



**Understanding Society
Working Paper Series**

No. 2020 – 15

October 2020

The Heterogeneous and Regressive Consequences of COVID-19: Evidence from High Quality Panel Data

Thomas F. Crossley, European University Institute

Paul Fisher, University of Essex

Hamish Low, University of Oxford

Non-technical summary

The onset of COVID-19 has caused a substantial contraction in economic activity, partly through individuals changing their behaviour in the face of the health risks and partly because of restrictions introduced by governments. In this paper we show that the scale of the economic impact of the COVID shock in the UK is very different for different people.

We find substantial impacts on labour markets, but these impacts do not show up in employment levels which changed very little between February 2020 and the end of May 2020. On the other hand, the fraction who are working a positive number of hours declined by 25 percentage points by the end of April, followed by a slight bounce back. The difference between the fraction employed and the fraction working positive hours highlights the effect of the Job Retention Scheme.

There is however substantial heterogeneity in the economic impact and mitigation strategies across groups. The young and those without any guaranteed hours of work experienced substantial falls in hours worked at the onset of COVID, and corresponding large falls in household earnings. But by the end of May, the decline in hours worked for these groups had been partly reversed, and further, household earnings showed less cumulative declines at this stage than for other groups. Mitigation of earnings losses by the young and precariously employed was partly through savings, but key components were finding new work and moving onto universal credit. Again, these mitigation strategies were very different from other groups. Minority ethnic groups experienced different labour market shocks: employment fell, in addition to the fraction working positive hours falling. Among those whose hours fell, there were 15 percentage points fewer from minority ethnic groups put on the Job Retention Scheme, and 13 percentage points more made unemployed.

The earnings losses that resulted were mitigated in different ways: the incidence of borrowing was higher, as was the incidence of transfers from family and friends. Finally, we show that the crisis has been regressive: those in the lowest long-run income quintiles have had the worst experiences. They have experienced the largest declines in the fraction working positive hours, and the largest declines in household earnings. For those in these lowest quintiles, these losses were mitigated by borrowing and by transfers from family and friends.

The Heterogeneous and Regressive Consequences of COVID-19: Evidence from High Quality Panel Data

Thomas F. Crossley, European University Institute, IFS, ESCoE and *Understanding Society*

Paul Fisher, University of Essex, ISER and *Understanding Society*

Hamish Low, University of Oxford, IFS and *Understanding Society*

Abstract: Using new data from the Understanding Society COVID-19 Study collected in two waves in April and in May 2020 in the UK, we make three contributions. First, Understanding Society is based on probability samples and the Covid-19 Study is carefully constructed to support valid population inferences. Second, the panel allows a long-run measure of income to characterise regressivity. Third, we have novel data on the mitigation strategies that households use. Our key findings are that those with precarious employment, under 30 and from minority ethnic groups face the biggest labour market shocks. Almost 50% of individuals have experienced declines in household earnings of at least 10%, but declines are most severe in the bottom income quintiles. Methods of mitigation vary substantially across groups: borrowing and transfers from family and friends are most prevalent among those most in need.

Keywords: COVID-19, job loss, inequality, mitigation, financial distress

The Understanding Society COVID-19 study is funded by the Economic and Social Research Council and the Health Foundation. Fieldwork for the survey is carried out by Ipsos MORI and Kantar. Understanding Society is an initiative funded by the Economic and Social Research Council and various Government Departments, with scientific leadership by the Institute for Social and Economic Research, University of Essex. The research data are distributed by the UK Data Service.

Data: University of Essex, Institute for Social and Economic Research. (2020). Understanding Society: COVID-19 Study, 2020. [data collection]. 3rd Edition. UK Data Service. SN: 8644, <http://doi.org/10.5255/UKDA-SN-8644-3>

The Heterogeneous and Regressive Consequences of COVID-19: Evidence from High Quality Panel Data *

Thomas F. Crossley[†] Paul Fisher[‡] Hamish Low[§]

August 10, 2020

Abstract

Using new data from the *Understanding Society* COVID-19 Study collected in two waves in April and in May 2020 in the UK, we make three contributions. First, *Understanding Society* is based on probability samples and the Covid-19 Study is carefully constructed to support valid population inferences. Second, the panel allows a long-run measure of income to characterise regressivity. Third, we have novel data on the mitigation strategies that households use. Our key findings are that those with precarious employment, under 30 and from minority ethnic groups face the biggest labour market shocks. Almost 50% of individuals have experienced declines in household earnings of at least 10%, but declines are most severe in the bottom income quintiles. Methods of mitigation vary substantially across groups: borrowing and transfers from family and friends are most prevalent among those most in need.

Keywords: COVID-19, job loss, inequality, mitigation, financial distress

JEL codes: C83, D31, J63

*We are grateful to Alexandra Brown for research assistance. The *Understanding Society* COVID-19 Study is funded by the Economic and Social Research Council (ES/K005146/1) and the Health Foundation (2076161). Fieldwork for the survey is carried out by Ipsos MORI and Kantar. Understanding Society is an initiative funded by the Economic and Social Research Council and various Government Departments, with scientific leadership by the Institute for Social and Economic Research, University of Essex. The research data are distributed by the UK Data Service. Further details and initial analysis of the *Understanding Society* COVID-19 Study circulated as Benzeval et al. (2020b). All errors are our own.

[†]European University Institute, IFS, ESCoE and and *Understanding Society* (email: thomas.crossley@EUI.eu).

[‡]ISER, University of Essex and *Understanding Society* (email: pfishe@essex.ac.uk).

[§]University of Oxford, IFS and *Understanding Society* (email: hamish.low@economics.ox.ac.uk).

1 Introduction

The onset of COVID-19 has caused a substantial contraction in economic activity, partly through individuals changing their behaviour in the face of the health risks and partly because of restrictions introduced by governments. In this paper we show that the scale of the economic impact of the COVID shock in the UK is very different for different people. These differences arise partly because the direct impact differs depending on individual characteristics and what sort of work people do, and partly because individuals take different steps to mitigate the shock. Further, government schemes to mitigate the impact may benefit some but not others. The aim of the paper is to highlight the idiosyncratic nature of the economic shocks and to show how heterogeneity across households mean the same economic shock has very different implications.

Our work makes three contributions to understanding the economic effects of COVID-19 and the role of the UK government in mitigating those effects. First, our results are based on a large, high quality longitudinal survey derived from probability samples. We use the first two waves of the *Understanding Society* COVID-19 Study alongside information from the long-running *Understanding Society* Main Study. We show the importance of using carefully modelled inverse probability weights and data derived from proper probability samples to capture accurately the differential effects of the aggregate shock; we also propose a statistical test of the efficacy of the weighting strategy. Second, the long panel data pre-COVID-19 provides a clear picture of the long-run incomes, and this provides crucial context of different household situations pre-COVID-19. Third, to address who is best able to mitigate the crisis and how the government has mitigated the crisis, the COVID-19 web surveys provide novel information on what mechanisms have mitigated losses for individuals. These three contributions enable us to provide a unique perspective on how COVID-19 has changed the economic reality faced by different individuals in the UK.

The backdrop to these changes in labour market status is how the UK government supported workers and households. Along with much of Europe, the UK government pursued a policy of explicitly protecting jobs through the Job Retention Scheme. This scheme allowed workers to be “furloughed” by their firms, which meant 80% of pay would be covered by a government subsidy, subject to a maximum of £2500, and was conditional on the worker not providing any hours of work. This is in contrast to the US where support operated through additional payments to the unemployed. The nature of support in the UK is crucial to understanding why increases in unemployment were so limited in the first stages of the crisis.

We split our analysis into showing differences in the economic shock to labour markets and then into showing differences in actions take to mitigate the economic shock. We find substantial

impacts on labour markets, but these impacts do not show up in employment levels which changed very little between February and the end of May. On the other hand, the fraction who are working a positive number of hours declined by 25 percentage points by the end of April, followed by a slight bounce back. The difference between the fraction employed and the fraction working positive hours highlights the effect of the Job Retention Scheme.

There is however substantial heterogeneity in the economic impact and mitigation strategies across groups. The young and those without any guaranteed hours of work experienced substantial falls in hours worked at the onset of COVID, and corresponding large falls in household earnings. But by the end of May, the decline in hours worked for these groups had been partly reversed, and further, household earnings showed less cumulative declines at this stage than for other groups. Mitigation of earnings losses by the young and precariously employed was partly through savings, but key components were finding new work and moving onto universal credit. Again, these mitigation strategies were very different from other groups.

Minority ethnic groups experienced different labour market shocks: employment fell, in addition to the fraction working positive hours falling. Among those whose hours fell, there were 15 percentage points fewer from minority ethnic groups put on the Job Retention Scheme, and 13 percentage points more made unemployed. The earnings losses that resulted were mitigated in different ways: the incidence of borrowing was higher, as was the incidence of transfers from family and friends.

Finally, we show that the crisis has been regressive: those in the lowest long-run income quintiles have had the worst experiences. They have experienced the largest declines in the fraction working positive hours, and the largest declines in household earnings. For those in these lowest quintiles, these losses were mitigated by borrowing and by transfers from family and friends.

Our results contribute to a fast-moving literature looking at the labour market consequences of COVID-19 and the impact of government support schemes. Much of the evidence comes from rapid surveys with quota samples, or convenience samples ex post calibrated to population totals on the basis of a limited number of observable characteristics like age and gender. For example, Adams-Prassl et al. (2020) report large effects in the U.K. on the young, on women and on those in insecure work, and (Belot et al., 2020) report a similar age effect. These studies have given initial indications of the effects of COVID-19 for the sample surveyed, but can only provide population estimates under very strong assumptions about sample inclusion and response.

For the US, the March 2020 US Current Population Survey is derived from reliable probability samples and it shows increased unemployment, decreased working hours, but little fall in wages (Béland, Brodeur and Wright, 2020). The labour market impacts have been shown to be bigger for

men, younger workers, Hispanics, and the less educated (Cortes, 2020); although (Montenovo et al., 2020) finds larger effects for women and those with larger families. Larger shocks for vulnerable populations have also been documented in data from probability samples for the Netherlands (von Gaudecker et al., 2020) and Germany (Schröder et al., 2020); and from population registers for Norway (Alstadsæter et al., 2020). The present paper is the first to report credible population estimates for the UK.

The paper proceeds as follows: in Section 2, we outline the data; Section 2 gives addition detail about our data and Section 3 discusses drawing population inferences from survey data. Section 4 reports the impact of these shocks on household earnings and further, documents the heterogeneity in how individuals are able to mitigate the impact of the shocks. Section 6 concludes.

2 Data

This paper is based on the first two waves of the *Understanding Society COVID-19* study (henceforth COVID-19 Study), fielded in, respectively, late April and late May of 2020; these surveys also collected retrospective information about February 2020.¹ The UK economy went into lockdown on 23rd March, while on 20th March, the UK government introduced the Job Retention Scheme, and subsequently the Self-Employment Support Scheme.

The COVID-19 study is built upon *Understanding Society*: the UK Household Longitudinal Study (henceforth Main Study) and uses monthly web surveys to capture the experiences and behaviour of Main Study participants during the initial phase of the COVID-19 pandemic. This means, first, that the COVID-19 Study inherits the properties of the Main Study that ensure reliable population inferences, and second, that data collected by the COVID-19 Study can be linked to data collected on the same participants, and their households, by past waves of the Main Study.²

With minor exceptions, all members of the Main Study who were aged sixteen or over in April 2020, and who belonged to active households, were invited to participate in the COVID-19 Study.³ Pre-notification letters introducing the study were sent to 42,330 Main Study members on 17 April. Of these, 32,596 had completed the Wave 9 annual interview (the latest released wave of the Main

¹Further information on the *Understanding Society* COVID-19 Study can be found in Institute for Social and Economic Research (2020b) and Institute for Social and Economic Research (2020a).

²*Understanding Society* (University of Essex Institute for Social and Economic Research, NatCen Social Research and Kantar Public, n.d.) is the UK's main longitudinal Household Survey, and one of the largest household panel studies in the world. It began in 2009 but carries on from the earlier *British Household Panel* survey which ran from 1991 to 2008. *Understanding Society* attempts to interview all adults in sample households annually and has a mixed mode design, with some panel members responding via a face-to-face interview and some completing a web interview.

³An active house is one that participated in at least one of the last two waves of the main study

Study). Respondents were offered a small financial incentive for each web survey. Subsequently, invitations to each web survey were sent by email and/or SMS text message, or by post. Each web survey had a 7-day fieldwork period and reminders were sent on days 2, 3, and 6.⁴ Each web questionnaire took approximately 20 minutes to complete.

Among those who had given a full adult interview in the Wave 9 annual interview, the response rates to April and May web surveys were 48.6% and 49.1 % respectively.⁵ These response rates are similar to the response rates of large government surveys in the UK.⁶ These are very good response rates for a voluntary web survey that attempts to contact a known set of individuals (so that non-respondents are identified: convenience and quota samples do not have knowable response rates). Nevertheless, this is significantly below the 85-90 % overall wave-on-wave retention rate that the Main Study achieves by following up web non-respondents by direct interviewer contact.

There were 17,452 respondents to the April web survey and 14,811 in May. Most of our analysis focuses on individuals that reached the end of the survey (“full respondents”). The COVID-19 study weighting strategy, which we describe in more detail in the next section and which is the basis for our population inferences, assigns a positive sample weight to respondents who also responded to Wave 9 of the main study and had a positive Wave 9 sample weight. This gives a basic analysis sample of 10,892 individual respondents. We further restrict our attention on respondents aged 20 to 65 in order to focus on the working age population and exclude a small number of respondents who provide incomplete February 2020 hours or employment information. This gives a final analysis sample of 7,404 individuals.

For our distributional analysis we created a measure of “long-run” income. This measure averages household net income across up to three previous waves of the main study, and assigns individual respondents to quintiles of long-run income on that basis. For this purpose income includes earned and unearned income, net of tax and inclusive of any benefits received, equivalised by household composition. It is important to note that the COVID-19 study is individual-based, and supports inferences about the distribution of income (for example) across adults rather than across households. Household income and other household variables are viewed as attributes of individuals.

⁴The first web survey was fielded on April 24th and the second on May 27th.

⁵5,519 fewer set main sample members were invited to the May survey, as some April non-respondents were issued to a telephone follow up survey instead, and a further group either entirely opted out of the COVID-19 study at the April invitation or were determined to be no longer eligible.

⁶For example, the Labour Force Survey - to which many web surveys with quota or convenience samples calibrate - has response rate of about 55% at the first wave, falling with subsequent and about 40 % overall. The Family Resources Survey which is the basis for official income statistics had a response rate of 52% in 2017/18.

3 Population Inferences

Distributional analysis is inherently about estimating finite population quantities and gradients. The ability of the Main Study, and by extension the COVID-19 Study, to deliver credible population estimates rests on the fact that the Main Study is based on probability samples, and on the use of carefully designed inverse-probability (IP) weights.⁷

A defining feature of probability samples is that every unit in the target population has a knowable, nonzero probability of selection (Valliant and Dever, 2018). This offers two important advantages over other types of samples (such as convenience or quota samples). First, the fact that all units in the target population have a nonzero probability of selection ensures that, with sufficiently large sample sizes, the full range of heterogeneity in the target population will be captured. Second, known selection probabilities mean that consistent estimates of population parameters and associated inferences can be obtained with well-established statistical methods involving IP weighting (see Wooldridge (2002) and the references therein).

Of course, real samples deviate from the theoretical ideal of a probability sample because of non-response, including, in the case of longitudinal studies, attrition. Nevertheless, there are multiple advantages to beginning from probability samples. First, while statistical adjustments may be needed to account for nonrandom attrition and non-response, such adjustments will be smaller if the initial selection probabilities are known. Second, when a study begins with a probability sample, useful information is often available on non-respondents. This is particularly true in longitudinal studies where rich information on individuals who attrit is available from past waves of the survey. When information is available on both respondents and non-respondents, the models of response probability that underlay IP weights can be estimated directly. In contrast, with convenience or quota samples information is only available for respondents, and the relationship between response probabilities and observable characteristics can only be inferred indirectly by comparing sample characteristics to external totals, ideally from a census or register, but often from a probability-sample based survey. This approach leads to what we refer to as “calibration weights” because the procedure calibrates a sample with entirely unknown inclusion probabilities to external totals. This procedure is less efficient, but more importantly, the set of variables used in the adjustment is typically very limited (for example, just age, education and gender). As discussed in Moffit, Fitzgerald and Gottschalk (1999); Wooldridge (2002), weights correct for selection on observables, and so the richness of observable predictors of response is critical. Finally, other types of samples

⁷The *Understanding Society* Main Study is a combination of four different probability samples. See University of Essex Institute for Social and Economic Research, NatCen Social Research and Kantar Public (n.d.) for more details.

may have a zero probability that certain parts of the target population will enter the sample. No weighting scheme can overcome the complete absence of a subgroup from the sample.

The Main Study employs state of the art methods to minimize non-response and attrition. It also provides carefully-modelled IP weights to account for the nonrandom nature of the remaining attrition.⁸ The extent to which the Main Study is “representative”, in the sense of supporting high quality inferences about population quantities, is continually evaluated: see Benzeval et al. (2020a) and the references therein.⁹ The study has been repeatedly judged to be of high quality: as just one example, Main Study income data aligns well with national statistics on the income distribution in the UK (Fisher et al., 2019). Given this robust evidence that (suitably weighted) waves of the Main Study provide reliable estimates of population quantities and gradients, the remaining issue is non-response to COVID-19 Study among respondents to Wave 9 of the Main Study.

IP weights are released with each wave of the COVID-19 Study.¹⁰ These weights were created via an adjustment to the cross-sectional weights available for Wave 9 of the Main Study. This means that the probability of response to each wave in the COVID-19 Study is modeled as the product of the conditional probability of response to that survey (given Main Study Wave 9 response) and the probability of Wave 9 response. The conditional probability of a response to a COVID-19 wave is modeled as function of information known at the time of issue to the COVID-19 Study. The resulting weights map the set of respondents to a given COVID-19 wave back to the target population at the time of Wave 9 (2017/18).¹¹

The choice set of predictors for response include basic demographics, household composition, economic variables and health variables, all drawn from the rich information collected by past waves of the Main Study. Note again that because the target sample is drawn from the Main Study, this information is available for both respondents and non-respondents to the COVID-19 Study. In addition, both the econometrics and survey statistics literatures (Moffit, Fitzgerald and Gottschalk, 1999; Nicoletti and Peracchi, 2005) emphasize the importance of including in weighting models variables that predict response and are potentially correlated with outcomes being studied, but are unlikely to be included in standard economic or social science models. Examples include previous survey outcomes, survey design variables and survey para data. Several such variables turn out to be good predictors of the conditional probability of response in the COVID-19 Study. These include indicator variables for the types of contact information the survey team held about the respondent prior to the COVID-19 Study (email address, mobile phone number, both, neither)

⁸The development of the Main Study weights is described in Lynn and Kaminska (2010)

⁹“Representative” is a widely-used but ill-defined term. See the discussion in Benzeval et al. (2020a)

¹⁰The weights were developed by the authors of this paper

¹¹Updated for subsequent mortality and emigration, but not immigration.

and the realized mode of previous waves of the main survey. The former may affect the salience of the survey request while the latter may be related to how easily the respondent would find it to complete a web survey. Either could quite plausibly be related to whether the respondent is employed or to the kind work they do. Variable selection for the final models from the initial set is done by LASSO.

In addition to the cross-sectional weights for the April and May waves, we also employ a weight for the balanced panel of individuals who responded to both the April and May waves. This weight was derived using the same methods.

It is common place to assess survey weights or weighted data by comparing summary statistics to some benchmark. We go further a propose a formal statistical test for whether the weights capture the probability of retention to wave t of a longitudinal data set, given response at wave $t - 1$, and then apply this test to the COVID-19 Study (given response to Wave 9 of the Main study.) While the literature contains a number of tests for whether panel attrition is random (see for example, Fitzgerald, Gottschalk and Moffitt (1998)), the test we propose tests instead whether the weights deal adequately with nonrandom attrition. To the best of our knowledge our proposed test is novel.

Let $Y_{t-1,i}$ be an observation of any variable of interest, Y for individual i in wave $t - 1$; $R_{t,i} = 1$ if responds to wave t (of the COVID-19 Study) and 0 otherwise; and similarly $R_{t-1,i} = 1$ if the individual responds to wave $t - 1$ (here Wave 9 of the Main Study). $X_{t-1,i}$ is a set of predictors of response observed for both respondents and nonrespondents, prior to the realization of $R_{t,i}$ (up to and including time $t - 1$). $X_{t-2,i}$ is defined analogously. Note that these may contain lagged values of Y . Let $w_{t-1,i}(X_{t-2,i}) < \infty$ be the wave $t - 1$ weight. This is the inverse of the wave $t - 1$ response probability. Analogously, $w_{t,i}(X_{t-1,i}) < \infty$ is the wave t weight. In this case, this is one of the COVID-19 Study IP weights. $s_{t-1,i} = \frac{w_{t-1,i}}{\sum w_{t-1,i}}$ is the wave $t - 1$ weight share and s_{ti} is the wave t weight share, defined analogously.

Under the joint null that

$$E[R_{t-1,i}|Y_{t-1}, X_{t-2}] = E[R_{t-1,i}|X_{t-2}] = w_{t-1,i}(X_{t-2})$$

and

$$E[R_{t,i}|X_{t-1}, Y_{t-1}] = E[R_{t,i}|X_{t-1}] = w_{t,i}(X_{t-1}),$$

i.e. the response to the relevant waves is independent of Y_{t-1} given pre-response observables:¹²

¹²We also require the technical condition $E[Y_{t-1,i}/w_{t-1,i}] < \infty$ and $E[Y_{t-1,i}/w_{t,i}] < \infty$; Wooldridge (2002).

$$E[s_{t-1,i}R_{t-1,i}Y_{t-1,i} - s_{t,i}R_{t,i}Y_{t-1,i}] = 0 \quad (1)$$

This moment condition captures the fact that under the joint null, either combination of respondents and associated weights provide a consistent estimate of $E[Y_{t-1}]$, and this provides a simple statistical test of the adequacy of the weights. All of the sets of weights developed for the Covid-19 Study are subjected to these tests, for a wide range variables of interest ($Y_{t-1,i}$). As an illustration, Table 1 reports the results of test of this type using the Wave 1 cross-section weight. The first column shows the estimated population mean of $Y_{t-1,i}$ using the Wave 9 (main study) response sample and associated weights. The next three columns use show estimates of the same mean using only respondents to the April COVID-19 Study, either unweighted (Column 2), with a crude calibration weight (Column 3), or with the full IP weights (Column 4). The calibration weight matches the April COVID-19 Study data to the Wave 9 Main Study data on the basis of set of cells defined by gender, age and education. It mimics the kind of calibration weights often employed with convenience samples, or the composition of a quota sample.¹³

The last two columns report test statistics based on Equation (1) and associated p-values, for the calibration weights (column 6) and full IP weights (Column 7). The test is reported for an illustrative set of variables, $Y_{t-1,i}$, with each row of the table corresponding to a different $Y_{t-1,i}$. These $Y_{t-1,i}$ variables are presented in two groups. The first group are variables that are included in the estimated model of response to wave 1 of the COVID-19 Study. These include measures of subjective financial satisfaction, housing tenure, occupation, savings behavior and financial arrears. Table 1 demonstrates that the null is rejected for calibration weights for all of these variables, but when the full IP weights are employed, the null is rejected only for the percentage owning a home with a mortgage. The magnitude of the moment is the difference between the weighted sample mean of the variable $Y_{t-1,i}$ and its target value in the Wave 9 Main Study. A comparison of the two columns shows that the reduction in the size of this difference is economically significant when moving from the calibration weights to the full IP weights. For example, using the calibration weights, we would overestimate the fraction of individuals reporting that they were living comfortably by 4 percentage points, and overestimate the fraction managing to save some of their income by 8 percentage points, and underestimate the fraction of individuals living in social housing by 9 percentage points.

The second set of $Y_{t-1,i}$ variables we consider are those that are *not* included in the the estimated model of response. These are indicators for being in poverty, in receipt of core benefits, in arrears

¹³See for example Adams-Prassl et al. (2020) and Belot et al. (2020)

Table 1: Statistical Testing of Survey Weights

	<u>Wave 9</u> Weighted	Unweighted	<u>Covid</u> Calibration weight	Full IP weight	<u>Test</u> Calibration	Full IP
<u>In Full IP weight only:</u>						
Subjective finances:						
<i>Living comfortably/ doing alright</i>	0.71	0.76	0.76	0.71	-0.04*** (0.000)	0.00 (0.632)
<i>Just about getting by</i>	0.21	0.18	0.18	0.21	0.03*** (0.000)	0.00 (0.616)
<i>Finding it quite/ very difficult</i>	0.07	0.06	0.06	0.08	0.01*** (0.000)	-0.01 (0.109)
Housing tenure:						
<i>Owned</i>	0.34	0.40	0.41	0.33	-0.06*** (0.000)	0.01 (0.076)
<i>Mortgage</i>	0.34	0.42	0.40	0.36	-0.07*** (0.000)	-0.02*** (0.000)
<i>Rented</i>	0.13	0.09	0.09	0.12	0.03*** (0.000)	0.00 (0.487)
<i>Social Housing</i>	0.19	0.09	0.10	0.19	0.09*** (0.000)	0.00 (0.587)
Low skill occupation	0.38	0.29	0.33	0.38	0.05*** (0.000)	-0.00 (0.802)
Any savings income	0.36	0.45	0.44	0.37	-0.08*** (0.000)	-0.01 (0.086)
Behind with some or all bills	0.05	0.03	0.03	0.05	0.02*** (0.000)	0.00 (0.685)
<u>In neither weighting model:</u>						
Poverty	0.15	0.11	0.12	0.14	0.03*** (0.000)	0.01 (0.320)
Receives core benefit	0.05	0.03	0.03	0.05	0.02*** (0.000)	-0.00 (0.755)
Behind with housing	0.09	0.06	0.06	0.09	0.03*** (0.000)	0.00 (0.730)
Smoker	0.15	0.09	0.10	0.13	0.05*** (0.000)	0.02** (0.002)
Long-standing illness	0.38	0.34	0.34	0.36	0.04*** (0.000)	0.02* (0.029)

Notes: P-values are reported in parenthesis. * p<0.05, ** p<0.01, *** p<0.001. IP weights are Inverse Probability weights.

'Core benefits' include Income Support, Job Seeker's Allowance and Universal Credit.

on rent or mortgage payments, being a current smoker, or having a long-standing illness. Again we see that the null is always rejected with the calibration weights, but less often when the full IP weights are employed. The calibration weights lead to an overestimate of all of the conditions, by between 2 and 5 percentage points. The point to stress is that IP weights reduce the bias in all cases, and eliminate it entirely in some. The results in Table 1 indicate that IP weights associated with the COVID-19 Study are very effective in adjusting for nonrandom attrition between Wave 9 of Main study and the COVID-19 Study, and that the IP weights provide much a more credible basis for population inferences than simple calibration weights.

Finally, the underlying *Understanding Society* samples are clustered and stratified random samples, and so the COVID-19 sample inherits this structure. We appropriately adjust standard errors for the resulting design effects.

4 Labour Market Shocks

In this section, we show the extent of heterogeneity and regressivity in the labour market shocks that individuals face, and how these shocks have evolved in the first three months of the pandemic. We describe labour market status using two main measures: whether an individual is employed and whether they are working a positive number of hours. We chose these measures because the UK Job Retention Scheme aims to maintain the employment relationship despite individuals not working any hours.

In Table 2, we show the fraction employed and the fraction working positive hours using reports on February, April and May 2020. The table disaggregates these measures by individual characteristics, including gender, ethnicity, age, long-run income quintile, household type and worker type. The first column shows (retrospective) numbers for the February 2020 employment "baseline" using the April respondent sample and associated cross-sectional weights. The remaining columns of the Table are based on the balanced panel of respondents to both April and May surveys, and the associated balanced panel weight. Comparing the first and second column confirms that the balanced panel (and associated weight) matches very closely the full cross-sectional numbers at a point in time.¹⁴

The table shows that, in aggregate, employment levels have changed very little between February and late May. This highlights clearly the effectiveness of the government Coronavirus Job Retention Scheme that aimed to preserve employment. On the other hand, the fraction of individuals who

¹⁴We present the February employment comparison as an illustration but this is also true of other measures and at other points in time.

Table 2: Labour Market Shocks by Individual Characteristics

	Empl. x-sec. (Feb)	Empl. (Feb)	Empl. (April)	Empl. (May)	+ve hours (Feb)	+ve hours (April)	+ve hours (May)
All	79	79	77	77	79	54	58
Gender:							
Men	83	83	80	81	82	57	61
Women	76	76	74	74	75	51	55
Ethnicity:							
Not BAME	80	80	78	79	80	54	59
BAME	71	70	63	64	69	48	50
Age:							
Age 20-29	78	76	69	72	75	44	52
Age 30-39	86	86	84	84	85	58	62
Age 40-49	85	86	85	85	85	63	67
Age 50-59	80	81	80	80	81	57	61
Age 60-65	59	57	55	55	57	37	40
Household type:							
Adult, no child	69	72	69	70	72	50	55
Adult, child	71	72	72	72	71	46	48
Adults, no child	78	77	74	74	76	52	56
Adults, child	85	85	83	84	84	57	63
Long-run income quintile:							
1	60	62	58	59	61	37	39
2	78	76	73	73	76	45	51
3	85	85	83	84	85	57	63
4	87	86	84	85	85	62	66
5	86	86	84	83	84	67	69
Worker type:							
Fixed hours	100	100	96	97	99	71	76
Flexible hours	100	100	95	95	98	73	74
Emp. sets (sure min.)	100	100	96	96	98	62	65
Emp. sets (no min.)	100	100	69	78	95	34	44
Self-employed	100	100	96	95	99	54	64

Notes: Employment includes both employees and self-employees. Sample sizes are 9531 (column 1) and 7404 (columns 2-7). Household type is measured in May and Worker type in February.

are working a positive number of hours fell 25 percentage points to April and only 54% of working age individuals were working positive hours. There was a slight recovery to 58% in the fraction working positive hours in May. However, the key point is that the difference between the fraction employed and the fraction working-positive-hours highlights starkly the potential unemployment problem in the coming months, as the main Job Retention Scheme tapers off from August, and comes to an end in October.

The absence of any impact on employment is difficult to reconcile with the substantial job losses reported by Adams-Prassl et al. (2020). However, our numbers are consistent with the only other probability-sample based employment data for the UK that we are aware of, the Office of National Statistic's Labour Force Survey. That data also show almost no change in employment up to May 2020.¹⁵

The breakdown by individual characteristics in Table 2 shows that while the labour market consequences have been felt across the board, there are some groups that have been particularly impacted, and others that were initially impacted but have rebounded more. Hardest hit initially were those individuals where the employer does not guarantee any minimum number of hours: of those employed on such a zero hours contract in February, the fraction working positive hours fell by the end of April from 95% to 34%. This was followed by a 10 percentage point bounce back in May. Sharp initial declines in the fraction working positive hours and in employment were seen for those aged under 30, but again has bounced back a little. This highlights the double-edged nature of the flexibility that comes with zero-hours contracts and the sort of jobs typically carried out by young workers. Looking at the labour market impact by ethnicity, individuals from minority ethnic groups experienced a substantially larger fall in employment than others, whereas the overall decline in the fraction working positive hours was similar. Across the distribution, the bottom three quintiles experienced the greatest reductions in the fraction working positive hours.

Table 3 shows reported reasons for the fall in hours worked for those who have experienced a decline by the end of May. In the population, 63% of individuals reported some decline in hours worked. This decline in hours may be caused directly by the health shock, indirectly by impacts on the economy due to the health shock, or for non-health related reasons. The first point to take from Table 3 is that the decline in hours is driven by the economic restrictions rather than directly by health or caring. Over 36% of those reporting a decline in hours were part of the Coronavirus Job Retention Scheme and so had substantial earnings replacement. In addition, 8% report the decline in their hours being caused by being made unemployed. By contrast, only 6% report health as a reason for the decline in hours, and 7% report caring for others. However, these averages mask considerable heterogeneity across different types of individual. Caring is more important for those with children; and the effect of health restricting work increases sharply with age.

Ethnicity is associated with a very different explanation for the hours decline: individuals from minority ethnic groups are 15 percentage points less likely to be supported by the Job Retention Scheme. Instead, they are 13 percentage points more likely to cite unemployment as the reason for

¹⁵<https://www.ons.gov.uk/employmentandlabourmarket/peopleinwork/employmentandemployeetypes/bulletins/uklabourmarket/june2020>

Table 3: Reasons for Decline in Hours by May

	% with hours fall	Emp. cuts	Furlough	Unemp.	Loss of self-emp. business	Health	Caring
Report an hours fall	63	10	36	8	10	6	7
Gender:							
Men	63	10	35	9	12	6	5
Women	63	10	36	8	9	6	8
Ethnicity:							
Not BAME	63	10	37	7	10	6	6
BAME	64	12	22	20	13	10	13
Age:							
Age 20-29	66	9	46	18	5	5	3
Age 30-39	61	9	36	7	8	4	15
Age 40-49	60	8	33	4	12	6	11
Age 50-59	63	13	32	6	12	7	2
Age 60-65	67	8	29	10	14	11	1
Household type:							
Adult, no child	59	15	28	14	9	8	1
Adult, child	60	13	46	2	10	2	13
Adults, no child	62	10	38	9	9	6	2
Adults, child	64	9	35	7	11	6	14
Long-run income quintile:							
1	69	13	35	12	14	9	8
2	69	10	44	8	9	6	8
3	62	9	44	5	8	5	6
4	59	9	32	9	7	5	6
5	58	9	22	9	13	6	7
Worker type:							
Fixed hours	56	11	43	8	0	5	5
Flexible hours	67	8	29	8	0	2	11
Emp. sets (sure min.)	72	19	49	6	0	6	4
Emp. sets (no min.)	89	11	49	35	0	4	3
Self-employed	86	2	6	6	53	11	14

Notes: The table refers to 6038 individuals employed in either February or May or both. Columns 2-7 refer to a percentage of the population experiencing a decline in weekly work hours by May. Respondents are allowed to report multiple reasons for an hours decline and so the rows do not sum to one. See table 2 notes.

their hours decline. This greater prevalence of unemployment among minority ethnic groups was also shown in Table 2 above.

Our overall conclusion is that the labour market effects are highly heterogeneous, particularly impacting the young, zero hours workers and minority ethnic groups, and regressive, penalizing most the lower quintiles of the long-run income distribution. The final point that Table 3 highlights is the importance of the Job Retention Scheme in the UK which has maintained many of those not actually working any hours notionally in employment. This is in marked contrast with the US where support operated through the extension of unemployment insurance without the same attempt to maintain attachment to the employer.

5 Earnings Losses and Mitigation

Section 4 showed that the effect of the COVID crisis on labour market outcomes varies substantially across individuals and across the distribution. The Job Retention Scheme moderates the link between hours and earnings for individuals, and the evidence in Section 4 shows that for many groups, it has been effective in doing so. Individual earnings losses may be cushioned by the earnings of other household members. Benefits, particularly Universal Credit, should then moderate the link between earnings and income. Finally, individuals and households may take steps to moderate the link between income and living standards. In this section, we start by considering how the labour market shocks documented above translate into a change in net household earnings. We then document the incidence of different mitigation strategies, including applications for universal credit, dis-saving, borrowing, transfers from friends and family and the use of foodbanks. The point we stress throughout this section is that the impacts on earnings are highly heterogeneous and regressive. Moreover, there are important differences in mitigation strategies across groups and across individuals.

Table 4 reports the impact on household earnings through April and May, and across the distribution. The measure of earnings we use is net, equivalised weekly household earnings of the individual respondents, including earnings from employment and self-employment.¹⁶ Average household earnings declined by 10% by the end of April, with a further 5% decline by the end of May. The right-hand side of Table 4 shows the distribution of the change in earnings, showing the 25th percentile, the median and the 75th percentile, and separately for April and for May. This highlights the extent of losses and the extent that these losses have worsened: the 25th percentile of the change was an 18% decline by April, but by May the 25th percentile was a 41% decline. The median earnings change has also deteriorated. On the other hand, the 75th percentile of the earnings change is positive in May.

When we consider the impact on earnings by long-run income level, the impact is increasingly severe the lower down the long-run income distribution. In the bottom quintile, the median fall in earnings was 13% by May, whereas in the top quintile, the median fall was only 2%. The differences are equally stark at the 25th percentile of the percentage change: in the bottom quintile, the 25th percentile fall was 60%. The impact of COVID on earnings has been highly regressive.

There is also substantial heterogeneity both between and within groups. Between groups, the

¹⁶In the COVID-19 Study, individual and household earnings are collected with single questions. This differs from the main study, which aggregates information from more detailed questions, but was necessitated by the strategy of brief, but frequent, web surveys. (Micklewright and Schnepf, 2010) is one assessment of such “single question” income data collection in surveys. Individual earnings are top-coded at £4,000 net per week, and household earnings is top-coded where the difference between household and individual earnings exceeds £4,000 net per week.

Table 4: Household Earnings Pre and Post Covid

	Feb	April	May	% change since Feb						
				April			May			
				Mean	p25	p50	p75	p25	p50	p75
All	549	501	478	-18	0	0	0	-41	-6	6
Gender:										
Men	573	519	490	-17	0	0	0	-44	-7	5
Women	526	484	467	-18	0	0	0	-38	-5	6
Age:										
Age 20-29	503	454	482	-20	0	0	0	-32	-2	20
Age 30-39	564	526	502	-14	0	0	0	-40	-5	5
Age 40-49	571	534	487	-16	0	0	0	-42	-8	1
Age 50-59	593	534	508	-18	0	0	0	-40	-5	5
Age 60-65	446	386	346	-30	0	0	0	-76	-13	9
Ethnicity:										
Not BAME	564	515	490	-17	0	0	0	-40	-6	5
BAME	411	370	368	-22	0	0	0	-42	-4	8
Household type:										
Adult, no child	560	471	461	-50	0	0	0	-51	-2	3
Adult, child	266	230	249	-20	0	0	0	-33	0	2
Adults, no child	588	540	518	-17	0	0	0	-42	-5	9
Adults, child	523	487	457	-17	0	0	0	-38	-8	4
Long-run income quintile:										
1	287	245	228	-31	0	0	0	-60	-13	4
2	395	356	365	-20	0	0	0	-36	-6	7
3	487	444	428	-15	0	0	0	-34	-3	4
4	664	593	559	-14	0	0	0	-43	-8	4
5	860	817	765	-12	0	0	0	-39	-2	8
Worker type:										
Fixed hours	624	598	556	-10	0	0	0	-30	-2	5
Flexible hours	704	660	616	-9	0	0	0	-36	-2	6
Emp. sets (sure min.)	479	455	433	-21	0	0	0	-42	-14	6
Emp. sets (no min.)	605	369	512	-44	-20	0	0	-35	0	59
Self-employed	551	390	466	-64	-25	0	0	-57	-24	6

Notes: Earnings are weekly, net and equivalised. Sample size: 6160 individuals (col. 1-3) and 5673 individuals (col. 4-9) reporting positive February earnings.

young initially experienced greater earnings changes, but by May, this position had reversed. Indeed, by May, in terms of earnings declines, young individuals show the least negative impact of any age group on household earnings. This rebound in earnings reflects the labour market changes over time shown in Table 2. Among the young, there was also a range of winners and losers: at the 75th percentile, the change in earnings was a 20% increase. Similar within group heterogeneity is

seen for those on zero-hour contracts and for the self-employed.

These declines in household earnings do not necessarily translate into declines in household income or living standards, partly because households have other sources of income, such as universal credit, but also because different households have different possibilities of mitigation. However, as with the earnings losses, there is substantial heterogeneity across groups, and across the income distribution.

Table 5 summarizes actions taken to mitigate earnings losses by individuals who have reported a decline in household earnings of 10% or more by May. Just under half of the population have experienced such a loss, but the extent of lost earnings would have been substantially larger without the protection of the Job Retention Scheme. Mitigation through self-insurance includes the use of savings, borrowing, additional work, or mortgage holidays. Mitigation from external sources includes transfers from other family or friends as well as making new applications for state benefits (through universal credit) or the use of food banks.

At this stage in the crisis, more individuals have used self-insurance, and in particular on their own savings, than have accessed external help to mitigate losses: more than a quarter have drawn down their savings. Significant numbers have also increased borrowing or asked for a mortgage holiday. New applications for Universal Credit were 8%,¹⁷ but overall less prevalent than savings or family transfers. However, different individuals have used very different mechanisms. Across the income distribution, borrowing and transfers from family and friends increase sharply as long-run income declines. Borrowing and transfers are also much more marked for single parents; and similarly for minority ethnic groups, where increased borrowing is three times as likely. By contrast, for those on zero hour contracts, a third report finding new work and a third report newly receiving universal credit.

The key question these results raise is how long support from family and friends, or from individuals' own saving and borrowing, can continue. The end of the Job Retention Scheme will lead to more widespread and deeper earnings losses potentially at the same time as these support mechanisms will potentially be exhausted.

¹⁷Note this implies over a million new applications.

Table 5: Mitigating Earnings Losses

	% with earn. loss	Used savings	Borrowed	New work	Mortgage hol.	New Universal credit	SEISS	Family transfer etc.	Used foodbank
All	45	26	8	3	8	8	8	12	2
Gender:									
Men	47	25	8	4	7	9	10	12	1
Women	44	27	8	3	10	7	7	13	4
Ethnicity:									
Not BAME	45	25	7	3	8	8	8	12	2
BAME	42	36	21	4	11	11	12	15	3
Age:									
Age 20-29	41	22	10	8	5	14	4	18	3
Age 30-39	45	23	9	4	13	9	6	16	2
Age 40-49	46	25	11	3	13	8	11	11	1
Age 50-59	45	32	6	2	6	9	10	12	1
Age 60-65	51	26	2	1	2	3	8	4	7
Household type:									
Adult, no child	44	25	7	6	5	21	9	22	4
Adult, child	37	35	22	5	5	5	5	40	12
Adults, no child	43	26	4	2	4	7	7	8	3
Adults, child	48	25	11	4	14	7	10	14	1
Long-run income quintile:									
1	52	31	16	4	6	14	11	27	8
2	46	32	11	2	10	9	11	17	2
3	40	25	6	5	8	9	6	9	2
4	47	21	4	4	9	6	4	7	0
5	42	23	4	2	10	4	9	4	0
Worker type:									
Fixed hours	40	21	6	2	9	4	0	9	1
Flexible hours	42	18	4	1	7	6	0	5	0
Emp. sets (sure min.)	53	22	11	6	5	17	0	16	1
Emp. sets (no min.)	41	25	3	31	16	36	0	11	1
Self-employed	59	53	15	7	12	23	54	17	1

Notes: Each cell refers to a percentage of individuals experiencing a household earnings loss of at least 10 percent between February and May. Methods of mitigation were collected in both April and May and respondents can report multiple methods of mitigation at each monthly interview. Sample size: 2617. SEISS refers to the “Self-employment Income Support Scheme”. See table 2 notes.

6 Conclusions

This paper shows that the aggregate effects of COVID-19 mask considerable differences in how individuals are affected both because they are differentially exposed to the labour market shocks and because they have different access to private and public support mechanisms. We present results from new high-quality UK data: the *Understanding Society* COVID-19 Survey. Given the goal of documenting the distribution of effects it is critical to be able to estimate reliably population and subpopulation quantities.

Two months after the “stay at home” policy was introduced at the end of March, unemployment

had barely increased in the UK. However, almost 50% of the working-age population are not working positive hours, with many being protected by the Job Retention Scheme which requires them not to work. This scheme has proved crucial at limiting losses to household earnings, though perhaps at the cost of delaying sectoral reallocation (Barrero, Bloom and Davis, 2020). Nonetheless, 45% of individuals had experienced at least a 10% decline in household earnings by May. The key point the data shows is that despite being an aggregate shock, and despite the far reaching policies introduced, the impacts on individuals are highly heterogeneous.

The welfare costs of the economic shocks depend both on the size of these direct shocks and also on the resources and mechanisms households have to mitigate the shocks. We show that, to date, self-insurance through using savings and additional borrowing is highly prevalent; external support from new welfare applications and transfers from friends and family. However, the largest economic shocks have fallen on those least able to mitigate. Those most affected are individuals from minority ethnic groups, single parents and those in the lowest quintile of long-run income.

These initial waves of the *Understanding Society COVID-19 Survey* map out the evolution of the labour market shocks and their impact. The first impact of the crisis on employment has been largely mitigated by the Job Protection Scheme, but as this scheme ends, those who are notionally employed but not working any hours will move into unemployment or need to adjust.

References

- Adams-Prassl, Abi, Teodora Boneva, Marta Golin, and Christopher Rauh.** 2020. “Inequality in the Impact of the Coronavirus Shock: Evidence from Real Time Surveys.” Institute of Labor Economics (IZA) IZA Discussion Papers 13183.
- Alstadsæter, Annette, Bernt Bratsberg, Gaute Eielsen, Wojciech Kopczuk, Simen Markussen, Oddbjorn Raaum, and Knut Røed.** 2020. “The First Weeks of the Coronavirus Crisis: Who Got Hit, When and Why? Evidence from Norway.” National Bureau of Economic Research Working Paper 27131.
- Barrero, Jose Maria, Nicholas Bloom, and Steven J Davis.** 2020. “Covid-19 is also a reallocation shock.”
- Béland, Louis-Philippe, Abel Brodeur, and Taylor Wright.** 2020. “The short-term economic consequences of Covid-19: exposure to disease, remote work and government response.”
- Belot, Michèle, Syngjoo Choi, Julian C Jamison, Nicholas W Papageorge, Egon Tripodi, and Eline Van den Broek-Altenburg.** 2020. “Unequal Consequences of Covid 19 across Age and Income: Representative Evidence from Six Countries.”
- Benzeval, Michaela, Christopher R. Bollinger, Jonathan Burton, Thomas F. Crossley, and Peter Lynn.** 2020a. “The representativeness of Understanding Society.” , (2020-08).
- Benzeval, Michaela, Jon Burton, Thomas F. Crossley, Paul Fisher, Annette JÄckle, Hamish Low, and Brendan Read.** 2020b. “The Idiosyncratic Impact of an Aggregate Shock The Distributional Consequences of COVID-19.” , (911).
- Cortes, Guido Matias.** 2020. “Heterogeneous Labor Market Impacts During the Early Stages of the Covid-19 Pandemic.” Rimini Centre for Economic Analysis Working Paper series 20-13.
- Fisher, Paul, Laura Fumagalli, Nick Buck, and Silvia Avram.** 2019. “Understanding Society and its income data.” , (2019-08).
- Fitzgerald, John, Peter Gottschalk, and Robert Moffitt.** 1998. “An Analysis of Sample Attrition in Panel Data: The Michigan Panel Study of Income Dynamics.” *Journal of Human Resources*, 33(2): 251–299.
- Institute for Social and Economic Research.** 2020a. *Understanding Society COVID-19 Study, 2020. [data collection]. 1st Edition.* UK Data Service. SN: 8644, 10.5255/UKDA-SN-8644-1.
- Institute for Social and Economic Research.** 2020b. *Understanding Society COVID-19 User Guide. Version 1.0.* University of Essex, Colchester.
- Lynn, Peter, and Olena Kaminska.** 2010. “Weighting strategy for Understanding Society.” , (2010-05).
- Micklewright, John, and Sylke V Schnepf.** 2010. “How reliable are income data collected with a single question?” *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 173(2): 409–429.
- Moffit, Robert, John Fitzgerald, and Peter Gottschalk.** 1999. “Sample attrition in panel data: The role of selection on observables.” *Annales d'Economie et de Statistique*, 129–152.

- Montenovo, Laura, Xuan Jiang, Felipe Lozano Rojas, Ian M Schmutte, Kosali I Simon, Bruce A Weinberg, and Coady Wing.** 2020. "Determinants of Disparities in Covid-19 Job Losses." National Bureau of Economic Research Working Paper 27132.
- Nicoletti, Cheti, and Franco Peracchi.** 2005. "Survey response and survey characteristics: microlevel evidence from the European Community Household Panel." *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 168(4): 763–781.
- Schröder, Carsten, Theresa Margareta Entringer, Jan Goebel, Markus M Grabka, Daniel Graeber, Martin Kroh, Hannes Kröger, Simon Kühne, Stefan Liebig, Jürgen Schupp, et al.** 2020. "Covid-19 is not affecting all working people equally."
- University of Essex Institute for Social and Economic Research, NatCen Social Research, and Kantar Public.** n.d.. "Understanding Society: Waves 1-9, 2009-2018 and Harmonised BHPS: Waves 1-18, 1991-2009. [data collection]. 12th Edition."
- Valliant, Richard, and Jill Dever.** 2018. *Survey weights: a step-by-step guide to calculation*. Stata Press College Station, TX.
- von Gaudecker, Hans-Martin, Radost Holler, Lena Janys, Bettina Siflinger, and Christian Zimpelmann.** 2020. "Labour Supply in the Early Stages of the COVID-19 Pandemic: Empirical Evidence on Hours, Home Office, and Expectations." IZA Institute of Labor Economics Working Paper 13158.
- Wooldridge, Jeffrey M.** 2002. "Inverse probability weighted M-estimators for sample selection, attrition, and stratification." *Portuguese Economic Journal*, 1(2): 117–139.