Improvement Analytics Unit
November 2021

COVID Oximetry @ home (CO@h): Feasibility analysis for evaluation of CO@h using regression discontinuity

Will Parry, Therese Lloyd

The Health Foundation

Key points

- COVID Oximetry @home (CO@h) is an England-wide programme to remotely monitor oxygen saturation levels of people diagnosed with COVID-19 and at risk of health deterioration, with the aim of improving patient outcomes: in particular, to avoid invasive ventilation, ICU admission and death.
- The CO@h standard operating procedure provided recommended guidance that all patients aged 65 and over, or patients who are classed as clinically extremely vulnerable (CEV), should be onboarded onto the service.
- A regression discontinuity design (RDD) is a robust quasi-experimental approach to causal inference and is appropriate in situations where eligibility for an intervention changes sharply at a predefined threshold of a 'running' variable (in this case, age). It takes advantage of the change in eligibility at the age cut-off by comparing the outcomes of patients who tested positive for COVID-19 just above and below the 65-year cut-off.
- This report documents the findings of the Improvement Analytics Unit (IAU)'s feasibility analysis to determine whether a regression discontinuity design would be an appropriate method for evaluating the causal impact of the CO@h programme.
- We found that the number and proportion of eligible patients that were onboarded onto the service was low, and that there was very little discontinuity (difference) in onboarding rates above and below the age threshold. For example, approximately 3.5% of eligible patients aged 65–69 were onboarded from January 2021, whereas approximately 2% of patients aged 60–64 were.
- The low rates of onboarding and the lack of discontinuity at the age cut-off would result in a lack of statistical power, virtually guaranteeing a null result, even if the intervention was successful in reducing emergency hospital outcomes. Furthermore, the lack of a discontinuity of the magnitude generally required for a high quality RDD would lead one to question the validity of the assumptions of this method relating to non-compliance.
- Therefore, we concluded that an RDD was not suitable for the evaluation of the CO@h programme in the current setting.
- A number of approaches to evaluate the CO@h programme are being pursued by evaluation partners, including a population-level analysis using generalised synthetic control methods by colleagues within the IAU. Each approach has different assumptions, strengths and weaknesses, but will together aim to provide reliable conclusions about the effect of CO@h on outcomes.

Introduction

COVID Oximetry @home (CO@h) is an England-wide programme to remotely monitor oxygen saturation levels of people diagnosed with COVID-19 and at risk of health deterioration. It aims to escalate cases of health deterioration earlier to avoid invasive ventilation, ICU admission and death. Following a pilot starting in March 2020, the scheme was rolled out nationally from November 2020, during the second wave of COVID-19 in England. The standard operating procedure for CO@h, published in November 2020, provided recommended guidance that all patients aged 65 and over or who are classed as clinically extremely vulnerable (CEV) should be onboarded onto the service.

The CO@h programme commissioned three different evaluation teams to evaluate the effect of this programme on outcomes ranging from patient experience to rates of ICU admission and death, using a range of different quantitative and qualitative methods.

The IAU planned two different quantitative approaches to estimating the causal impact of this programme, one of which was a local randomisation approach to Regression Discontinuity Design (RDD).[1,2] This analysis takes advantage of the change in eligibility at age 65 by comparing outcomes of patients who tested positive for COVID-19 just above and below the 65-year threshold (cut-off). This is a robust quasi-experimental approach to causal inference and is appropriate in situations where eligibility for an intervention changes sharply at a predefined threshold of a 'running' variable (in this case, age). This design largely avoids problems of observed or unobserved confounding, as the expected outcomes of patients just below and above the cut-off are assumed to be equivalent in the absence of the intervention.

This approach, however, requires certain assumptions/criteria to be met. Therefore, the first stage of the study was a feasibility analysis to determine whether this study design could be implemented in a robust and reliable way.

Further details of the planned analysis, including background, study cohort and methodology, can be found in the statistical analysis protocol.[3]

The feasibility analyses were performed in July 2021 using pseudonymised person-level data received from NHS Digital on 28 May 2021, which included the available data on patients who tested positive for COVID-19 between 1 October 2020 and 30 April 2021, on COVID-19 testing, patient characteristics, deaths and onboarding onto CO@h.

We also had access to information collected from each CCG documenting conversations with CCGs as to whether they had submitted 'complete' onboarding data by 27 May 2021. However, those CCGs flagged as complete, had indicated that they had submitted all the data they could access, rather than the data being complete in terms of all onboarded patients being recorded. Particularly in the early stages of implementation, not all patients onboarded may have been captured.

As not all CCGs had submitted complete onboarding data by 28 May 2021, a further and final set of data was made available to evaluation partners on 23 July 2021, but this was not available at the time when a decision on the feasibility of the study needed to be taken. However, preliminary analysis of this final set of data showed that despite there being approximately 25% more records of onboarded patients with a first COVID-19 positive test in the age range 60–69, this would not make any substantive difference to the results of the analysis.

Aims of the feasibility analysis

The analysis presented here aimed to identify whether an evaluation of CO@h was feasible using an RDD. The original objectives were to:

- determine the extent of any discontinuity in onboarding at age 65, the age cut-off specified in the standard operating procedure for CO@h
- determine whether there is sufficient statistical power for any evaluation to be able to detect the effect of CO@h on outcomes, given the discontinuity in onboarding.

Additionally, as the CO@h programme team and other evaluation partners identified differences in uptake, roll-out, eligibility criteria and completeness of data submissions between CCGs, we also investigated:

- whether there was a discontinuity at age 50 and the numbers of patients onboarded either side of this cut-off, as some areas were using this as their threshold for onboarding
- whether there was a discontinuity later in the programme (from January 2021 onwards), as some time may have been required to establish the programme properly, and data recording and submission was considered to have improved over time
- whether there was a discontinuity for CCGs recorded as having submitted 'complete' onboarding data (ie having submitted all onboarding data that they could access) according to roll-out monitoring information
- discontinuities and numbers of patients onboarded for each CCG individually, to determine whether there was a subset of CCGs that more closely followed the CO@h standard operating procedure guidelines.

Eligibility of patients for inclusion

Figure 1, below, shows a flow chart of the number of patients that fulfilled the study eligibility rules[3] and the number of eligible patients that were available for the different feasibility analyses conducted. At the outset, there were 20,750 patients in total who were recorded as onboarded in the period from 1 October 2020 to 30 April 2021. Considering the number of oximeters it is thought were purchased by NHS England and NHS Improvement for the programme[4] (200,000), and the number of first COVID-19 positive tests (lateral flow and/or PCR) for patients aged 65 and over during the period[5] (approximately 400,000), this was a much lower number than anticipated. After restricting to patients with a first COVID-positive test record, this fell to 16,832. After removing patients who were CEV or living in care homes, this fell to 12,996.

Restricting the data set to exclude age ranges >5 years either side of the cut-off (which would likely not meet the criteria for a local randomisation RDD) and removing a small number of patients whose ages could not be exactly determined around the cut-off (due to the weekly precision of their dates of birth), resulted in

2,904 onboarded patients aged between 60 and 69 (ie one day below 70) years old. Where the sample was restricted further to only include patients that tested positive from January 2021 onwards, the number of patients onboarded fell to 2,439. Numbers were similar for the group aged 45–54 (ie one day below 55) years old (Figure 1).

The number of patients onboarded was very low compared to the numbers of patients testing positive for COVID-19. For those aged between 60-69 years that were not CEV or living in care homes, across the whole period, there were 193,111 patients testing positive, compared to only 2,904 onboarded, for example. It is striking that restricting the data to patients testing positive from January 2021 onwards resulted in a large drop in the number of patients that tested positive in the data set (to 91,207 for those aged between 60-69 years). This suggests that although a considerable proportion of onboarded patients aged 60–69 (that were not CEV or living in care homes) were onboarded from January 2021 onwards, far fewer patients in this group were testing positive in this period than previously. A similar pattern was observed in the patients aged 45–54 years old.

Restricting the data set to CCGs that were recorded in monitoring information as having submitted 'complete' onboarding data resulted in the number of CCGs dropping from 135 to 38. For those aged between 60-69 years (that were not CEV or living in care homes and were onboarded from January 2021 onwards), there were 1,297 onboarded patients out of 21,487 patients testing positive for these CCGs with complete onboarding data. A similar pattern was observed in the patients aged 45–54 years.

Figure 1: Flowchart of patients onboarded and total patients (in parentheses) after application of each eligibility rule

Cleaned GDPPR data set:
Onboarded = 20,750 (2,236,300)

↓

Patients with a first positive test:
Onboarded = 16,832 (2,232,382)

↓

Patients who are not a CEV:
Onboarded = 13,124 (2,013,730)

↓

Patients who are not lving in a care home:
Onboarded = 12,996 (1,993,797)

↙ ↘

Eligible patients
Patients aged between 60-69 years, excluding indeterminate ages at the cut-off,
Onboarded = 2,904 (193,111)
Number of CCGs = 135

↓

Eligible patients
Patients aged between 60-69 years, excluding indeterminate ages at the cut-off, tested in Jan 2021 onwards:
Onboarded = 2,439 (91,207)
Number of CCGs = 135

↓

Eligible patients
Patients aged between 60-69 years, excluding indeterminate ages at the cut-off, tested in Jan 2021 onwards, including only CCGs with complete* data:
Onboarded = 1,297 (21,487)
Number of CCGs = 38

Eligible patients
Patients aged between 45-54 years, excluding indeterminate ages at the cut-off,
Onboarded = 2,810 (391,409)
Number of CCGs = 135

↓

Eligible patients
Patients aged between 45-54 years, excluding indeterminate ages at the cut-off, tested in Jan 2021 onwards:
Onboarded = 2,323 (180,859)
Number of CCGs = 135

↓

Eligible patients
Patients aged between 45-54 years, excluding indeterminate ages at the cut-off, tested in Jan 2021 onwards, including only CCGs with complete* data:
Onboarded = 1,289 (42,500)
Number of CCGs = 38

* Note: 'complete' data indicated that the CCG had submitted all the data they could access, rather than the data being complete in terms of all onboarded patients being recorded.

Onboarding and testing over time

Figure 2 below shows the number of eligible patients (that were not CEV or living in care homes) aged between 60 and 69 years across all CCGs that were onboarded, over calendar time (with the data binned/aggregated by week). Around January 2021, there was a substantial increase in the number of patients onboarded, followed by a gradual decline.

Figure 2: Eligible patients aged 60–69 onboarded over time

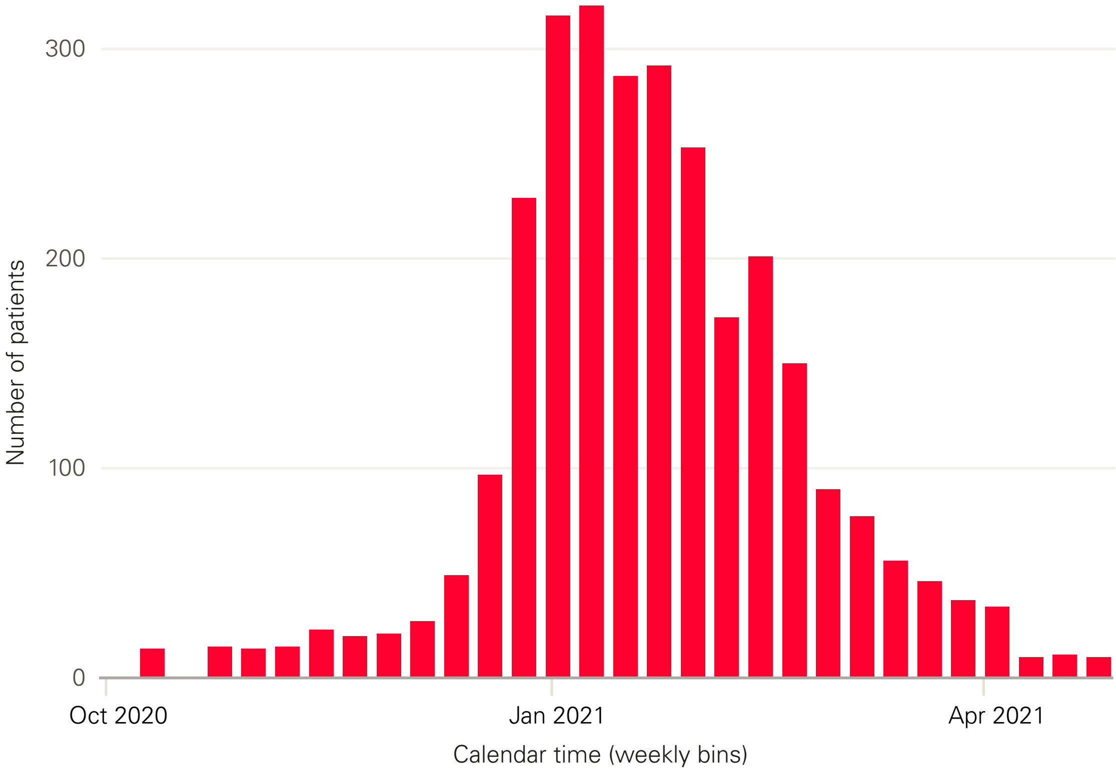

Note: some bars have been removed due to low numbers of patients.

Figure 3 shows the number of eligible patients (that were not CEV or living in care homes) aged between 60 and 69 years across all CCGs that tested positive, over calendar time (with the data binned by week). There were many patients testing positive around November 2020 and January 2021. Taken together, these two charts suggest that the lack of onboarding (despite many positive tests) prior to January 2021 would negate finding a discontinuity in onboarding prior to this time, thus providing motivation for looking at January 2021 onwards in addition to the entire period.

Figure 3: Eligible patients aged 60–69 testing positive over time

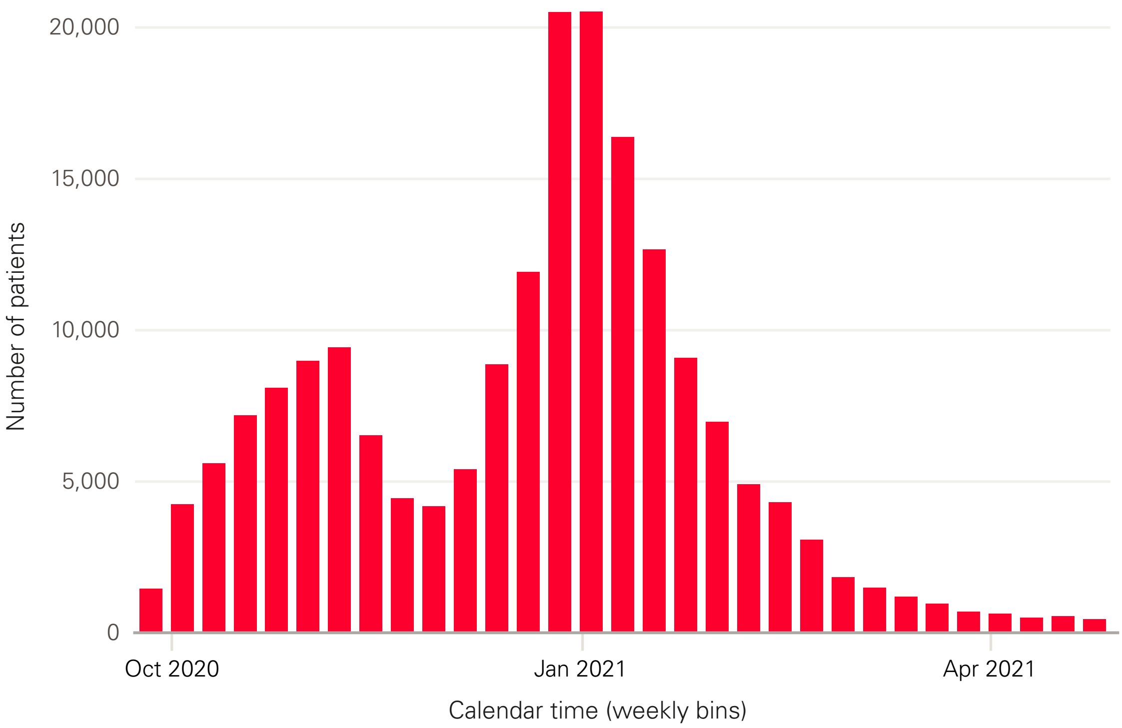

Similar charts to Figures 2 and 3 for patients aged between 45 and 54 years are available in the appendix (Figures A1 and A2). These show very similar patterns to the 60–69 age group.

Identifying any discontinuity

The precision with which we could identify the patients' dates of birth was weekly (to ensure individual patients could not be identified by their dates of birth). Therefore, the discontinuity plot below shows the percentage of eligible patients (that were not CEV or living in care homes) aged between 60 and 69 years , across all CCGs, onboarded by weeks of age either side of the cut-off (in this case, age 65). Due to low numbers, the data were binned by fortnights. The ideal situation in RDD is for the discontinuity to be close to 100% – ie below the cut-off no patients are onboarded, and above the cut-off all patients are onboarded. In this case, the discontinuity is almost imperceptible. The percentage of patients below the cut-off is very low (~1 to 2%), but so is the percentage above the cut-off (~2 to 3%). There is also a great deal of noise.

Figure 4: RDD plot showing percentage of eligible patients onboarded, by distance from the cut-off in weeks at age 65

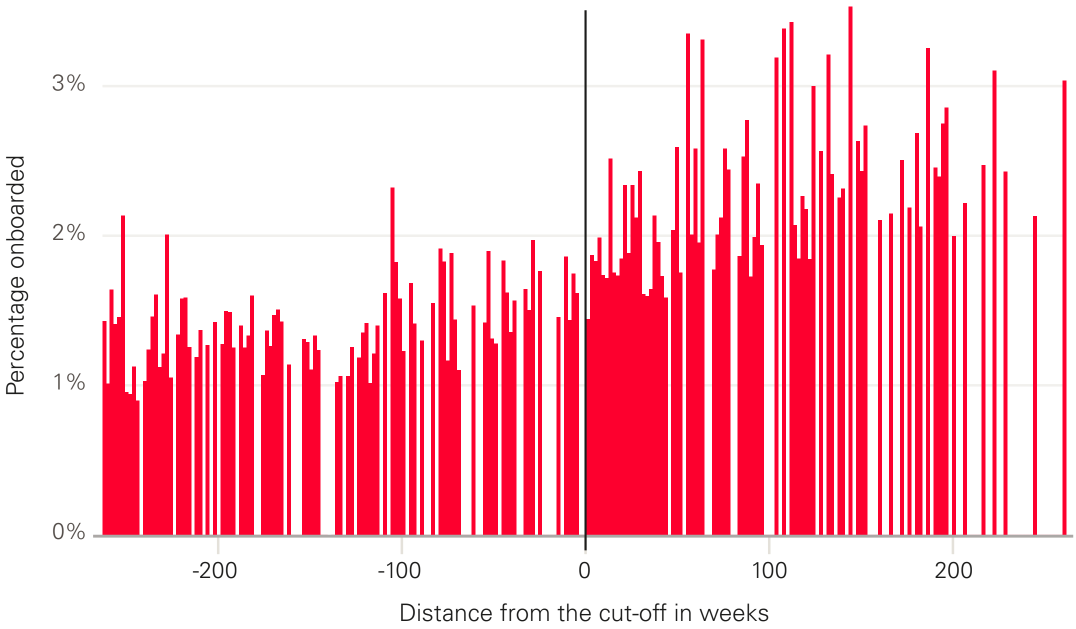

Note: some bars have been removed due to low numbers of patients.

The plot was smoothed by binning the data in years either side of the cut-off (Figure 5). This did not improve the situation greatly, although it shows the small change in onboarding more clearly either side of the cut-off.

Figure 5: RDD plot showing percentage of eligible patients onboarded, by distance from the cut-off in years at age 65

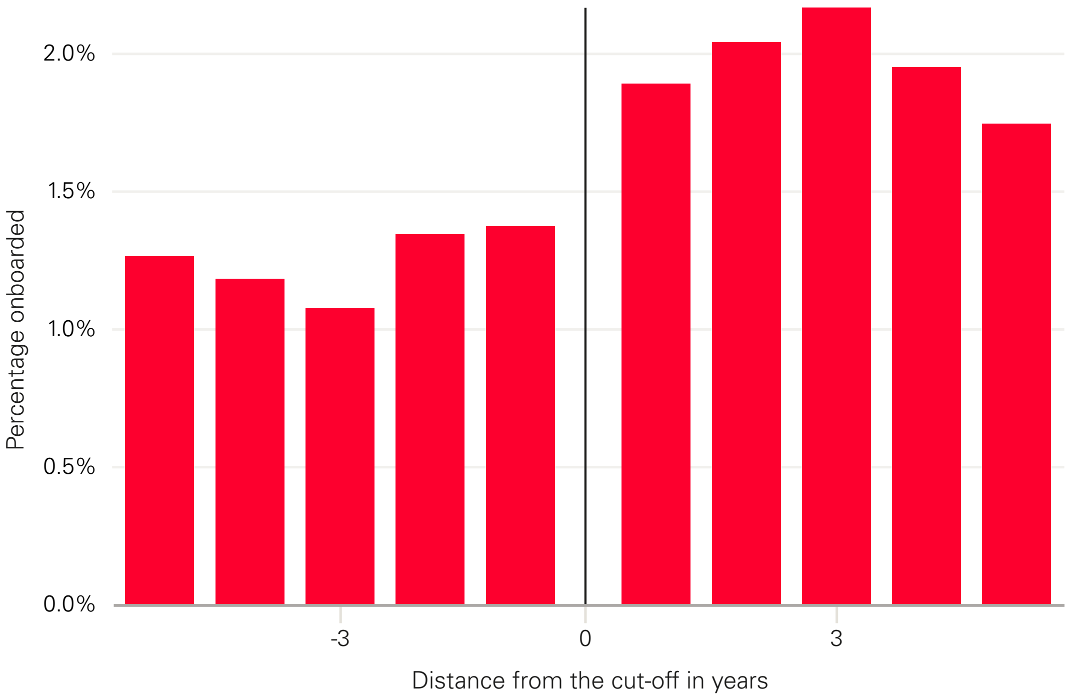

To see whether the situation improved for patients testing positive from January 2021 onwards, the discontinuity plot was redrawn for this period across all CCGs (Figure 6). The proportion of patients onboarded increased slightly, but there was still very little in the way of a discontinuity present.

Figure 6: RDD plot showing percentage of eligible patients onboarded, testing positive from January onwards, by distance from the cut-off in years at age 65

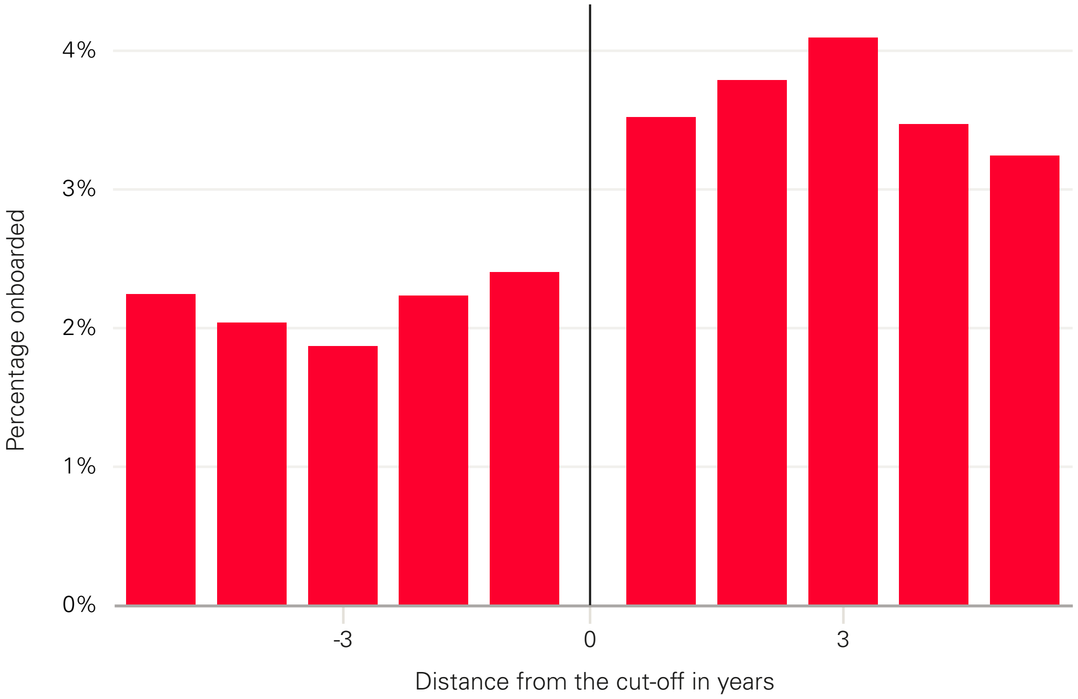

Finally, the plot was redrawn only including patients from those CCGs identified as having submitted 'complete' onboarding data (Figure 7). Again, this resulted in an increase in the proportion of patients onboarded either side of the cut-off, but there was still very little in the way of a discontinuity present.

Similar plots were created for those aged 45–54 years (see Appendix, Figures A3, A4, A5). These produced comparable results in terms of the discontinuity, although with lower proportions of patients onboarded either side of the cut-off. Ultimately, it did not seem that many patients over the threshold ages were onboarded, compared to those with a first COVID-19 positive test.

Figure 7: RDD plot showing percentage of eligible patients onboarded, testing positive from January onwards, for CCGs with 'complete' onboarding data, by distance from the cut-off in years at age 65

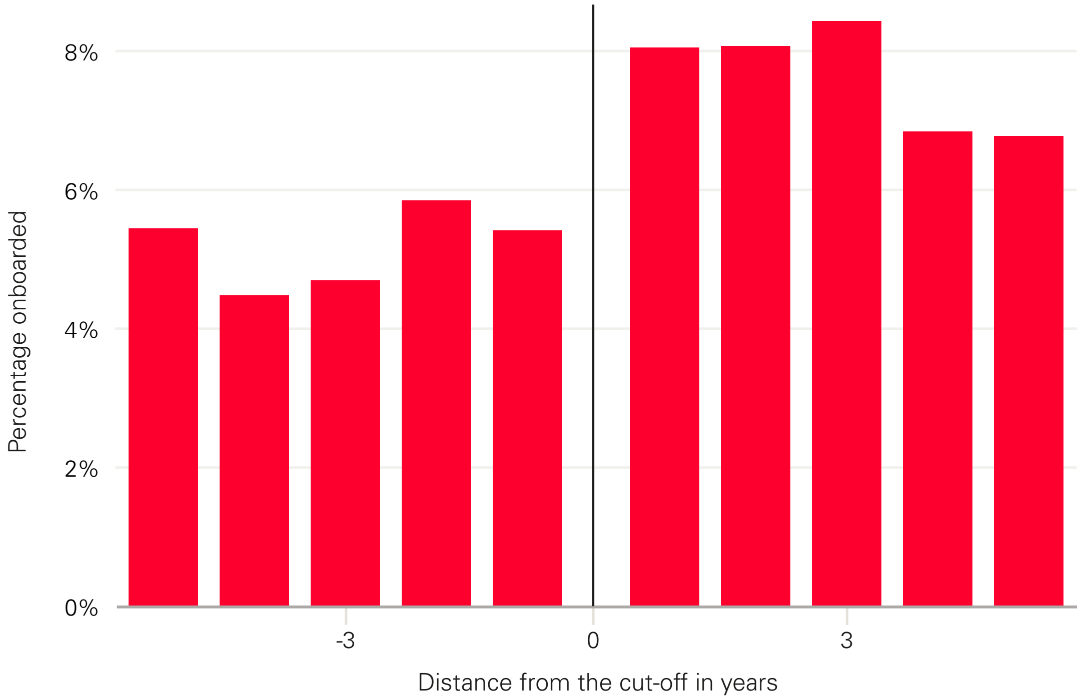

Onboarding either side of the cut-off by CCG

To identify whether there might be a subset of CCGs for which the onboarding rate was reasonably high above the cut-off but low below it, scatter plots of onboarding rates either side of the cut-off by CCG were produced. In these plots, the x-axis shows the onboarding rate below the cut-off and the y-axis shows the onboarding rate above the cut-off. The size of the points represents the CCG sample size (N). Ideally, in an RDD study, one would hope to find that most of the CCGs are clustered in the top left-hand corner, with close to 100% onboarding on the y-axis and close to 0% onboarding on the x-axis.

Figure 8 shows this plot for eligible patients 5 years either side of the age 65 cut-off, from January 2021 onwards. Although there are a few CCGs with greater differences between the onboarding rates either side of the cut-off, most CCGs had very low rates of onboarding and there is mostly little difference in the rates either side of the cut-off. The maximum difference in onboarding rate within a CCG was ~27%.

Figure 8: Plot of onboarding rate either side of the cut-off (age 65) for eligible patients testing positive from January onwards, by CCG

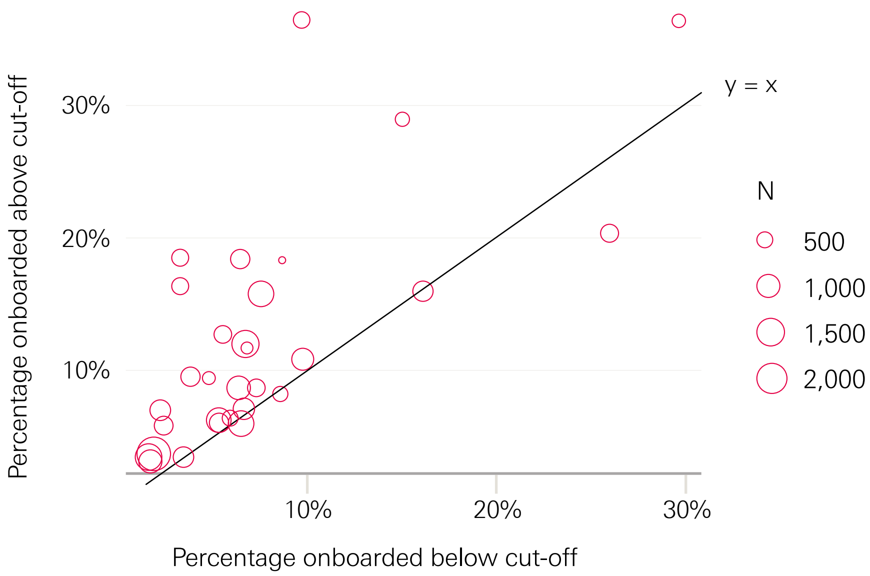

Note: CCGs with <10 patients onboarded have been removed.

Figure A6 in the Appendix shows a similar plot for those aged 45–54 years tested from January 2021 onwards and shows that the maximum difference within a CCG was only ~17%.

Power analysis

To identify whether there might be a subset of CCGs for which there may be sufficient power to identify a possible effect of CO@h on health outcomes, a series of power analyses were conducted. During the period October 2020 to April 2021, there were approximately 300,000 first COVID-19 positive tests in the age group 65–84.[5] During the same period, approximately 30,000 patients in this age group were admitted to hospital,[6] giving a rough estimate of 10% for the probability of being admitted given a first positive test. Table 1 shows the estimated power of the analysis if the baseline risk of the outcome (in the absence of the intervention) was assumed to be 10%, eg for an admission to hospital. The risk in the treated group (due to the effect of CO@h) was assumed to be 5% (ie a reduction in the baseline risk of 50%). Alpha was assumed to be 5% and a two-sided test was used.

The table shows the results of the power calculations (calculated using the {pwr} package in R), at both age 65 and age 50, for all CCGs that have a certain minimum discontinuity at the cut-off, using patients tested from January 2021 onwards. The sample size (N) above and below the cut-off shows the number of eligible patients within the groups. The effect size shows Cohen's h.[7] As a rule of thumb, h = 0.2 is regarded as a 'small' effect size, h = 0.5 is a 'medium' effect size, and h = 0.8 is a 'large' effect size. The resulting estimates of effect size are all far below 0.2 (the maximum being approximately 0.03). The power estimate shows the probability of finding a statistically significant result if there is an underlying mean difference between the groups of at least the specified 50%. For example, the first row of the table shows that if we include in the analysis all patients aged 60–69, tested from January

2021 onwards and eligible for the study, in CCGs where there is a minimum discontinuity of 0.05 (ie 5%), then the power is 0.10. This means that there is a 10% probability of identifying an effect if there is a reduction in baseline risk of 50%. In no combination of scenarios tabulated is the power higher than 10%. For rarer outcomes such as admission to ICU or death,[8] the power would be even lower.

It is clear that statistical power would be severely limited in these examples, meaning that undertaking the analysis would most likely result in null findings, even if there was a positive effect of the intervention on hospital admissions and mortality.

Table 1: Power analysis under various conditions (two-sided test, alpha = 5%, baseline risk = 10%, treated risk = 5%)

Cut-off	Minimum discontinuity	N below cut-off	N above cut-off	Effect size (Cohen's h)	Power
65	0.05	4,781	2,433	0.0167	0.1026
65	0.10	1,666	838	0.0271	0.0981
50	0.05	928	981	0.0229	0.0792
50	0.10	527	586	0.0293	0.0777

Note: power calculations undertaken using the {pwr} package in R.

Conclusions

An RDD design would have allowed for a robust and reliable evaluation of CO@h, avoiding some of the risks of observed and unobserved biases which are common in observational studies, if CO@h had been implemented in line with the standard operating procedure. However, this was not the case and the analyses presented here have shown that an RDD is not suitable for the evaluation of the CO@h programme in the current setting. The number of patients onboarded, combined with the lack of discontinuity at the age cut-off, lead to a lack of statistical power, virtually guaranteeing a null result, even if the intervention was successful in reducing emergency hospital outcomes. Additionally, the lack of discontinuity of the magnitude generally required for a high quality RDD would lead one to question the validity of the assumptions of this method relating to non-compliance.[9]

The results of this feasibility analysis highlight some of the difficulties in embedding robust quantitative evaluation into the implementation of new interventions. Although an RDD design would have allowed us to robustly evaluate the causal impact of the programme if it had been implemented according to the standard operating procedure guidelines, these were not mandatory, and local teams will have implemented CO@h in the manner that best reflected their local population's needs and their limited resources during the height of the pandemic. Furthermore, several areas had implemented CO@h services before the standard operating procedure was made available.

A number of approaches to evaluate the CO@h programme are being pursued by the evaluation partners, each with different assumptions, strengths and weaknesses, which will together aim to provide reliable conclusions about the effect of CO@h on outcomes. Understanding the effect of CO@h will help inform the use of pulse oximeters in any future waves of COVID-19[10] as well as provide learning for the remote monitoring of other conditions in the future.

References

1. Cattaneo MD, Idrobo N, Titiunik R. *A Practical Introduction to Regression Discontinuity Designs: Foundations.* Cambridge University Press; 2020 (https://www.cambridge.org/core/elements/abs/practical-introduction-to-regression-discontinuity-designs/F04907129D5C1B823E3DB19C31CAB905).
2. Cattaneo MD, Idrobo N, Titiunik R. *A Practical Introduction to Regression Discontinuity Designs: Volume II.* Draft. Cambridge University Press; 2018 (https://cattaneo.princeton.edu/books/Cattaneo-Idrobo-Titiunik_2018_CUP-Vol2.pdf).
3. Lloyd T, Parry W. Statistical analysis protocol for an evaluation of COVID Oximetry @ home using a Regression Discontinuity Design. The Health Foundation; 2021 (https://health.org.uk/sites/default/files/2021-03/SAP_Oximetry_Regression_Discontinuity_Design.pdf).
4. Torjesen I. Covid-19: Patients to use pulse oximetry at home to spot deterioration. *BMJ*. 2020; 371 (doi:10.1136/bmj.m4151).
5. UK Government. Download data | Coronavirus in the UK [webpage]. *GOV.UK Coronavirus (COVID-19) in the UK;* 14 September 2021. Accessed 15 September 2021. (https://coronavirus.data.gov.uk/details/download).
6. NHS England. COVID-19 Hospital Activity [webpage]. *NHS England;* 9 September 2021. Accessed 15 September 2021 (https://www.england.nhs.uk/statistics/statistical-work-areas/covid-19-hospital-activity).
7. Cohen J. *Statistical Power Analysis for the Behavioral Sciences.* 2nd ed. Routledge; 1988 (https://www.taylorfrancis.com/books/mono/10.4324/9780203771587/statistical-power-analysis-behavioral-sciences-jacob-cohen).
8. Gray WK, Navaratnam AV, Day J, et al. Variability in COVID-19 in-hospital mortality rates between national health service trusts and regions in England: A national observational study for the Getting It Right First Time Programme. *EClinicalMedicine.* 2021; 35 (doi:10.1016/j.eclinm.2021.100859).
9. Angrist JD, Imbens GW, Rubin DB. Identification of Causal Effects Using Instrumental Variables. *Journal of the American Statistical Association.* 1996; 91(434): 444-455 (doi:10.2307/2291629).
10. Gurdasani D, Drury J, Greenhalgh T, et al. Mass infection is not an option: we must do more to protect our young. *The Lancet.* 2021; 398(10297): 297–298 (doi:10.1016/S0140-6736(21)01589-0).

Appendix

Figure A1: Eligible patients aged 45–54 onboarded over time

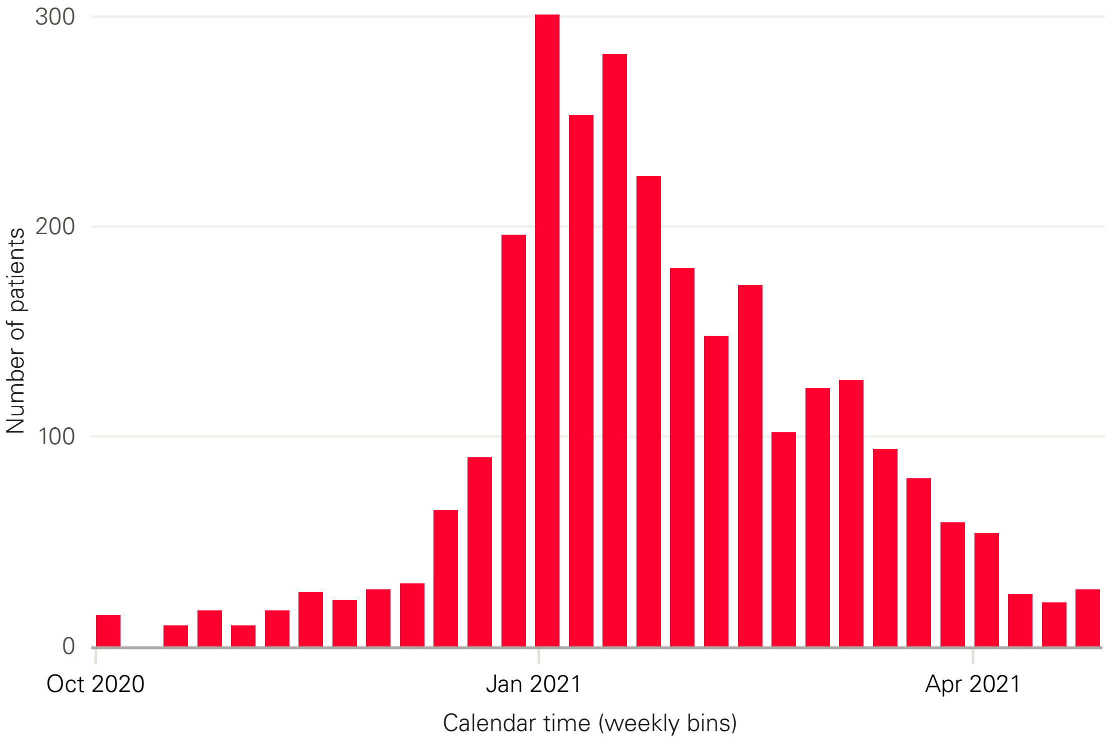

Note: some bars have been removed due to low numbers of patients.

Figure A2: Eligible patients aged 45–54 testing positive over time

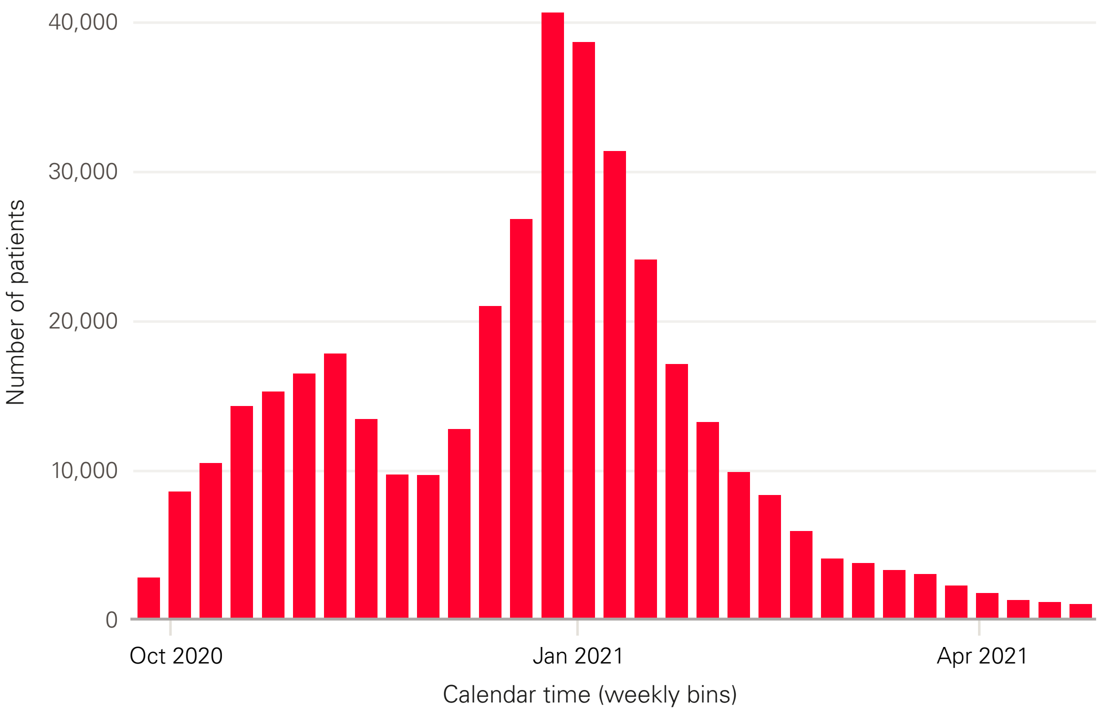

Figure A3: RDD plot showing percentage of eligible patients onboarded, by distance from the cut-off in years at age 50

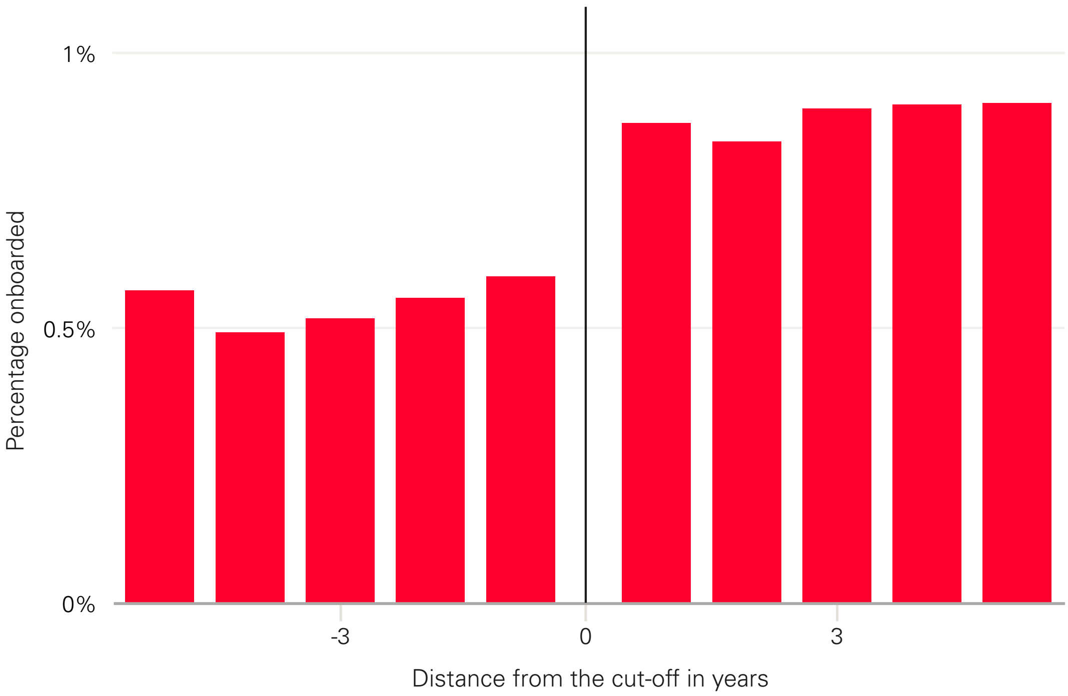

Figure A4: RDD plot showing percentage of eligible patients onboarded, testing positive from January onwards, by distance from the cut-off in years at age 50

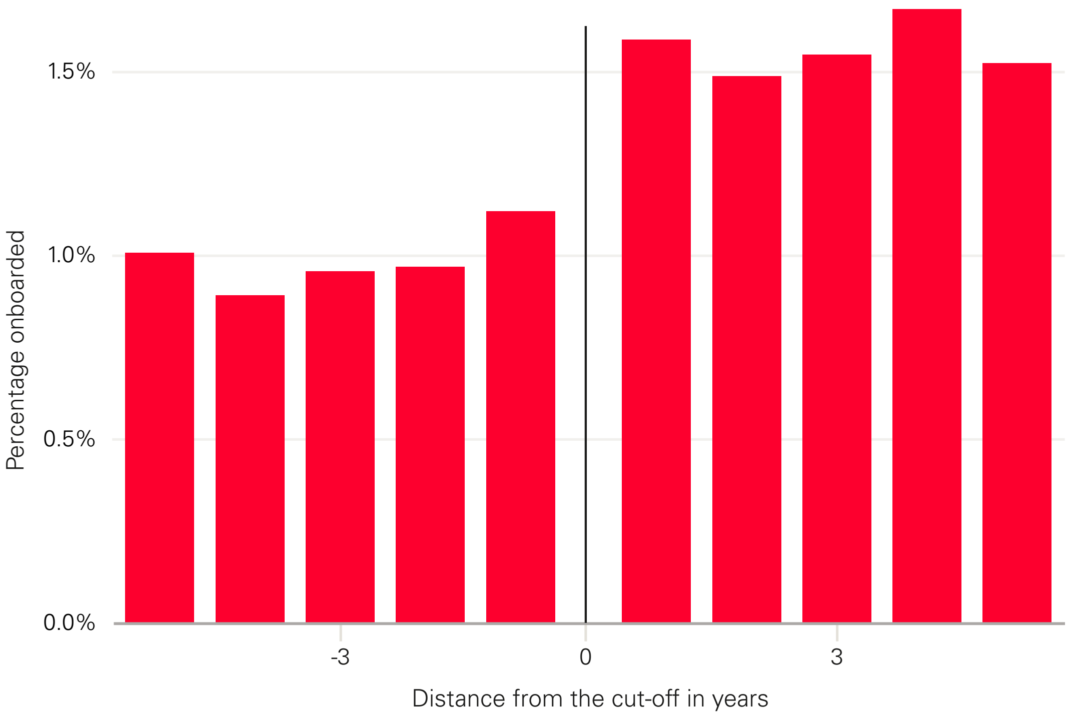

Figure A5: RDD plot showing percentage of eligible patients onboarded, testing positive from January onwards, for CCGs with 'complete' onboarding data, by distance from the cut-off in years at age 50

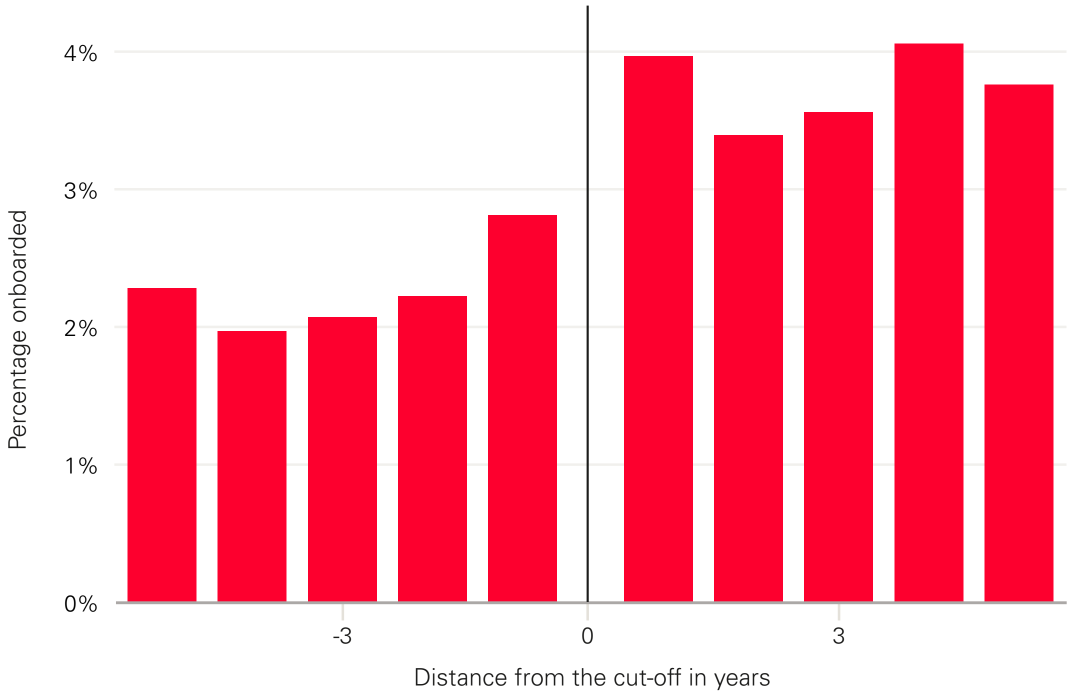

Figure A6: Plot of onboarding rate either side of the cut-off (age 50) for eligible patients testing positive from January onwards, by CCG

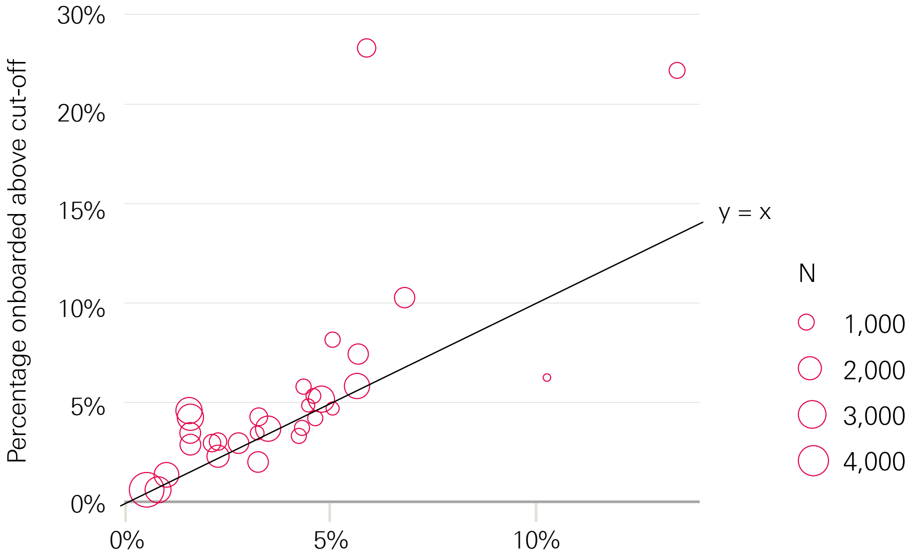

Note: CCGs with <10 patients onboarded have been removed.

The Health Foundation
8 Salisbury Square, London EC4Y 8AP
T +44 (0)20 7257 8000
E info@health.org.uk
@HealthFdn
www.health.org.uk

The Health Foundation is an independent charity committed to bringing about better health and health care for people in the UK. Our aim is a healthier population, supported by high quality health care.

ISBN: 978-1-911615-69-9
COVID Oximetry @home (CO@h): Feasibility analysis for evaluation of CO@h using regression discontinuity
Registered charity number: 286967
Registered company number: 1714937