which allowed the results to be released in the last quarter of 2019.

In Table A.3, we compare our spending indicator with the information collected by the surveys by the ONS and Understanding Society. The first row discusses the release delay. To illustrate this, consider spending data that occurred during the month of March 2019. This information is collected by the MDB automatically and with no delay via the app and thus can be processed and released a few days after 31 March 2019, say – for sake of concreteness – by 3 April 2019. In contrast, the LCFS and USS releases come only after eleven and six-to-twelve months: for instance, data about March 2019 became publicly available on 19 March 2020 and in the month of December 2019 respectively.

Another relevant dimension of comparison is the frequency of data release. The MDB data allows us to observe every transaction that takes place on a daily basis and so we are able to generate weekly (and even daily) spending indicators. On the other hand, the data release of the ONS' Living Cost and Food Survey is yearly while the one of Understanding Society occurs at quarterly frequency.

As for method of data collection, it is worth emphasizing that the fintech app minimizes measurement errors as both the exact amount, and the exact time and

²The latest release for the financial year of April 2018 to March 2019 is in <https://www.ons.gov.uk/peoplepopulationandcommunity/personalandhouseholdfinances/expenditure/bulletins/familyspendingintheuk/april2018tomarch2019>.

³Information on the methodology and the coverage of the Understanding Society Surveys can be found in <https://www.understandingsociety.ac.uk>.

date of transaction are electronically recorded in real-time via the app. In contrast, the LCFS is based on a mixture of diary records (on smaller and more frequent expenses, especially groceries) and a retrospective interview (for larger and less frequent purchases) whereas the USS uses only a retrospective interview. As discussed by [Attanasio, Hurst, and Pistaferri \(2014\)](#) and [Pistaferri \(2015\)](#), retrospective interviews are subject to recollection biases and diary records typically suffer from severe non-classical measurement errors that may seriously compromise inference.

In terms of sample size, our MDB sample covers more than 8,000 individuals, which will likely increase further over time as the user base grows and existing users refresh their accounts, while ONS and Understanding Society cover up to 6,000 households and 40,000 individuals, respectively. However, the MDB data may not be as representative of the British population as the other sources, as the users self-select in using the app. Therefore, we will analyze the its representativeness in detail in [Appendix B](#).

It is useful at this point to discuss the means of payment. In the MDB data, we observe separately the amount of electronic payments and ATM cash withdrawals, though for the latter we do not know which category this is spent on. On the other hand, the LCFS and the USS include both electronic and cash payments. However, as discussed in [Section 2](#), the share of cash transactions is less than a quarter of the total transactions and has been declining over time. Furthermore, the MDB data allows us a detailed coverage of different spending types. This is similar to the LCFS but different from the USS, which only covers some expenses types such as groceries, and alcohol and tobacco (but not other relevant categories such as transport, restaurant or travel).

All methods include information to explore regional heterogeneities as well as those arising from housing tenure, income and age. Finally, the unit of observations in the LCFS (and the USS) is a household In contrast, the unit of observations in the MDB is a user, who through the app may be linking either individual or joint accounts, with the associated credit and debit cards.

In sum, all three methods and sources have advantages and disadvantages. The most appealing feature of our method is timeliness. As such, it may represent a useful addition that complements the LCFS and the USS, which are based on more representative samples of the British population.

Table A.3: Comparison of different consumption indicators

	Real-time Consumption Indicator	Office for National Statistics	Understanding Society
Release delay <i>e.g. data for March 2019</i>	Money Dashboard 3 days 3 Apr 2019	Living Costs and Food Survey 11 months after the financial year ends 19 Mar 2020	The UK Household Longitudinal Study 6 months after data collection ends for the past 2 years Dec 2019
Frequency of data release	Weekly (Daily)	Yearly	Quarterly
Data collection method	Automated via the app	Diary and interview	Interview
Transactions type	Electronic payments & ATM cash withdrawals	Electronic & cash payments	Electronic & cash payments
Sample size	34,601 individuals	1,500 to 6,000 households	34,959 (Wave 9) - 40,000 (Wave 1) individuals
Representativeness	Partial	Full	Full
Spending categories	<ul style="list-style-type: none"> • Nondurable • Durable • Services • Retail • Groceries • Alcohol and Tobacco • Home • Online shopping • Transport and Fuel • Restaurant • Travel • Recreation and Culture 	<ul style="list-style-type: none"> • • • • • • • • • • • • 	<ul style="list-style-type: none"> • • • • • • • • • • • •
Heterogeneity	<ul style="list-style-type: none"> • Geographical • Housing Tenure • Income • Age 	<ul style="list-style-type: none"> • • • • 	<ul style="list-style-type: none"> • • • •


Notes: The table compares different sources of inferring consumption in the U.K.: real-time consumption indicator presented in this paper using Money Dashboard transaction level data; Office for National Statistics' Living Costs and Food Survey; Understanding Society UK Household Longitudinal Study.

B Descriptive statistics and representativeness

B.1 Demographic and regional distribution

In this section, we provide some descriptive features of the data. The age profile of our sample is presented in Figure B.1. Figures B.2 and B.3 present the number of users and the share of users with respect to the total population of the regions they live in. Finally, Figure B.4 shows the gender profile of our sample.⁴ We can see that MDB users are younger and more likely to be based in the Greater London area than the British population. Furthermore, men appear to be over-represented in our sample, even though this may partly reflect the fact that shared accounts are more likely attributed to men.


Figure B.1: Age profile of users



Note: Age profile of the users in the final dataset.


⁴It is important to note that gender is a derived variable, provided by MDB, and not by the user itself.

Figure B.2: Number of users by regions




Note: Total number of users with respect to the regions they live in. NA: users with no postcode information.

Figure B.3: Number of users by regions normalized by region population



Note: Share of total number of users to the total population of the regions they live in. Excluding those with no postcode information.

Figure B.4: Share of users by gender



B.2 Representativeness

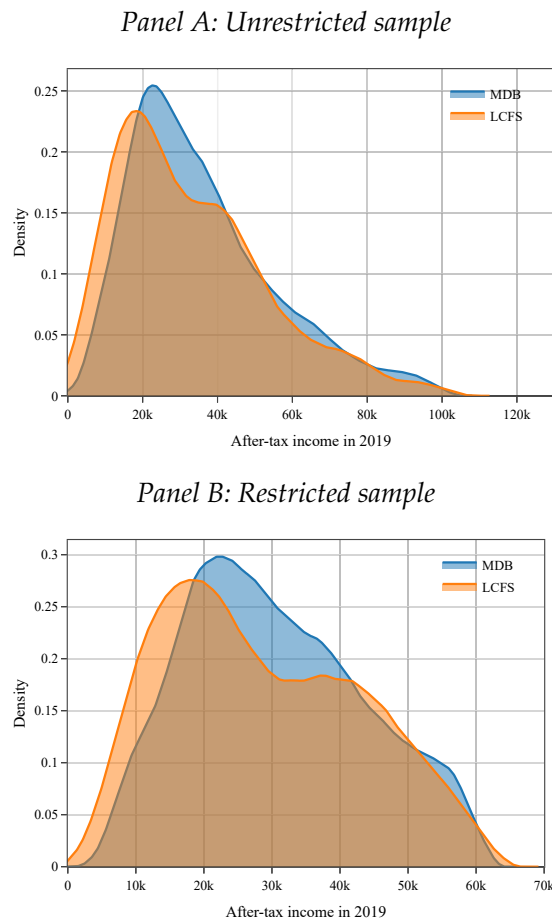
As discussed in more detail in Appendix [A.3](#), there exist some important differences between the MDB data and other sources and national statistics in particular. The MDB has many desirable properties especially concerning measurement and timeliness. However, this comes at the cost of a potentially less representative sample, as the users self-select in using the app and are not selected to ensure representativeness. In fact, as we have seen above there are some important demographic and regional differences between the MDB sample and the British population. Therefore, it is crucial to analyze whether our sample may still bear some representativeness for consumption and income of the British population.

In this appendix, we thus investigate the representativeness of our MDB sample along two dimensions. First, we study how the pre-pandemic income and consumption distribution compares to other, representative sources. Second, we study how aggregated expenditure from the MDB data compares over time with the official data from the ONS on final household expenditure.

Cross-sectional dimension. We start by analyzing the representativeness along the cross-sectional dimension. Our benchmark will be the 2019 data from the 2018/2019 wave of the LCFS survey. The derived households characteristics data contains all the data we need, i.e. total expenditure and disposable income as well as more detailed expenditure on COICOP categories. We aggregate the normal weekly income and consumption measures to annual measures by multiplying them by 52.

First, we look at how the pre-pandemic income distribution in the MDB data compares to the distribution in the LCFS. Figure [B.5](#) presents the results. In the top panel we study the unrestricted income distribution while the bottom panel shows the distribution after we impose our additional restrictions (we impose the same restrictions in both the MDB and LCFS data). We can see that the income distribution in the MDB data is fairly similar to the LCFS, both in the unrestricted sample as well as in our main sample where we exclude very high income individuals. The MDB distribution is slightly shifted to the right relative to the LCFS sample, however, the match is very close.

Figure B.5: Income distribution in the MDB and LCFS data




Notes: The figure shows distribution of the annual after-tax income in 2019 in the MDB and LCFS sample. As the LCFS is currently only available for the first quarter 2019, annual income is computed based on the normal disposable income of the sample in that quarter. The pdfs' are estimated using an Epanechnikov kernel. The top panel shows the unrestricted income distribution in the MDB and the LCFS data, only excluding incomes over 100k that are top-coded in LCFS data. The bottom panel shows the income distribution in the MDB and LCFS data once we impose our additional restrictions (i.e. dropping very low and high incomes). As explained in Appendix A, we drop very high incomes to further minimize the risk to include users who use MDB for business-related purposes (users with an annual after-tax income in 2019 above 60k). However, our results are robust to using higher thresholds.

In Figure B.6, we compare the distribution of COICOP total expenditures (including housing) in the MDB and the LCFS data.⁵ Reassuringly, the two distributions are virtually identical.

⁵Note that this definition does not include mortgage interest payments and council tax, which are included in LCFS under 'other expenditure'.

Figure B.6: Consumption distribution in the MDB and LCFS data




Notes: The figure shows distribution of the annual total consumption expenditures including housing (COICOP definition) in 2019 in the restricted MDB and LCFS samples. As the LCFS is only available for the first quarter of 2019, annual expenditure is computed based on the COICOP total consumption expenditure measure of the sample in that quarter. The pdfs are estimated using an Epanechnikov kernel.

Finally, we also study how spending across the main COICOP categories in the MDB data matches up with the LCFS. For this purpose, we have constructed 12 expenditure categories ranging from food and non-alcoholic drinks to miscellaneous goods and services, trying to match the definitions of the corresponding COICOP categories in the LCFS data as closely as possible.

The results are shown in Figure B.7, which shows the annual average expenditure in 2019 on the main COICOP categories as a share of total expenditure (including housing). We can see that the shares in the MDB data match the LCFS data fairly well. We observe the largest differences for 'food and non-alcoholic drinks', 'alcoholic drinks, tobacco and narcotics' as well as 'recreation and culture'. However, note that the total share on food and alcohol is comparable across the MDB and LCFS data. Given that the tags in the MDB data may be sometimes not specific enough to identify purchases of alcohol (if a user buys alcohol at a supermarket in combination with other foods, this will likely show up under broader tags than 'alcohol'/'tobacco', e.g. 'supermarket' or 'foods, groceries, household', which are included in the food category). Similarly, the high share on 'recreation and culture' in the MDB data is driven by the tag 'Enjoyment', which could potentially also contain expenditure on other COICOP categories, which we are however not able to disentangle in the MDB data. By way of summary, it seems that the expenditure on different COICOP categories in the MDB data is fairly consistent with the LCFS data.

Figure B.7: Spending on main COICOP categories in the MDB and LCFS data



Notes: The figure shows annual average expenditure on the main COICOP categories as a share of total expenditure including housing in the restricted MDB and LCFS samples. As the LCFS is only available for the 2019Q1, annual expenditure is computed based on the COICOP expenditure measures of the sample in that quarter.

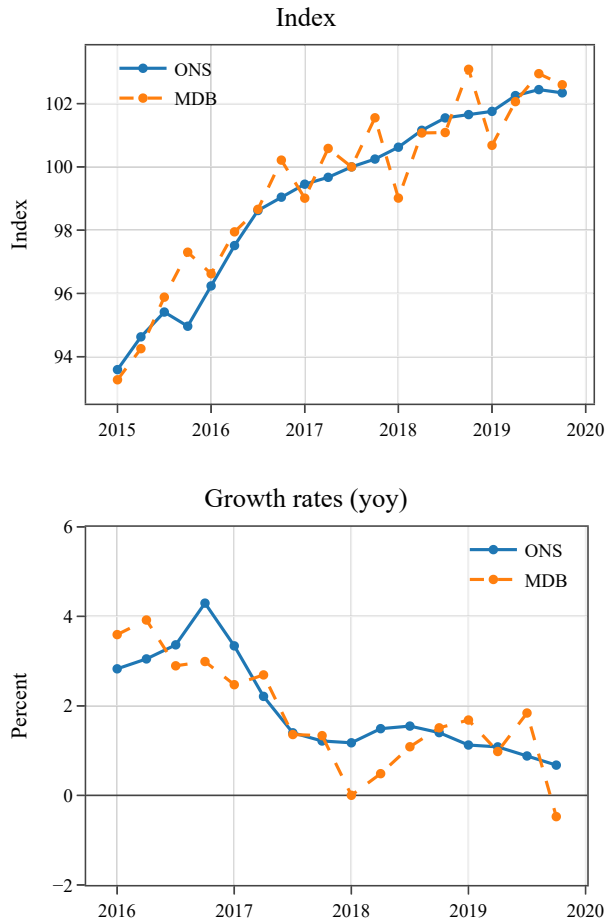
Time-series dimension. We also study how the MDB compares with national statistics in the time-series dimension. To this end, we construct an aggregate consumption expenditure measure based on the MDB data over the period 2015-2019. In this context, it is problematic to follow the same set of users, as we do for our main results, because of the following reasons. First, as the MDB users age and gain more experience, this will naturally be reflected in their expenditure and income. This is very different from national statistics, which reflect such changes only to the extent the overall demography changes. Second, it constrains the user pool considerably, as the number of users has grown over time and only a smaller share of users have used the app consistently from the start.

To address these issues, we follow an alternative approach and use a rotating panel of users. In particular, we focus on users that we consistently observe for at least three consecutive years. This leaves us with a rotating panel of around 45,000 users. For these users, we compute monthly total consumption expenditure and rents. After removing outliers, we compute monthly averages of total

⁶To remove outliers in total expenditure, we drop the top and bottom 5% of observations, first within every month to remove outliers in a given month and then also across the entire sample to remove global outliers. For rents, we remove observations below 200 and above 6,000 GBP.

expenditure and rents.⁶ Using this data, we construct a measure of total expenditure including imputed rents, by adding the average monthly rent to the average total expenditure. This simple approach of imputing rents essentially amounts to impute to all homeowner occupiers the same average rent paid by the renters in our sample. Finally, we deflate the nominal expenditure series by the U.K. CPI and aggregate the monthly deflated series to the quarterly frequency.⁷

Figure B.8: Comparison to national statistics



Notes: The figure compares quarterly measures of aggregate household consumption expenditure from the MDB data (non-SA) and the official statistics from the ONS (SA). The top panel shows an index of the series, normalized to 100 in mid-2017. The bottom panel shows year-over-year growth rates for the period 2016-2019.

Figure B.8 compares the MDB aggregate to the measure of real final consump-

⁷Importantly, the evolution of aggregate expenditure based on this representativeness sample is very similar to aggregate expenditure based on our main sample in 2019, the year the two samples overlap, when accounting for the aforementioned age and experience effects.

tion expenditure per capita from the ONS. Because the MDB data does not include cash payments, the level of our series will be somewhat lower. Therefore, normalize both series to an index, choosing mid-2017 as the base period. We can clearly see that the two series follow a very similar trend. We can also see that the MDB data features a strong seasonal pattern. However, given the relatively short sample period, it is challenging to remove the seasonality from the series. In the bottom panel, we thus show the yoy growth rates for 2016-2019. Even though the MDB growth rates are somewhat more erratic, they track the growth rates of the ONS data reasonably well. The correlation between the two growth rates stands as high as 79%.

The cross-sectional evidence in combination with the time-series evidence suggests that the MDB sample bears quite some resemblance with the UK population in terms of income and consumption – alleviating concerns regarding the representativeness of the MDB data. There is also external evidence on the representativeness of the MDB data, in particular by [Bourquin et al. \(2020\)](#) who investigate the representativeness of the data along a battery of measures, including demographics, income, earnings, the number of salaries in the household, the number of bank accounts and credit cards, and pre-crisis indicators of financial distress.

C Additional analysis and charts

In this Appendix, we present some additional results on the heterogeneous spending response to the pandemic. We begin by grouping MDB users according to their housing tenure, age gender.⁸ The results are presented in Figure C.1.

Looking at the role of housing tenure status, we can see that mortgagors (in orange) and renters (in green) have cut their expenditure by more than owners (in blue). The recovery in June is least pronounced for mortgagors, which could possibly reflect some form of financial distress. Turning to age, we find a larger drop in spending for younger people (green line), closely followed by middle aged users (orange line). For older people in our sample (blue line), we record a smaller decline in consumption. Finally, women in our sample cut expenditure somewhat less than men.

In Figure C.2, we look at how expenditure varies depending on whether a user receives benefits, a pension or has children.⁹ Consistent with our results in the main text for users for whom we observe an increase in benefits, we find that benefits recipients (in orange) cut their expenditure by about 10% less than non-recipients (in blue).¹⁰ Looking at pensioners (orange line), we find that the fall in expenditure has been comparable in April and May to other users (in blue) but the recovery for this group in June appears more muted. This stands in contrast to the results we have for the older users in the working age population, for which we observe a smaller response than for any other age group. Finally, the bottom panel of Figure C.2 reveals that users with children (in orange) cut their expenditure by less than users without (in blue), possibly reflecting the government support through child-related benefits and tax credit.

⁸We identify mortgagors and renters based on whether we observe rent or mortgage payments for a given user. Users for which we do not observe either are classified as owners.

⁹Pensioners are identified based on pension payments and the age of a given user. Note that we exclude pensioners from our main analysis but report the results for them here separately. Finally, we identify users with children based on whether they receive family benefits and/or have significant child related expenses.

¹⁰We also tried to identify “likely unemployed” adults by conditioning on users who receive government benefits in March to June 2020. The results are very similar to the ones for the users for which we observe an increase in benefits over this period.

Figure C.1: Spending by housing tenure, age, and gender



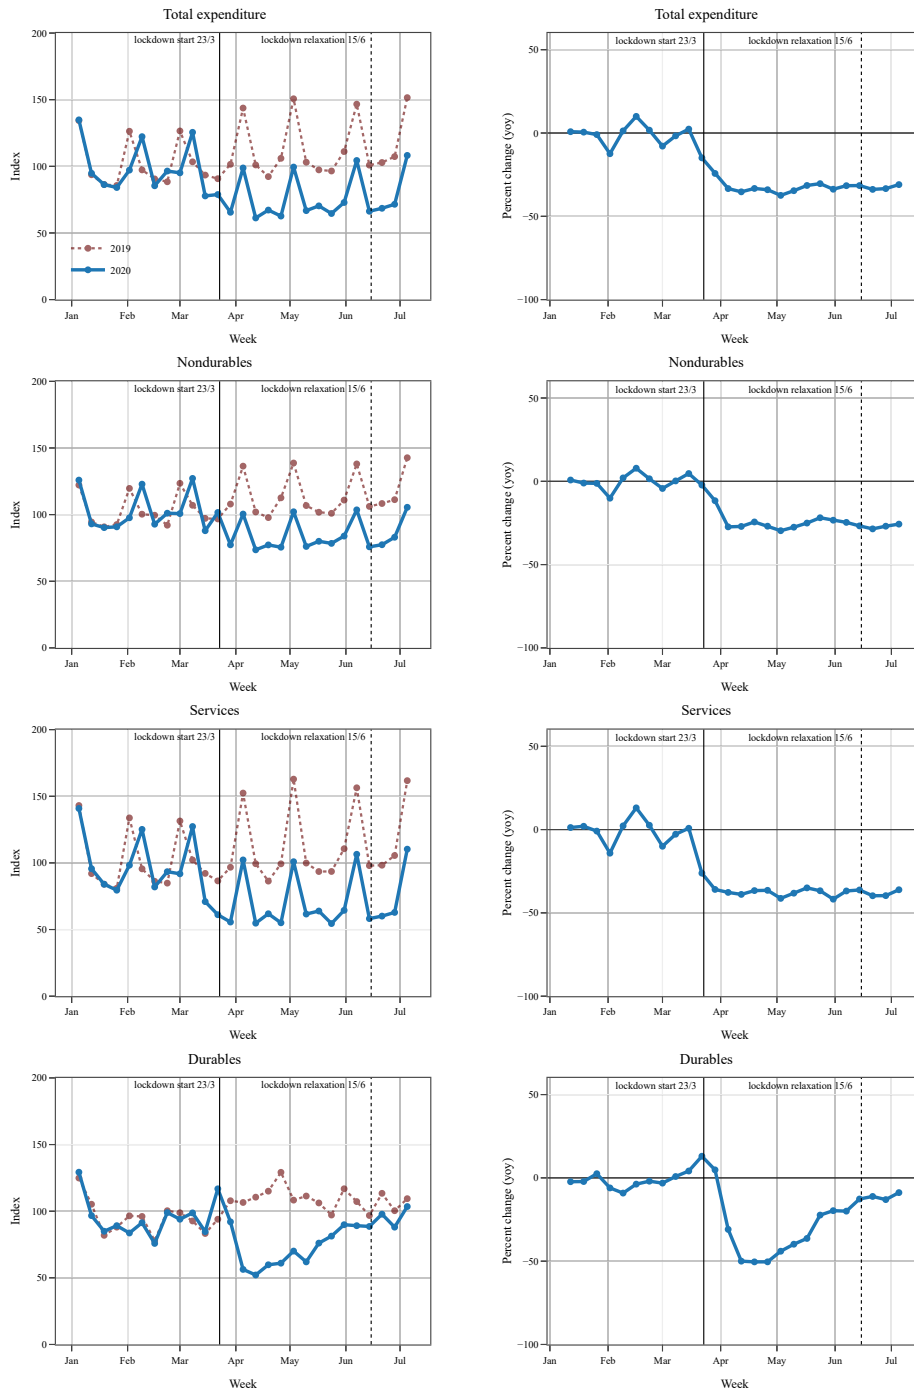


Figure C.2: Spending for benefit recipients, pensioners, and users with children



For completeness, in Figure C.3, we present the average weekly expenditure measures of Section 4 but including recurring bills. The results are qualitatively very similar to the ones excluding recurring bills in the main text. However, by abstracting from recurring bills and focussing on discretionary spending, the results are visually easier to interpret, especially for the indices in the left column.

Figure C.3: Average weekly expenditures by main categories




Notes: The figure shows the average weekly total expenditure, as well as spending in non-durable goods, services and durable goods normalised to 100 in January (left); year on year changes (right). All series have been deflated by the U.K. CPI. Normalising both series to 100 in January nets out regular growth and age effects.

D Robustness

In our main analysis, we have exclusively relied on the automatic tags provided by MDB. To mitigate concerns that the patterns in expenditure and income that we uncover are driven by the performance of the tagging algorithm, we report here the share of outgoing and incoming transactions that are tagged over time. Figure D.1 presents the results. From Panel A, we can see that the algorithm is generally more successful in tagging expenditure than income, with around 80% of the outgoing transactions tagged and 55% of the incoming transactions. Importantly, the share of tagged transactions turns out to be remarkably stable across time. From Panel B, we can see that this does not only hold in the aggregate but also by the four income groups we consider.

Figure D.1: Share of untagged transactions in 2020 (percent)




Notes: The figure shows the share of untagged transactions over the course of 2020, in the aggregate (Panel A) and by income groups (Panel B). The shares of untagged outgoing transactions is on the left and the share of incoming transactions is on the right.

Given that a non-negligible share of transactions cannot be tagged by the algorithm, we have also performed some robustness checks regarding the tag type and the treatment of untagged transactions.

In Figure D.2, we analyze how our main results are affected if we rely on the user precedence tag as opposed to the automatic tags. Interestingly, the results turn out to be quite consistent to using the automatic tags.

Figure D.2: Robustness with respect to tag type

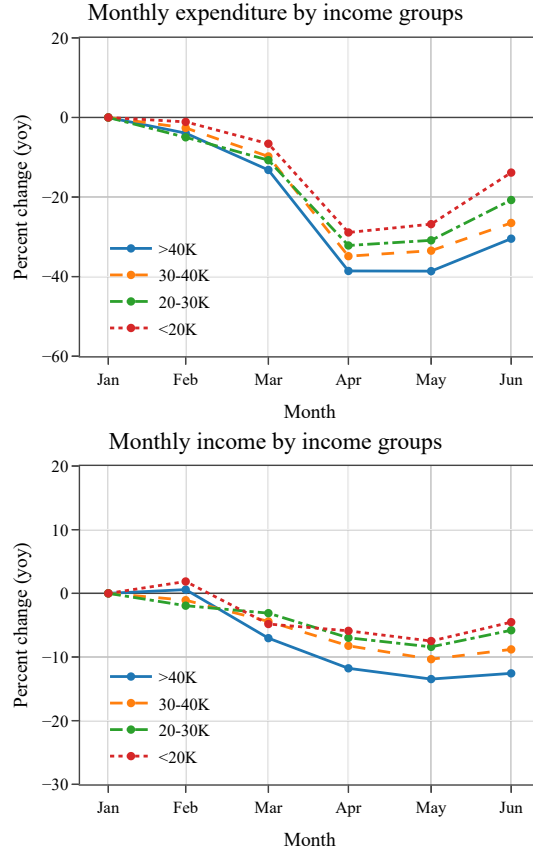


Notes: The figure compares the results on income and expenditure based on the automatic tags (left column) to the results based on user tags (right column).

Finally, we have also constructed expenditure and income measures that are based on all incoming/outgoing flows into the user's current (and credit for spending) accounts, including untagged transactions, net of transfers. Reassuringly, the results turn out to be qualitatively very similar to our baseline results, see Figure D.3. However, identifying transfers or refunds among the untagged

transactions is very difficult. Therefore, we view our approach as more conservative as relying on the tagged transactions avoids inflating spend and income by erroneously attributing transfers as actual expenditure/income.

Figure D.3: Including untagged transactions



Notes: Monthly expenditure and income by income groups based on measures including untagged transactions. The measures are constructed based on the outgoing/income flows into a user's current (and credit for spending) accounts, net of transfers. For the expenditure measure, we use all tagged transactions except transactions with a tag related to transfers, savings or credit card repayments (to avoid double counting). Furthermore, we include untagged transactions except transactions that are above 100 and exact multiples of 50, (as these are likely untagged transfers). For the income measures, we include all transactions with a relevant income tag as well as untagged transactions. Among the untagged transactions, we drop transactions with amounts below 100 pounds as well as transactions with amounts that are exact multiples of 50. This approach is similar to the one used in Bourquin et al. (2020).