HEDG

HEALTH, ECONOMETRICS AND DATA GROUP

THE UNIVERSITY of York

WP 21/14

The impact of Covid-19 vaccination for mental health

Kausik Chaudhuri and Peter Howley

August 2021

http://www.york.ac.uk/economics/postgrad/herc/hedg/wps/

The impact of Covid-19 vaccination for mental health

Kausik Chaudhuri [1] and Peter Howley [2]

Abstract

We examine the impact of vaccination against Covid-19 for mental health. Our estimates suggest that vaccination leads to a significant and substantive improvement in the mental health of those most at risk of hospitalisation and death from Covid-19. Our proposed explanation is that in the absence of vaccination, anxiety about contracting COVID-19 has a deleterious impact on the mental health of this cohort. On the other hand, our findings suggest that vaccination does not materially impact the mental health of those least at risk from Covid-19, namely younger cohorts. The lack of any significant impact for this cohort may explain vaccine hesitancy amongst young people. For this group, a lack of uptake may be principally due to a lack of perceived benefits for their own well-being as opposed to vaccine hesitancy.

Key words: covid-19; mental well-being; vaccination; propensity score matching

JEL Classification: I12, I31

Declaration of Interests: None to declare.

[1] Senior Lecturer in Economics, Leeds University Business School, University of Leeds, LS2 9JT, Leeds, United Kingdom. Email: K.Chaudhuri@lubs.leeds.ac.uk

[2] Professor of Economics and Behavioral Science, Leeds University Business School, University of Leeds, LS2 9JT, Leeds, United Kingdom. Email: p.howley@leeds.ac.uk.

The impact of covid-19 vaccination for mental health

Abstract: We examine the impact of vaccination against Covid-19 for mental health. Our estimates suggest that vaccination leads to a significant and substantive improvement in the mental health of those most at risk of hospitalisation and death from Covid-19. Our proposed explanation is that in the absence of vaccination, anxiety about contracting COVID-19 has a deleterious impact on the mental health of this cohort. On the other hand, our findings suggest that vaccination does not materially impact the mental health of those least at risk from Covid-19, namely younger cohorts. The lack of any significant impact for this cohort may explain vaccine hesitancy amongst young people. For this group, a lack of uptake may be principally due to a lack of perceived benefits for their own well-being as opposed to vaccine hesitancy.

Key words: covid-19; mental well-being; vaccination; propensity score matching

1. Introduction

There is an extensive literature documenting the economic impacts of the COVID-19 pandemic. To fully appreciate the consequences of this pandemic for human welfare we need to assess its impact on our mental health as well as economic well-being. With this in mind, a nascent literature is beginning to carefully detail the consequences of the pandemic for people's mental health. A consistent finding in the emerging literature on this topic is that the Covid-19 pandemic, particularly the social distancing restrictions, is associated with a substantive rise in psychological distress (see Banks et al., 2021 for a recent review of this literature). To date, the main approach to quantifying the mental health burden in this literature is to rely on surveys which ask individuals to report their mental well-being at various stages during the pandemic. A significant number of these studies have similar measures of well-being pre-pandemic allowing people to track changes in mental health over time.

While this research suggests that the impact of the pandemic for mental health can be substantive, these effects are not uniformly distributed. As an illustration, using data from Understanding Society, Daly et al. (2020) reported that 18–34 year-olds were disproportionately impacted, relative to other groups. Further studies have reported that women, ethnic minorities and those faced with financial insecurity are also disproportionately negatively impacted (Anaya et al. 2021; Banks & Xu, 2020:2021; Cheng et al., 2021; Daly et al., 2020; Etheridge & Spantig, 2020;

Giovanis & Ozdamar, 2020; Niedzwiedz et al., 2021; Proto & Quintana-Domeque, 2021; Shen & Bartram, 2021; Swaziek & Wozniak, 2020; Zamarro & Prados, 2021).

Some potential limitations with this existing literature is that even when data is available before the pandemic, these studies are not able to precisely identify an appropriate counterfactual. This can be important as, for example, some mental health measures have been trending downwards prior to the pandemic and therefore before and-after comparisons may overstate the mental health burden (Banks et al., 2021). Additionally, there are seasonal patterns to reported mental health measures, meaning that any reported difference in mental health pre-and post-pandemic may be partly confounded with seasonal trends. A further consideration when assessing the existing literature is the relevant interpretation. By comparing the self-reported well-being from individuals surveyed during certain time periods, the estimated disutility associated with the pandemic will be sensitive to the dates selected. The question here is when does the pandemic start/end and what constitutes the reference time period.

We offer an alternative way of quantifying the consequences of the pandemic for mental health. Specifically, we look to estimate the impact of receiving a Covid-19 vaccine for mental health as captured by the General Health Questionnaire (GHQ). Our central hypothesis is that vaccination will largely remove health related anxiety for vaccinated adults as they have been shown to be remarkably effective in reducing illness and death. Recent estimates, for instance, suggest an upwards of 90% reduction in the probability of hospitalisation and death once vaccinated (NHS 2021). While the pandemic may still negatively impact the mental health of people through other channels such as social distancing, health related anxiety (at least as applied to oneself) should we suggest be largely eliminated once vaccinated. Accordingly, for the 'treated', our estimates of the impact of vaccination for mental health can be seen as a reasonable lower-bound estimate of health-related anxiety associated with the pandemic itself. These estimates should supplement the emerging body of work quantifying the mental health impact of lockdowns and other less restrictive interventions such as social distancing guidelines aimed at minimising the spread of Covid-19.

In determining the impact of vaccination for mental health, we compare the mental health of individuals who have received the Covid-19 vaccine with the mental health of those who have not. Here we take advantage of wave 7 and 8 of Covid-19 surveys by the UK Household Longitudinal Study as these surveys record a measure of people's mental health as captured by the General Health Questionnaire, their vaccine status and also their willingness to take a vaccine. At the time these surveys were undertaken, about 43% of the sample population had taken at least one dose of the vaccine. We are able to take advantage of significant variation in the actual number as well as characteristics of people who have taken the vaccine in these surveys. One of the reasons for such variation is that there was significant spatial variation in the uptake of vaccines in the UK, mainly for logistical reasons. Simply put, some local authorities were in a better position to distribute vaccines than others. Second, the interview dates are randomised across each wave.

One potential problem with this direct comparison approach is that both groups labelled hereafter for ease of writing as treatment (taken at least one dose of vaccine) and control group (have not taken the vaccine) may differ in ways that matter when it comes to measuring the impact of the vaccine for mental health. For example, vaccines were distributed based on clinical need with clinical need being principally determined by age and the presence of underlying health conditions. Therefore, a raw estimate of the difference in mental health between both groups may confound the beneficial impact of the vaccine with other characteristics such as age and health.

Another relevant consideration is that while individuals could not choose to take a vaccine as availability was based on clinical need, they could *choose* not to take one. It is possible that individuals who choose not to take a vaccine may differ in relevant background characteristics which could also be related with mental health. This is however less likely to be an issue in the UK than elsewhere as general acceptance of the vaccine is high as compared to similar European countries. In short, the unbalanced distributions of characteristics between the two groups may create a problem with selection bias and, in turn, a biased estimate of the average treatment effect for the treated.

Our approach to dealing with this potential bias is to employ various matching techniques, such as covariate matching (CVM), propensity score matching (PSM) and entropy balancing. The basic idea behind matching approaches is to find in a large group of individuals who are not exposed to the treatment, those individuals who are most similar to treated units in all relevant pre-treatment characteristics. In practice it involves reweighting or simply discarding units in order to ensure that the treatment and control group have similar distributions of characteristics so that the treatment variable becomes closer to being independent of the background characteristics (Hainmueller 2012). By helping to ensure that both groups have equal distributions of characteristics much like what would be observed under random assignment, matching can lead to more reliable estimates of 'treatment' effects in the presence of non-random assignment. In this case, after matching, those in the control group have a similar probability of being offered the vaccine as those observed in the treatment group. This is facilitated by matching on the basis of observed characteristics such as age and health status, factors we know will predict the likelihood of being offered a vaccine.

Fortunately, in our dataset, we also have a variable which asks individuals if they are willing to take a vaccine. This means we can match individuals not only when it comes to their individual characteristics which predict the likelihood of them being offered a vaccine but also their willingness to take one.[3] In such a scenario we argue that the differences between both groups after matching can be seen as an unbiased estimate of the average treatment effect (ATT) for the treated (people who have taken the vaccine) because the treatment group does not on average differ systematically from our matched control group consisting of those who have not taken the vaccine.

[3] In our sample, around 7.5% of the respondents were unwilling to take the vaccine. For unknown reasons 5% of the vaccinated respondents indicated that they were not willing to take vaccine. Our results remain qualitatively the same if we drop these observations.

Our main findings suggest that vaccination can lead to a significant and substantive improvement in mental health. This is principally evident, however, for those most at health risk from Covid-19, namely older demographics and/or those with underlying health conditions. In contrast to older demographics, the benefits of vaccination for younger cohorts in terms of their own mental well-being are negligible. This we posit may be an important factor underpinning comparatively lower rates of vaccine uptake amongst young adults. For this demographic, vaccine hesitancy could be the result of a rational calculation of what they perceive as the benefits for their own well-being versus the actual costs (e.g. time and risks however small of serious side effects) of vaccination.

2. Dataset and key variables

We employ data from the UK Household Longitudinal Study (UKHLS) also known as Understanding Society. The UKHLS is a household panel that captures, among other things, information from adults about their economic and social circumstances, socio-demographics and mental health. Our key outcome variable of mental health contained in this survey dataset is the 12-item version of the General Health Questionnaire (GHQ-12). This 12-item scale is designed to assess somatic symptoms, anxiety and insomnia, social dysfunction, and general happiness (see Table A1 in appendix for a list of all question items). It is possibly the most commonly used measure of subjective well-being in the literature (Jackson 2007). The GHQ offers an advantage over single question measures of subjective well-being, such as happiness and life satisfaction as it is based on responses to 12 separate questions. Each of the 12 items is scored on a four-point scale. The overall GHQ score can take values from 0 to 36, with 36 representing the lowest level of psychological well-being. For ease of interpretation, we reversed the overall score so that a value of 36 represents the highest level and going forward we refer to this variable simply as mental health or mental well-being.

Beginning in April 2020, participants of the UKHLS were asked to complete a short online survey on the impact of the COVID-19 pandemic. The survey consists of all eligible consented individuals aged 16 years and over in eligible households. The survey was undertaken monthly between April-July 2020 and then every two months thereafter. Beginning in Wave 7 of these monthly surveys, respondents were questioned in relation to whether they had taken a vaccine for covid-19 as well as their willingness to take one in the absence of vaccination. In the analysis that follows we focus on wave 7 and 8[4] of these special Covid-19 surveys given the availability of information pertaining to vaccine status, vaccine willingness and whether people were classified as clinically vulnerable to Covid-19. We discuss the structure of these variables in the following section.

Key variables

The National Health Service (NHS) in the UK created a two-level priority list classifying those most at risk from serious illness if they are infected by COVID-19. People in the first list were deemed clinically extremely vulnerable[5] to Covid-19 and individuals belonging to this category included, amongst others, those with serious health conditions such as receiving cancer treatment, those who have a lung condition such as cystic fibrosis, severe asthma or COPD, or a compromised immune system. The second 'priority' list for vaccination consisted of individuals deemed at high risk, at least relative to the rest of the population. This included care home residents and care home workers, frontline health and social care workers, those with underlying health conditions that put them at greater risk (e.g. diabetes) and people aged 50 and over. We used a binary indicator to indicate whether individuals belong to one of these priority lists and thus at comparatively higher risk from hospitalisation and death from Covid-19. We label this variable as 'clinically vulnerable'.

As noted above, health and social care workers were placed in a priority list for vaccination. Other 'key workers' were not placed in the priority list but the UKHLS does capture whether individuals are classified as key workers during the pandemic. This included people working in

[4] Wave 7 began late January 2021 and wave 8 began in March 2021.
[5] More information on the characteristics of individuals belonging to this group is contained here https://www.nhs.uk/conditions/coronavirus-covid-19/people-at-higher-risk/whos-at-higher-risk-from-coronavirus).

Education and childcare, Key public services, Local and national government, Food and other necessary goods, Public safety and national security, Transport, Utilities, Communications and financial services. We derived a binary indicator classified whether individuals work in one of these employment categories which we label as key worker. Finally, the survey asked assessed vaccine hesitancy through the following question: “Imagine that a vaccine against COVID-19 was available for anyone who wanted it. How likely or unlikely would you be to take the vaccine?” The respondent answer options are “very likely, likely, unlikely and very unlikely”. Using this survey measure we are able to classify unvaccinated people according to their willingness to take one, and therefore importantly match people according to vaccine hesitancy.

3. Matching

Table 1 shows that prior to matching our treatment (vaccinated) and control group (not vaccinated) are dissimilar on a variety of variables that are predictive of vaccine status and also potentially mental health. This includes age, presence of underlying health conditions and their willingness to take a vaccine. To ensure balance between our treatment and a control group we implement a variety of matching approaches. The first is known as covariate matching which simply looks to match individuals based on the similarity of their covariates. In practice, it is impractical to match directly on covariates because of the curse of dimensionality and so typically with covariate matching, distance measures like the Mahalanobis distance are employed to calculate the similarity of two individuals in terms of their covariate values and the matching, in turn, is done on these distances (Zhao 2004). With this approach, non-treated and treated groups become only randomly different based on covariates so that the outcomes of the matched non-treated and treated groups, which keep the originally observed values, are comparable under the matched covariate condition.

In addition to covariate matching, we also employ propensity score matching. With this approach we estimate a propensity score for each individual in the dataset using a logit model where a dummy variable capturing whether individuals have been vaccinated is our outcome variable. These estimates are presented in Table A2 in the online appendix.[6] The propensity score is the probability of treatment assignment (i.e. being assigned a vaccine) conditional on observed baseline covariates (Rosenbaum and Rubin 1983). It is a balancing score which means that conditional on the propensity score, the distribution of characteristics across treated and untreated subjects should be similar (see Austin 2011b for further discussion on this point). What this means in practice is that, similarly to covariate matching, once we match individuals on the basis of their estimated propensity score, the distribution of observed baseline characteristics (our covariates) should be the same across our treated units (those who have received a vaccine) and our control units (those who have not received a vaccine).

The covariates we selected for our matching models include age and a dummy variable indicating whether an individual is defined as clinically vulnerable as these variables are the strongest predictors of whether individuals would have been offered a vaccine at the time these surveys were carried out and these variables may also be correlated with mental health. To that, we add a variable which indicates how responsive individuals are to receiving a vaccine and whether they were classified as a key worker. We also include a regional identifier as there was some regional heterogeneity in the distribution of vaccines and a wave dummy. Finally, we include some socio-demographic characteristics such as gender, relationship status (whether respondents have a partner or not) and whether they were born in UK. The results were qualitatively similar whether we included or excluded these sociodemographic controls.

Once we have a propensity score for each individual, the next step is to match sets of treated and untreated individuals. There are a plethora of techniques that are available to the researcher when conducting this process and there is as yet no clear guidelines available as to which one should be most preferred (e.g. see Baser 2006 for a discussion of this issue). The same is true when it comes

[6] We also report p-value for the balancing test with regard to the propensity score (p-value is 0.091) implying that the matching procedure implemented by kernel matching is balanced. In Figure A1, we also show that after matching, the density of the estimated propensity score is remarkably similar for the treated and the control group.

to which distance metrics to employ in order to facilitate matching on covariates. Our response was not only test how sensitive our results were to different matching techniques (i.e. covariate versus propensity score matching) but also to test how sensitive our results were to different distance metrics and matching algorithms. The specific matching techniques we applied were the following: *Nearest Neighbour*, *Mahalanobis-distance, Kernel and Radius*. The first two can be seen as covariate matching based on an estimated distance metric and the last two as propensity score matching based on an estimated propensity score.

While covariate and propensity score matching are the two most common matching techniques, entropy balancing is another approach that has emerged quite recently as an alternative (see Hainmueller 2012). Entropy balancing is a data preprocessing procedure that looks to directly incorporate covariate balance in the estimation by reweighting a dataset. The preprocessing is based on a maximum entropy reweighting scheme that assigns weights to each unit such that the covariate distributions in the reweighted data satisfy a set of moment conditions specified by the researcher. We reweight the mean from the control group to match the target mean from the treatment group such that the reweighted data can be used to analyze the average treatment effect on the treated.

With nearest neighbour we simply match each individual in the treatment group (hereafter referred to as treated unit) with an individual in the control group (hereafter referred to as control unit) based on the closest distance between their estimated Mahalanobis-distance metric (defined later) *with* replacement. Matching with replacement keeps bias low at the cost of larger variance. Bias is lower when employing matching with, as opposed to without replacement, as it means particularly suited control units can be used more than once (with replacement) as matches for treated units. Such a strategy in turn ensures a better balance between our treatment and control group as we are not forcing a match between our treated unit and what could be a dissimilar control unit. Next we employed Mahalanobis-distance kernel matching where matching is based on a distance metric that measures the proximity between observations in a vector of explanatory variables (X) where the distance matrix (MD) is defined as $MD(X_i, X_j) = \sqrt{(X_i - X_j)' \Sigma^{-1} (X_i - X_j)}$ where Σ is the covariance matrix of X.[7]

Thirdly, we employed radius matching which is a type of propensity score matching approach. With radius matching, each treated unit is matched only with a control unit whose propensity score falls in a predefined range of the propensity score of the treated unit. Here we adopt a predetermined caliper (bandwidth) of 0.2*standard deviation of logit of the estimated propensity score (see Austin 2011a for further details relating to the choice of appropriate bandwidth). Our second propensity score matching approach was kernel matching with Epanechnikov kernel and data dependent automatic bandwidth selection.[8] This weighting ensures that closer matches gets more weight in the matching process and has the advantage of allowing more flexibility in the accepted matches. Finally we employed entropy balancing (with means of all the covariates to be balanced) where the balancing uses a reweighting scheme directly incorporating covariate balance into the weight function applied to the sample units.[9] Entropy balancing (Hainmueller 2012) identifies weights for the control sample to equalize the distribution of determinants across treatment and control samples. All estimated standard errors are bootstrapped with 100 replications.

[7] We use Epanechnikov kernel with an automatic data dependent bandwidth which is estimated as 0.581 by cross-validation with respect to the means of the covariates.

[8] Bandwidth is estimated to be equal to 0.010 by cross-validation with respect to the mean of the propensity score.

[9] If we use both means and variances of all the covariates to be balanced, the results reported in the paper remains qualitatively unchanged.

Table 1: Descriptive Statistics

Variable	Obs	Mean	Std. Dev.	Obs	Mean	Std. Dev.	p-value for t-test
	Did not get Vaccine			Did get Vaccine (either one or both)			
GHQ-12	12,423	23.14	6.15	9,562	24.04	5.64	0.00
Age	12,423	49.19	15.64	9,562	61.67	13.84	0.00
Born in UK	12,423	0.87	0.34	9,562	0.90	0.30	0.00
Clinically vulnerable Dummy	12,423	0.34	0.47	9,562	0.56	0.50	0.00
Male	12,423	0.42	0.49	9,562	0.41	0.49	0.20
Key worker	12,423	0.25	0.43	9,562	0.25	0.43	0.62
Couple	12,423	0.69	0.46	9,562	0.73	0.44	0.00
Vaccine willingness	12,423	0.91	0.29	9,562	0.95	0.22	0.00

Note: p-value for t-test reports whether the mean/proportion of the variables in the two groups are equal.

4. Matching results

Our main results are presented in Table 2. In columns 1 -5, we present our matching estimates using the different matching approaches described above. Looking first at column 1, we can see that the estimated mental health benefit for this sample from being treated (receiving the vaccine) comes to 0.938 ($p < .01$) units in our GHQ measure when using nearest neighbour matching. In column 2, we can see that we obtain similar estimates when using Mahalanobis distance matching as the estimated treatment effect for the treated is 0.810 ($p < 0.01$). In columns 3 and 4 we can see that both our propensity score matching approaches, namely radius and kernel predict similar treatment effects (0.94 & 1.07, $p < 0.01$). Finally, we observe that the estimated treatment effect with entropy balancing comes to 1.811 ($p < 0.05$). In sum, our matching estimates suggests that for this group of vaccinated adults, the impact of vaccination is positive with the estimated impact ranging between a 0.8 and 1.8 unit increase in mental well-being as captured by the GHQ. Most of the estimates however converge between 0.8 and 1 units with the estimates when using entropy matching somewhat of an outlier. We adopt a conservative approach in what follows by focusing our attention on the estimates obtained using the other matching approaches.

Irrespective of which matching approach is most preferred, a pertinent question is how large are these estimated effects? That is, all the estimates are statistically significant but does vaccination generate a substantive improvement in mental health and if so how substantive? As a way to illustrate these impacts we can compare these estimated effect sizes with the estimated impacts of other major life events in the wider ‘economics of happiness’ literature. Fortunately, as the GHQ is a widely used outcome measure, there are a wide array of studies quantifying the impact of various life events for mental health as captured by this outcome variable. An estimated mental health gain of between 0.8 to 1 units in our GHQ measure would, for example, be approximately one half to two thirds of the estimated disutility associated with unemployment (see Howley and Knight 2021; Flint et al. 2013). Unemployment alongside disability are perhaps the two factors noted in the ‘economics of happiness’ literature as having the most substantive and long-lasting negative impacts for mental health. Other commonly identified negative correlates include divorce and widowhood and our estimates would suggest that the estimated benefit of vaccination for our treated group far exceeds the estimated negative mental health impacts associated with either of these life events. Clearly then for this group of vaccinated individuals, our estimates would suggest that vaccination led to a substantive as well as statistically significant improvement in mental health.

It is important to note that these estimates relate to the treatment group which in our case is just under half of the sample population. This half of the population are also those most likely to be at health risk from Covid-19 given that distribution of the vaccine was based on risk of hospitalisation and death. Earlier we posited that the vaccine will be beneficial for mental health as it will reduce health related anxiety. If this holds true then we would expect that the estimated impact of vaccination to be greater for those most at risk from hospitalisation and death. As a way to test this hypothesis, we divided our sample of treated individuals into sub-groups based on age and the presence of underlying health conditions. These results can also be seen in Table 2. With age, we simply divided our sample into two groups, namely those above and below the median age in our survey which is 56. Understandably the estimates are less precisely estimated but we can see that the mental health benefits associated with vaccination is largely concentrated on the group who are 56 and over. When looking at those under the age of 56 the estimated impacts are comparatively small and not statistically significant across most matching estimates. In a similar vein, we can also see that the mental health benefits are much more pronounced for those with underlying health conditions.

As a final check we estimated the mental health benefit for those who were both 56 and over and/or belong to the clinically vulnerable sub-group. These estimates can be seen in the last two rows in Table 2. What is perhaps notable here is that when looking at the sub-group of individuals who are both under the age of 56 and not classified as clinically vulnerable, none of the estimates were statistically significant. Overall we can see that the mental health benefits from vaccination can be substantive and match if not exceed that of many major life events, but principally for only those who are deemed at significant health risk, namely older and/or clinically vulnerable sub-groups. For younger groups without underlying health conditions, any positive impact associated with vaccination is negligible.

Table 2: Impact of Vaccination on mental wellbeing (either one or both doses)

	Distance Metric		*Propensity Score*		
	Nearest Neighbour	Mahalanobis -distance kernel matching	*Radius*	*Kernel*	Entropy
Whole sample, N = 21985	0.938*** (0.322)	0.810*** (0.213)	0.940*** (0.267)	1.075*** (0.401)	1.811** (0.744)
Age equal to or higher than Median age, N = 11502	0.821** (0.352)	1.169*** (0.422)	0.804* (0.466)	1.632* (0.831)	3.115*** (0.877)
Age less than Median age, N = 10483	0.288 (0.231)	0.386* (0.226)	0.165 (0.197)	0.188 (0.186)	0.079 (0.167)
Clinically Vulnerable, N = 9644	1.367*** (0.520)	1.562*** (0.385)	1.518*** (0.450)	1.817*** (0.712)	3.065*** (1.054)
Clinically Not Vulnerable, N = 12341	0.370 (0.288)	0.293 (0.199)	0.773*** (0.236)	0.864*** (0.258)	0.250 (0.196)
Clinically Vulnerable and age equal to higher than median age, N = 6928	1.035** (0.516)	2.058*** (0.577)	0.736* (0.432)	1.699* (0.944)	3.720*** (0.663)
Clinically Not Vulnerable and age lower than median age, N = 2716	0.168 (0.261)	0.500 (0.370)	0.062 (0.382)	0.201 (0.461)	0.097 (0.388)

Note: Standard errors are calculated using 100 bootstraps. We present the results for Average treatment effect on treated (ATET). *** denotes significance at 1%,** denotes significance at 5%,* denotes significance at 10% level.

4. Conclusion

The objective of this study was to estimate the benefits of being vaccinated for Covid-19 for mental health. One potential problem with estimating the impact of vaccination is that vaccines were not distributed randomly. We sought to minimise this potential source of bias by employing a range of matching techniques. With matching, we look to ensure that the distribution of characteristics across treated (have taken a vaccine) and control units are equivalent so that we can more reasonably infer that any differences in mental well-being between them is due to vaccination. Given the unique features of the survey at our disposal we are able to match individuals both when it comes to their probability of being offered a vaccine as well as their level of vaccine hesitancy. While this does not completely eliminate all sources of bias, coupled with the consistency of our estimates across a wide array of matching approaches it does point towards the robustness of our results.

There is an extensive literature developing concerned with estimating the burden of lockdowns and other social restrictions for mental health. This study can be seen as an initial first step in quantifying the mental health impact of the other major weapon used by governments to protect public health, namely investment in the development of vaccines. By focusing on the impact of vaccination, our estimates can also be seen as capturing the impact of health related anxiety for people's well-being. Our intuition is that given the success of Covid-19 vaccines when it comes to reducing the risk of hospitalisation and death, health related anxiety should we suggest be largely eliminated once vaccinated. As such, our estimates can be seen as lower bound estimates of the psychological distress associated with people's concern for their health during the pandemic.

We find that the mental health benefit from vaccination is significant for those most at risk of hospitalisation and death, namely the over 55's and/or those with underlying health conditions. Apart from being statistically significant, the estimated gain in mental health is substantive and in keeping with major life events. This, in turn, would suggest that prior to vaccination, general anxiety about becoming ill with covid-19 was a significant and substantive source of psychological distress for many people. On the other hand, we find no evidence to suggest that vaccination materially impacts the mental health of those who are at low risk of hospitalisation and death, namely younger cohorts without any underlying health conditions. It is perhaps notable here that young adults are the group most likely not to seek vaccination. In England for example, which has had a successful vaccination drive overall with, at the time of writing, three-quarters of the population having received both doses, vaccination has stalled in recent times, particularly amongst younger cohorts. While undoubtedly a factor, vaccine hesitancy does not seem to be the main driving force as recent estimates by the Office for National Statistics suggest that less than 10% of 16-29 year olds report that they are hesitant about getting vaccinated (ONS 2021). As we have shown, there are limited if any benefits to this group personally in terms of own mental well-being from vaccination. Therefore, for this group, this decision to not get vaccinated could be due simply to a rational evaluation of the costs and perceived benefits of vaccination as opposed to vaccine hesitancy per say.

The lack of any personal benefit coupled with positive externalities in terms of herd immunity does suggest a role for policy intervention to encourage vaccination amongst this cohort. One approach could be financial incentives which have been shown to be effective in encouraging vaccination for the flu among young people (Bronchetti et al. 2015) as well as to access health services (Onza 2021). Providing financial incentives do not change beliefs, rather change the perceived benefit but as we noted earlier the majority of young people who are unvaccinated do not express hesitancy about vaccination and as such measures which increase perceived benefits may be all that is needed for much of this cohort. Some state and local governments particularly in the US have already started to employ financial incentives. One such example is Philadelphia's vaccination lottery which has a grand prize of $50,000.

While financial incentives have been shown previously to be effective in encouraging health behaviours such as flu vaccination, what is less certain is the appropriate size and nature of any reward (e.g. cash, vouchers, contribution to charity etc.). This could be an important consideration as financial incentives in certain instances can crowd out people's intrinsic motivation to get vaccinated (Frey and Jegen 2001). In addition to increasing the perceived benefit, alleviating the barriers to vaccination may also be an effective strategy towards encouraging vaccination amongst young adults. This could involve setting up vaccination centres in locations that are convenient for young people such as schools, universities or where large numbers are congregated such as music festivals.

References

Anaya, L. Howley, P., Waqas, M and Yalonetzky, G. (2021). Locked down in Distress: a Causal Estimation of the Mental-Health Fallout from the COVID-19 Pandemic in the UK (June 25, 2021). Available at SSRN: https://ssrn.com/abstract=3873973 or http://dx.doi.org/10.2139/ssrn.3873973

Austin P.C. (2011). Optimal caliper widths for propensity-score matching when estimating differences in means and differences in proportions in observational studies. Pharmaceutical Statistics. 10(2):150-61. doi: 10.1002/pst.433.

Austin P.C. (2011). An introduction to propensity-score methods for reducing the effects of confounding in observational studies. *Multivariate Behavioral Research*, 399-424.

Banks, J., Fancourt, D., & Xu, X. (2021). *Mental health and the COVID-19 pandemic*. Institute for Fiscal Studies. Available at: https://ifs.org.uk/publications/15368

Banks, J., & Xu, X. (2020). The Mental Health Effects of the First Two Months of Lockdown during the COVID-19 Pandemic in the UK. *Fiscal Studies*, *41*(3), 685–708. https://doi.org/10.1111/1475-5890.12239

Baser O. (2006).Too much ado about propensity score models? Comparing methods of propensity score matching. *Value Health,* 9:377–85.

Bronchetti, E. T., Huffman, D. B., and Magenheim, E. (2015). Attention, intentions, and follow-through in preventive health behavior: Field experimental evidence on flu vaccination. *Journal of Economic Behavior & Organization*, 116, 270–291

Cheng, Z., Mendolia, S., Paloyo, A. R., Savage, D. A., & Tani, M. (2021). Working parents, financial insecurity, and childcare: mental health in the time of COVID-19 in the UK. *Review of Economics of the Household*, *19*(1), 123–144. https://doi.org/10.1007/s11150-020-09538-3.

Daly, M., Sutin, A., & Robinson, E. (2020). Longitudinal changes in mental health and the COVID-19 pandemic: Evidence from the UK Household Longitudinal Study. *Psychological Medicine*. https://doi.org/10.1017/S0033291720004432

Etheridge, B., & Spantig, L. (2020). The gender gap in mental well-being during the Covid-19 outbreak: evidence from the UK. *ISER Working Paper Series*. https://www.iser.essex.ac.uk/research/publications/working-papers/iser/2020-08.

Flint E, Shelton N, Bartley M, et al. (2013) Do local unemployment rates modify the effect of individual labour market status on psychological distress? *Health & Place* 23: 1–8.

Frey, B.S., and Jegen, R. (2001). Motivation crowding out theory. *Journal of Economic Surveys*, 15(5), 589–611.

Giovanis, E., & Ozdamar, O. (2020). *Income Losses and Subjective Well-Being: Gender and Ethnic Inequalities during the Covid-19 Lockdown Period in the UK*. https://doi.org/10.21203/rs.3.rs-132600/v1.

Hainmueller, J. (2012). Entropy Balancing for Causal Effects: A Multivariate Reweighting Method to Produce Balanced Samples in Observational Studies. Political Analysis 20(1), 25-46. DOI: 10.1093/pan/mpr025

Howley, P. and Knight, S. (2021) Staying down with the Joneses: Neighbourhood differences in the well-being effects of unemployment. *Work, Employment and Society*, 1-21.

Jackson, C. (2007). The General Health Questionnaire. *Occupational Medicine*, *57*(1), 79–79. https://doi.org/10.1093/occmed/kql169

Niedzwiedz, C. L., Green, M. J., Benzeval, M., Campbell, D., Craig, P., Demou, E., Leyland, A., Pearce, A., Thomson, R., Whitley, E., & Katikireddi, S. V. (2021). Mental health and health behaviours before and during the initial phase of the COVID-19 lockdown: Longitudinal analyses of the UK Household Longitudinal Study. *Journal of Epidemiology and Community Health*, *75*(3), 224–231. https://doi.org/10.1136/jech-2020-215060.

Office for National Statistics (2021) Coronavirus and vaccine hesitancy, Great Britain. Available at: https://www.ons.gov.uk/peoplepopulationandcommunity/healthandsocialcare/healthandwellbeing/bulletins/coronavirusandvaccinehesitancygreatbritain/9august2021.

Onza, A. (2021). Studies probe how payouts affect U.S. vaccination rates. Science 373 (6555), 611.

Proto, E., & Quintana-Domeque, C. (2021). COVID-19 and mental health deterioration by ethnicity and gender in the UK. *PLOS ONE*, *16*(1), e0244419. https://doi.org/10.1371/journal.pone.0244419.

Rosenbaum P.R., Rubin D.B. (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika* 70:41–55

Shen, J., & Bartram, D. (2021). Fare differently, feel differently: mental well-being of UK-born and foreign-born working men during the COVID-19 pandemic. *European Societies*, *23*(sup1), S370–S383. https://doi.org/10.1080/14616696.2020.1826557

Swaziek, Z., & Wozniak, A. (2020). Disparities Old and New in US Mental Health during the COVID-19 Pandemic*. *Fiscal Studies*, *41*(3), 709–732. https://doi.org/10.1111/1475-5890.12244.

University of Essex Institute for Social and Economic Research. Understanding Society: COVID-19 Study, 2020-2021. 2021. doi:http://doi.org/10.5255/UKDA-SN-8644-9.

University of Essex Institute for Social and Economic Research. Understanding Society COVID-19: User Guide. 2021.

Zamarro, G., & Prados, M. J. (2021). Gender differences in couples' division of childcare, work and mental health during COVID-19. *Review of Economics of the Household*, *19*(1), 11–40. https://doi.org/10.1007/s11150-020-09534-7

Zhao, Z (2004). Using Matching to Estimate Treatment Effects: Data Requirements, Matching Metrics, and Monte Carlo Evidence, *The Review of Economics and Statistics*, 86(1), 91–107.

Online Appendix

Table A1: Statements that make up the GHQ

1. Have you recently been able to concentrate on whatever you're doing?
2. Have you recently lost much sleep over worry?
3. Have you recently felt that you were playing a useful part in things?
4. Have you recently felt capable of making decisions about things?
5. Have you recently felt constantly under strain?
6. Have you recently felt you couldn't overcome your difficulties?
7. Have you recently been able to enjoy your normal day-to-day activities?
8. Have you recently been able to face up to problems?
9. Have you recently been feeling unhappy or depressed?
10. Have you recently been losing confidence in yourself?
11. Have you recently been thinking of yourself as a worthless person?
12. Have you recently been feeling reasonably happy, all things considered?

Table A2: Logit Estimates for having vaccination

Variable	Coefficient
Age	0.098*** (0.002)
Born in UK dummy	0.011 (0.064)
Clinically Vulnerable Dummy	0.598*** (0.043)
Male Dummy	-0.390*** (0.040)
Couple dummy	0.054 (0.044)
Key worker dummy	1.103*** (0.048)
Vaccine Willingness	0.619*** (0.080)
No. of obs.	21985
Psuedo-R2	0.434
p-value for balancing test	0.094

Note: Results are for the pooled model for propensity score estimation with kernel matching and common support. Region and wave dummies were also included but left unreported for parsimony. p-value for balancing test indicates that with regard to the propensity score, the matching procedure implemented by kernel matching is balanced, so we can trust matching results.

Figure A1: Kernel Density Plot of Propensity Score before and after matching with kernel and common support