American Journal of Epidemiology **Submitted Manuscript**[1]

Special Collection: Mental Health

Title: Machine Learning Detects Heterogeneous Effects of Medicaid Coverage on Depression

Authors: Ryunosuke Goto, MD1; Kosuke Inoue, MD, PhD2; Itsuki Osawa, MD3; Katherine Baicker, PhD4; Scott L. Fleming, BS5; Yusuke Tsugawa, MD, PhD6,7

ORCiD IDs: N/A

Correspondence Address: Ryunosuke Goto, MD. Department of Pediatrics, The University of Tokyo Hospital Address: 7-3-1 Bunkyo-ku, Hongo, Tokyo 113-8655, Japan. Email: rgoto@m.u-tokyo.ac.jp. Phone: +81-3-3815-5411. ORCID: 0000-0002-4164-0990

Joint Authorship: N/A

[1] N/A indicates not applicable.

Affiliations: 1. Department of Pediatrics, The University of Tokyo Hospital, Tokyo, Japan; 2. Department of Social Epidemiology, Graduate School of Medicine, Kyoto University, Kyoto, Japan; 3. Department of Emergency and Critical Care Medicine, The University of Tokyo Hospital, Tokyo, Japan; 4. University of Chicago, Chicago, IL; 5. Department of Biomedical Data Science, Stanford University, Stanford, CA; 6. Division of General Internal Medicine and Health Services Research, David Geffen School of Medicine, University of California, Los Angeles, CA; 7. Department of Health Policy and Management, Fielding School of Public Health, University of California, Los Angeles, CA

Key words: machine learning, causal machine learning, causal forest, generalized random forest, causal inference, high-benefit approach, mental health, psychiatry, depression, Oregon Health Insurance Experiment, Medicaid

Acknowledgments[2]**:** This study was presented at the Society for Epidemiologic Research 2023 Annual Meeting

Funding: RG receives funding from the Japan Foundation for Pediatric Research (22-001) and the Chernobyl-Fukushima Medical Fund for other work not related to this study. KI receives funding from the Japan Society for the Promotion of Science (21K20900 and 22K17392), the Japanese Endocrine Society, and the Program for the Development of Next-generation Leading Scientists with Global Insight (L-INSIGHT) sponsored by the Ministry of Education, Culture, Sports, Science and Technology (MEXT), Japan for other work not related to this study. IO receives funding from the Toranomon Hospital, Tokyo, Japan for other work not related to this study. KB receives funding from the National Institute of Health (NIH)/National Institute on Aging (P01AG005842,

[2] Study investigators, conference presentations, preprint publication information, thanks.

R01AG034151) for other work not related to this study, and serves on the board of directors of Eli Lilly and Mayo Clinic and on advisory panels of the Congressional Budget Office and National Institute for Health Care Management. SLF receives funding from a National Defense Science and Engineering Graduate Fellowship and a Stanford Graduate Fellowship for other work not related to this study. YT receives funding from the NIH/National Institute on Aging (R01AG068633 & R01AG082991), the NIH/National Institute on Minority Health and Health Disparities (R01MD013913), and Gregory Annenberg Weingarten GRoW @ Annenberg for other work not related to this study, and serves on the board of directors of M3, Inc. The funders had no role in considering the study design or in the collection, analysis, interpretation of data, writing of the report, or decision to submit the article for publication.

Conflict of Interest: Authors declare that they have no competing interests.

Disclaimer: N/A

Data Availability Statement: All data used in this study are available online from the National Bureau of Economic Research's Public Use Data Archive and can be accessed at https://www.nber.org/research/data/oregon-health-insurance-experiment-data. The code will be available upon reasonable request to the corresponding author.

ABSTRACT

In 2008, Oregon expanded its Medicaid program using a lottery, creating a rare opportunity to study the effects of Medicaid coverage using a randomized controlled design (Oregon Health Insurance Experiment). Analysis showed that Medicaid coverage lowered the risk of depression. However, this effect may vary between individuals, and the identification of individuals likely to benefit the most has the potential to improve the effectiveness and efficiency of the Medicaid program. By applying the machine learning causal forest to data from this experiment, we found substantial heterogeneity in the effect of Medicaid coverage on depression; individuals with high predicted benefit were older and had more physical or mental health conditions at baseline. Expanding coverage to individuals with high predicted benefit generated greater reduction in depression prevalence than expanding to all eligible individuals (21.5 vs. 8.8 percentage point reduction; adjusted difference [95%CI], +12.7 [+4.6,+20.8]; P=0.003), at substantially lower cost per case prevented ($16,627 vs. $36,048; adjusted difference [95%CI], -$18,598 [-$156,953,-$3,120]; P=0.04). Medicaid coverage reduces depression substantially more in a subset of the population than others, in ways that are predictable in advance. Targeting coverage on those most likely to benefit could improve the effectiveness and efficiency of insurance expansion.

Introduction

Assessing the effects of health insurance on health can be challenging, because insured individuals differ from uninsured individuals in ways that may themselves directly affect health outcomes. In 2008, the state of Oregon allocated limited spots in its Medicaid program for low-income adults through a lottery, allowing researchers to assess the effects of health insurance coverage on health outcomes, healthcare utilization, and financial strain using a randomized controlled design.[1] Results from this Oregon Health Insurance Experiment (OHIE) showed that Medicaid coverage reduced financial strain[1] and increased healthcare use across settings, including emergency department (ED) use[2] and primary care visits.[1] The effects on physical health were mixed: self-reported health improved, but there were no detectable changes in physical health outcomes.[3] The effect on mental health, however, was substantial: Medicaid enrollees had a 10% lower probability of screening positive for depression,[3] a 50% lower likelihood of undiagnosed depression, and a 60% lower probability of untreated depression than the control group.[4]

These findings have important implications, as depression is one of the leading causes of disability in the US,[5] representing a major unmet health need for low-income populations, and those gaining insurance were much more likely to have their depression diagnosed and treated.[4] Health insurance can thus play a critical role in improving mental health. However, health insurance expansion comes with a substantial price tag, as insured people use more healthcare than the uninsured, and budgets for public insurance programs like Medicaid and Medicare impose a growing strain on state and federal budgets.[6] Evaluation of the effectiveness of expansions must incorporate both the costs and the benefits.[7]

The average benefits of Medicaid expansion in treating depression seen in the OHIE may mask substantial heterogeneity, with some people benefitting much more than others. In this post hoc analysis of the OHIE, we assess the degree of response heterogeneity and the extent to which it is predictable *ex ante*. By applying a novel machine learning method recently introduced in the econometrics literature, the causal forest,[8] we delineate the characteristics of individuals with high or low predicted benefit and evaluate both the health benefits and efficiency of an approach for targeting health insurance coverage on those most likely to benefit—an approach called the "high-benefit approach."[9]

Methods

Study sample

We analyzed data from the OHIE, a randomized-controlled trial of the effects of health insurance coverage. Multiple institutional review boards have approved the OHIE, and written informed consent was obtained from all participants in in-person data collection. The OHIE took advantage of the random allocation of a Medicaid program for low-income (below 100% of the federal poverty level), uninsured, able-bodied adults in Oregon in 2008. Details on the lottery are described elsewhere.[1]

To assess various outcomes, a total of 12,229 participants in Portland, Oregon were given in-person surveys an average of 25 months after the lottery began. The in-person survey contained questions on healthcare utilization, health insurance coverage, and medications. Additionally, several anthropometric and blood-pressure measurements were taken, and dried blood spots were also obtained. Depression was assessed using the eight-question version of the Patient Health Questionnaire (PHQ-8).[10] The details of the in-person data collection are described elsewhere.[3] Of these participants, we included individuals who responded to in-person surveys with outcome, treatment, and baseline variables (including select variables on ED utilization at baseline) available (**Figure S1**).

Variables

The primary outcome was whether or not an individual screened positive for depression (a binary outcome), defined as a PHQ-8 score of 10 or higher. We also evaluated the annual healthcare cost per case of depression prevented, calculated by dividing the total annual healthcare spending (for any healthcare service utilization) in our sample by the expected number of depression cases prevented. The expected number of depression cases prevented was calculated by multiplying the size of our sample by the average treatment effect of Medicaid coverage on depression. The average annual healthcare spending was estimated by multiplying the individual-level numbers of prescription drugs, self-reported office visits, emergency department visits, and hospital admissions by the average estimated cost for each type of utilization (methods for calculating the healthcare spending are described in prior work on OHIE).[3] Whether an individual won a lottery for Medicaid was used as an instrumental variable to estimate the health benefit of Medicaid coverage.

The following baseline covariates were used in the analyses: gender; age; educational level (more than a high school diploma or not); race and ethnicity (non-Hispanic White, non-Hispanic Black, Hispanic, or other); whether the interview was conducted in English or not (in Spanish or through an interpreter of another language); diagnoses before the lottery (hypertension, diabetes, high cholesterol, asthma, heart attack, congestive heart failure, empyema/chronic obstructive pulmonary disease (COPD), kidney failure, cancer, and depression); ED utilization (the number of ED visits, having had any ED visits for mood disorders, having had any ED visits for psychiatric conditions or substance abuse); and hospital and ED spending (sum of total hospital charges, and sum of total ED charges prior to randomization). As for gender, participants were asked at the time of the survey whether their gender was male; female; transgender: male to female; or transgender: female to male (we acknowledge that this classification is not inclusive and needs to be revised, but given the post hoc nature of this study, we used this classification). Since there were very few individuals who answered “transgender,” we dichotomized the gender variable, including individuals who answered "transgender: male to female” in the female gender group and "transgender: female to male” in the male gender group. As individuals who won the lottery won the eligibility for health insurance coverage for all members of their households, all models included the number of household members on the lottery list.[3] When constructing the causal forest model, categorical covariates were converted to dummy variables, and a total of 23 covariates were used in the model.

Data on ED visits and charges were taken from visit-level data for all ED visits to 12 hospitals in the Portland area in the pre-randomization period, defined as January 1, 2007 to March 9, 2008. These data were truncated at twice the 99th percentile of the original distribution to ensure de-identification.[3] Additional utilization data came from self-reports of office visits and a catalog of prescription drugs taken during the in-person data collection. All other data were obtained from information provided by the participants when they signed up for the lottery (prior to randomization) and from self-reported in-person surveys conducted from August 31, 2009 until October 13, 2010.

Statistical analyses

We estimated the individual treatment effects (ITEs), defined as the treatment effect for each individual, conditional on the individual's observed characteristics, of Medicaid coverage on the probability of a positive screening for depression using a causal forest, a machine learning-based model that predicts the treatment effects for individuals based on their covariates.[8] The causal forest algorithm extends the regression tree and random forest algorithms to estimating the treatment effects for different subgroups, conditional on their observed characteristics.[8] Whereas traditional subgroups analyses are limited to subgroups specified a priori,[11] the causal forest allows for improved characterization of treatment effect heterogeneity by searching across the full spectrum of individual characteristics.[12] To avoid overfitting, the causal forest model uses a randomly-selected proportion of the entire sample to build each tree, which is further split into a subsample for determining the tree structure (the splitting subsample) and a subsample for estimating the treatment effect in each leaf (the estimating subsample), a property called "honesty."[8] We used cross-validation to tune the proportions of these subsamples, along with the number of variables considered for each split, the minimum number of samples each node should contain, the proportion of the data used for determining splits, whether the estimation sample tree should be pruned such that no leaves are empty, the maximum imbalance of a split, and the penalty for imbalanced splits. In addition, we constructed the causal forest model and estimated the ITEs using cross-fitting with 10 folds, which has been shown to be an efficient form of data-splitting.[13] For each fold k, this procedure fits the causal forest on observations not included in fold k and predicts the ITEs of the observations in fold k.[14] The calibration of the causal forest was evaluated by ranking the ITEs into quintiles within each of the folds, calculating the average treatment effect of individuals in each quintile with the causal forest, and comparing them to the ordinary least squares estimates. For a model with good calibration, the average treatment effects estimated with the causal forest and ordinary least squares for each quintile will be similar, and will incrementally increase across quintiles. In addition, the calibration and the heterogeneity of the model were evaluated using the best linear projection of the ITEs, following the approach by Semenova and Chernozhukov.[15] The best linear projection evaluates whether the average prediction of the ITEs is correct ("mean prediction" in **Table S1**) and whether the forest adequately captures the heterogeneity in ITEs ("differential prediction" in **Table S1**). The ITEs were represented in percentage point reduction in prevalence of depression (the signs of the estimates were flipped and multiplied by 100 so they can be interpreted as percentage point reduction; a positive ITE represents decreased risk of screening positive for depression). All covariates listed above were used in the causal forest. To construct the causal forest model, we first used whether an individual won the lottery as an instrumental variable for Medicaid coverage and performed an intention-to-treat analysis as a supplemental analysis.[16] We performed the instrumental variable analysis in a manner similar to the two stage least squares approach, as with the original OHIE studies (**Appendix S1**).[1-4]

Using the predicted ITEs, we compared the characteristics of individuals with high predicted benefit and low predicted benefit from Medicaid coverage, defined as those with predicted ITEs above vs. below (or equal to) the median of the full sample, by computing standardized absolute mean differences for each covariate. Additionally, we plotted the ITEs across age and the number of comorbidities, variables chosen based on the comparison between high and low ITE groups.

Next, we estimated the average treatment effect of Medicaid coverage on depression for two separate scenarios: (1) expanding coverage to individuals with high predicted benefit from health insurance (defined as individuals with predicted ITEs above the median), and (2) expanding coverage to all individuals in the sample. The average treatment effects were estimated with instrumental variable regressions. The mean difference in the average treatment effects calculated using the two approaches, its 95% confidence interval, and P-value were obtained using percentile bootstrapping with 10,000 replications. We compared two outcome variables: the health benefit (the depression cases averted) and the healthcare spending (the annual healthcare spending per case of depression prevented). We performed robustness checks by dividing the sample into high and low predicted benefit groups using the median of each fold (instead of the median of the full sample) as the cutoff (**Appendix S2**).

Finally, using the important predictors of ITEs (i.e., covariates with the greatest predictive value of ITEs) identified with the causal forest, we investigated whether we could identify individuals with high ITEs using a small number of variables. In particular, we estimated the average treatment effect of providing Medicaid coverage to individuals selected based on age, the variable identified as the most important predictor.

As supplemental analyses, we compared ITEs across race and ethnicity, stratified by age. ITEs were estimated using the causal forest model, fixing variables other than age and race and ethnicity at the median, and were represented as a heatmap. In addition, we compared the number of depression cases averted by race and ethnicity, for the high-benefit approach and the population approach. The number of depression cases averted for each approach was calculated by multiplying the total number of individuals by the treatment effect in the subgroup. Differences in the number of depression cases averted were estimated using bootstrapping with 10,000 replications. All analyses were conducted using R version 4.1.1 using the package grf (R Project for Statistical Computing).[17]

Results

Basic characteristics

A total of 10,068 low-income individuals met the inclusion criteria. Of these individuals, 5,274 were lottery winners and 4,794 were in the control group. The distributions of baseline characteristics were similar between the lottery winners and the control group (**Table 1**).

The causal forest model for predicting the individual treatment effects

The causal forest model of the effects of Medicaid coverage on screening positive for depression showed good calibration (**Table S1, Figure S2**). There was significant heterogeneity in the treatment effect of Medicaid coverage on depression based on the best linear projection of the ITEs (**Table S1**). The ITEs showed a bimodal distribution (**Figure 1**). The median of the predicted ITE was 11.6 percentage points (pp) reduction, and the cutoff between high and low predicted benefit was set at this value. The variable importance plot showed that age was frequently split on in the causal forest (**Figure S3**). The causal forest for the intention-to-treat analysis similarly showed good calibration (**Table S2, Figure S4**).

Characteristics of individuals with high vs low ITEs

Comparing the characteristics of individuals with high vs. low predicted benefit from Medicaid coverage, we found that individuals with high predicted benefit were older and more likely to have physical or mental health conditions at baseline (**Table 2**). We did not observe large differences in gender, educational level, or ED visits at baseline, although those with low predicted benefit were more likely to be Hispanic. The weighted prevalence of those who screened positive for depression in the control group for the high ITE vs. low ITE groups were

36.4% and 23.2%, respectively, and in the treated group were 30.3% and 23.5%, respectively. Older individuals tended to have higher predicted ITEs and more comorbidities (**Figure 2**). While the ITEs were relatively constant across younger individuals up to the mid-30s, the predicted ITEs increased drastically from the mid-30s to the mid-50s (**Figure 2**). The differences between high vs. Low ITE groups were similar when the cutoff for high predicted benefit was defined as the median for each of the 10 folds (**Table S3**).

Expanding Medicaid coverage to individuals with high predicted benefit

Expanding Medicaid coverage to individuals with estimated ITEs above the median achieved greater average reduction in prevalence of depression compared to expanding coverage to all individuals in the sample (average treatment effect, 21.5 vs. 8.8 pp reduction; adjusted difference [95% CI], +12.7 [+4.6, +20.8]; P=0.003; **Table 3**). The healthcare spending required to prevent a case of depression was lower when Medicaid expansion targeted those individuals with high estimated ITEs compared to covering all eligible individuals (annual healthcare spending per case of depression prevented, $16,627 vs. $36,048; adjusted difference [95% CI], -$18,598 [-$156,953, -$3,120]; P=0.04). We obtained similar results when the cutoff for high predicted benefit was defined as the median for each of the 10 folds (**Table S4**).

Targeting Medicaid coverage expansion by age

The results of the causal forest analysis indicated that that age is the covariate most predictive of ITEs. Therefore, we conducted a post hoc analysis using only the information on age. In particular, we evaluated the scenario of expanding Medicaid coverage to individuals aged 50 years or older, a cutoff we chose based on **Figure 2**. We found that this approach was associated with a larger reduction in prevalence of depression (23.1 vs. 8.8 pp reduction; adjusted difference [95% CI], +14.3 [+1.0, +27.6]; P=0.03; **Table S5**) in depression and was more efficient (annual healthcare spending per case of depression prevented, $16,430 vs. $36,048; adjusted difference [95% CI], -$18,598 [-$157,845, +$17,312]; P=0.09) compared to expanding coverage to all individuals in the sample. We also found that this approach was as effective (23.1 vs. 21.5 pp reduction; adjusted difference [95% CI], +1.6 [-8.6, +11.9]; P=0.75; **Table S6**) and efficient ($16,430 vs. $16,627; adjusted difference [95% CI], -$260 [-$9,456, +$22,372]; P=0.94; **Table S6**) as expanding coverage to individuals with high estimated ITEs based on the causal forest.

Comparison of effects by race and ethnicity

Based on our comparison of high vs. low ITE groups, we found that a larger number of Hispanic people in the low ITE group than in the high ITE group. More specifically, the proportion of people categorized as high ITE was 35.1% among Hispanic people, compared to 57.6% among the non-Hispanic Black population and 52.5% among the non-Hispanic White population. Comparing the number of depression cases averted by race and ethnicity, we found that the number of depression cases averted was higher in the high-benefit approach for the non-Hispanic White population, but not in the non-Hispanic Black or Hispanic populations (**Table S7**). However, predicted ITEs (stratified by age, the most important determinant of treatment effect heterogeneity in our model) were similar across race and ethnicity groups (**Figure S5**), indicating that race and ethnicity *per se* was unlikely to have been an important determinant of treatment effect heterogeneity.

Discussion

In this post hoc analysis of the OHIE using the machine learning causal forest, we found substantial – and predictable – heterogeneity in the effect of Medicaid coverage on depression. Those who experienced large improvements in depression were older and had more physical or mental health conditions at baseline. We found that providing Medicaid coverage to individuals with high likelihood of benefit as predicted using *ex ante* information—an approach known as the "high-benefit approach"[9]—reduced depression by a three-fold greater margin than providing coverage to all low-income individuals (the population approach).[18] This high-benefit approach was not only effective in preventing the depression cases, but also more cost-effective than broader expansions as captured by the healthcare spending per each case of depression averted. Such an approach may prove useful especially when expansions to all low-income individuals is not practical due to resource limitations. Taken together, our findings suggest that it is possible to use baseline information to prioritize coverage expansion to those who are likely to benefit the most.

The OHIE underscored the importance of health insurance in addressing the unmet mental health needs of a population by reducing the prevalence of undiagnosed and untreated depression.[3,4] We used a novel method for incorporating existing information to predict the heterogeneous effects of health insurance coverage, and find that in this case age is a key driver of the effect of Medicaid coverage even when other factors are considered. In addition, we provide new information about the relationship between age and the effect of Medicaid coverage: we show that the effect of Medicaid coverage increases drastically from the mid-30s and peaks for individuals in the mid-50s and above. Our analyses using the causal forest identified age as the most important predictor of ITEs (the fact that age was a strong predictor was not known *ex ante*, nor was the functional form of that relationship), and our post hoc analysis revealed that providing coverage to individuals with age 50 and achieves similar effectiveness and efficiency to more complicated eligibility criteria. This may be because the effects of socioeconomic adversity on depression are said to accumulate over time: that is, the longer the exposure to the negative consequences of socioeconomic status, the more lower socioeconomic status contributes to worse mental health.[19] In the context of our study, if we assume that older individuals were exposed to lower socioeconomic status for longer periods of time, then insurance coverage likely helped older individuals out of the negative consequences of poverty (such as the financial strain of getting healthcare and the distress from not being able to afford it), thereby alleviating their mental health burden, more so than younger individuals. This should be confirmed in future studies. Ultimately, our findings highlight the importance and utility of evaluating the heterogeneity in treatment effects across the full spectrum of individual-level demographic and health characteristics as well as the intricate interactions among them using the causal forest. Future studies could use other ITE estimators to explore whether our results can be replicated.[20-22]

Importantly, our approach facilitates the policy option of prioritizing coverage and treatment plans based on predicted benefit. Though their application in healthcare is scarce, several models have been developed for estimating treatment effects at the individual level and detecting treatment effect heterogeneity: the causal forest,[8] double/debiased machine learning,[14] and orthogonal random forest[23] to name a few, which have been gaining attention in the econometrics literature. These methods have shown promise not only in randomized trials but also in observational data,[15,24] suggesting their value for policy evaluation in experimental and observational settings alike. As our study suggests, applying these methods to exploring a policy's treatment effect heterogeneity and determining the optimal coverage based on the predicted benefit could be a new avenue for precision policy making.

These promising findings should not undermine the importance of addressing disparities in healthcare, especially in light of the possibility that algorithm-based healthcare coverage may exacerbate disparities if the estimated treatment effect is smaller among minoritized populations. First, it is possible that the data used to develop the algorithms may be biased if minoritized patients were more or less likely to have coded diagnosis of certain conditions.[25,26] However, it is important to note that most variables available in the OHIE data were collected using surveys, which is less sensitive to biased coding practice than the variables collected from administrative data such as claims and electronic health records.

In addition, our findings indicated that Hispanic individuals are less likely to be categorized in the high predicted benefit group. Though we found that race and ethnicity *per se* was unlikely to have been an important determinant of treatment effect heterogeneity and the observed heterogeneity by race and ethnicity was likely to be due to different distribution of age across race and ethnicity groups, it is also possible that some race and ethnicity groups enjoy smaller benefit in other algorithm-based healthcare allocation scenarios. As such, policymakers could use our approach to delineate the characteristics of individuals at risk of not receiving sufficient benefit from the intervention and to make sure they are not marginalized by building strategies that are beneficial to them.[27] Thus, the causal forest approach could help reveal disparities in healthcare by evaluating the heterogeneity in treatment effects across the full spectrum of individual-level demographic and health characteristics. These disparities should be addressed in future studies investigating the relationship between health disparities and algorithm-based healthcare allocation.

Limitations

Our study has limitations. The causal forest can only detect heterogeneity across covariates included in the model, and there may also be other variables not included in the OHIE that drive treatment effect heterogeneity of Medicaid coverage. Second, our study may have limited external generalizability to low-income adults in settings other than Oregon. Third, individuals in the OHIE gained an average of 17 months of Medicaid coverage,[3] and the long-term effects of insurance coverage or the effects of a different type of coverage might be different. Fourth, although lottery assignment was random and thus a good instrumental variable for Medicaid coverage, lottery assignment was not blinded, and thus could potentially have affected mental health directly, moving bias away from the null. Fifth, conclusions on the benefit of Medicaid coverage on outcomes other than depression should not be made based on our study. Sixth, the there is no way to directly address the ethical issues of providing insurance coverage to those with high ITEs. Thus, any policymaker using the high-benefit approach needs to simulate its impact on disparities before implementation and look out for unintended consequences after implementation. Finally, our estimates of the difference in cost-per-case are crude; we use an average cost-per-visit, and visits for those with high ITE may involve different costs or intensity. These stylized figures should thus be interpreted as illustrative.

REFERENCES

1. Finkelstein A, Taubman S, Wright B, et al. The Oregon Health Insurance Experiment: Evidence from the First Year. *The Quarterly Journal of Economics*. 2012-08-01 2012;127(3):1057-1106. doi:10.1093/qje/qjs020

2. Taubman SL, Allen HL, Wright BJ, Baicker K, Finkelstein AN. Medicaid increases emergency-department use: evidence from Oregon's Health Insurance Experiment. *Science*. 2014;343(6168):263-268.

3. Baicker K, Taubman SL, Allen HL, et al. The Oregon experiment—effects of Medicaid on clinical outcomes. *New England Journal of Medicine*. 2013;368(18):1713-1722.

4. Baicker K, Allen HL, Wright BJ, Taubman SL, Finkelstein AN. The effect of Medicaid on management of depression: evidence from the Oregon Health Insurance Experiment. *The Milbank Quarterly*. 2018;96(1):29-56.

5. Mokdad AH, Ballestros K, Echko M, et al. The State of US Health, 1990-2016. *JAMA*. 2018-04-10 2018;319(14):1444. doi:10.1001/jama.2018.0158

6. Poisal JA, Sisko AM, Cuckler GA, et al. National Health Expenditure Projections, 2021–30: Growth To Moderate As COVID-19 Impacts Wane: Study examines National Health Expenditure Projections, 2021-30 and the impact of declining federal supplemental spending related to the COVID-19 pandemic. *Health Affairs*. 2022;41(4):474-486.

7. Sommers BD, Gawande AA, Baicker K. Health insurance coverage and health—what the recent evidence tells us. Mass Medical Soc; 2017. p. 586-593.

8. Wager S, Athey S. Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*. 2018;113(523):1228-1242.

9. Inoue K, Athey S, Tsugawa Y. Machine-learning-based high-benefit approach versus conventional high-risk approach in blood pressure management. *International Journal of Epidemiology*. 2023:dyad037.

10. Kroenke K, Strine TW, Spitzer RL, Williams JB, Berry JT, Mokdad AH. The PHQ-8 as a measure of current depression in the general population. *Journal of affective disorders*. 2009;114(1-3):163-173.

11. Kent DM, Steyerberg E, van Klaveren D. Personalized evidence based medicine: predictive approaches to heterogeneous treatment effects. *Bmj*. 2018;363

12. Davis JMV, Heller SB. Using Causal Forests to Predict Treatment Heterogeneity: An Application to Summer Jobs. *American Economic Review*. 2017-05-01 2017;107(5):546-550. doi:10.1257/aer.p20171000

13. Chernozhukov V, Chetverikov D, Demirer M, et al. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*. 2018-02-01 2018;21(1):C1-C68. doi:10.1111/ectj.12097

14. Chernozhukov V, Chetverikov D, Demirer M, Duflo E, Hansen C, Newey W. Double/debiased/neyman machine learning of treatment effects. *American Economic Review*. 2017;107(5):261-65.

15. Semenova V, Chernozhukov V. Debiased machine learning of conditional average treatment effects and other causal functions. *The Econometrics Journal*. 2021;24(2):264-289.

16. Athey S, Tibshirani J, Wager S. Generalized random forests. *Annals of Statistics*. 2019;47(2):1148-1178.

17. Tibshirani J, Athey S, Wager S, et al. Package 'grf'. Comprehensive R Archive Network; 2018.

18. Rose G. Sick Individuals and Sick Populations. *International Journal of Epidemiology*. 1985-01-01 1985;14(1):32-38. doi:10.1093/ije/14.1.32

19. Everson SA, Maty SC, Lynch JW, Kaplan GA. Epidemiologic evidence for the relation between socioeconomic status and depression, obesity, and diabetes. *Journal of Psychosomatic Research*. 2002-10-01 2002;53(4):891-895. doi:10.1016/s0022-3999(02)00303-3

20. Caron A, Baio G, Manolopoulou I. Estimating Individual Treatment Effects using Non-Parametric Regression Models: a Review. *Journal of the Royal Statistical Society Series A: Statistics in Society*. 2022-07-01 2022;185(3):1115-1149. doi:10.1111/rssa.12824

21. Künzel SR, Sekhon JS, Bickel PJ, Yu B. Metalearners for estimating heterogeneous treatment effects using machine learning. *Proceedings of the National Academy of Sciences*. 2019-03-05 2019;116(10):4156-4165. doi:10.1073/pnas.1804597116

22. Hahn PR, Murray JS, Carvalho CM. Bayesian regression tree models for causal inference: Regularization, confounding, and heterogeneous effects (with discussion). *Bayesian Analysis*. 2020;15(3):965-1056.

23. Oprescu M, Syrgkanis V, Wu ZS. Orthogonal random forest for causal inference. PMLR; 2019:4932-4941.

24. Athey S, Wager S. Estimating treatment effects with causal forests: An application. *Observational Studies*. 2019;5(2):37-51.

25. Mhasawade V, Zhao Y, Chunara R. Machine learning and algorithmic fairness in public and population health. *Nature Machine Intelligence*. 2021-07-29 2021;3(8):659-666. doi:10.1038/s42256-021-00373-4

26. Parikh RB, Teeple S, Navathe AS. Addressing Bias in Artificial Intelligence in Health Care. *JAMA*. 2019-12-24 2019;322(24):2377. doi:10.1001/jama.2019.18058

27. Cintron DW, Adler NE, Gottlieb LM, et al. Heterogeneous treatment effects in social policy studies: An assessment of contemporary articles in the health and social sciences. *Annals of epidemiology*. 2022;70:79-88.

TABLES AND FIGURES

Table 1. Descriptive statistics of controls and lottery winners.

Characteristic	Controls (n=4794)	Lottery winners (n=5274)	Standardized absolute mean difference
Female gender (%)	56.9	55.6	0.02
Age	40.8 ± 11.7	40.8 ± 11.7	0.00
Education (%)			
High school diploma or less	66.6	65.8	0.02
Post-high school	33.4	34.2	
Race and ethnicity (%)			
Non-Hispanic White	63.7	63.7	0.00
Non-Hispanic Black	11.0	10.2	0.03
Hispanic	17.4	17.5	0.00
Other	7.9	8.7	0.03
Interview conducted in English (%)	90.8	90.2	0.02
Diagnosis before lottery (%)			
Hypertension	18.1	17.8	0.00
Diabetes	7.4	6.9	0.02
High cholesterol	12.8	11.8	0.03
Asthma	19.6	18.8	0.02
Heart attack	2.0	1.6	0.03
Congestive heart failure	1.0	1.1	0.02
Emphysema/COPD	2.5	2.3	0.00
Kidney failure	1.8	1.7	0.00
Cancer	4.0	4.1	0.00
Depression	34.3	33.3	0.02
Number of ED visits, pre-randomization	0.8 ± 1.8	0.8 ± 1.8	0.00
Any ED visits for mood disorders, pre-randomization (%)	1.5	1.6	0.00
Any ED visits for psychiatric conditions or substance abuse, pre-randomization (%)	3.2	3.0	0.01
Sum of total charges, pre-randomization (US dollars)	2156 ± 8693	1805 ± 7318	0.00
Sum of ED charges, pre-randomization (US dollars)	893 ± 2439	841 ± 2368	0.00

COPD, chronic obstructive pulmonary disease; ED, emergency department.
Categorical variables are expressed as proportions and continuous variables are expressed as mean ± standard deviation.

Table 2. Comparison of individuals with high vs. low individual treatment effects. The cutoff for high individual treatment effect was set at the median of the ITEs.

Characteristic	Low ITE group (n=5034)	High ITE group (n=5034)	Standardized absolute mean difference
Female gender (%)	58.2	54.3	0.08
Age	30.8 ± 5.6	50.9 ± 6.3	0.35
Education (%)			
High school diploma or less	68.2	64.2	0.09
Post-high school	31.8	35.8	
Race and ethnicity (%)			
Non-Hispanic White	60.9	66.5	0.12
Non-Hispanic Black	9.0	12.1	0.10
Hispanic	22.6	12.3	0.27
Other	7.6	9.0	0.05
Interview conducted in English (%)	87.9	93.0	0.18
Diagnosis before lottery (%)			
Hypertension	7.5	28.4	0.57
Diabetes	2.8	11.6	0.35
High cholesterol	4.9	19.7	0.46
Asthma	19.6	18.7	0.02
Heart attack	0.3	3.4	0.23
Congestive heart failure	0.2	1.8	0.16
Emphysema/COPD	0.4	4.4	0.26
Kidney failure	1.3	2.2	0.06
Cancer	2.0	6.1	0.21
Depression	30.7	36.9	0.13
Number of ED visits, pre-randomization	0.8 ± 1.8	0.8 ± 1.8	0.00
Any ED visits for mood disorders, pre-randomization (%)	1.4	1.7	0.03
Any ED visits for psychiatric conditions or substance abuse, pre-randomization (%)	3.1	3.1	0.01
Sum of total charges, pre-randomization (US dollars)	1523 ± 6335	2444 ± 9432	0.00
Sum of ED charges, pre-randomization (US dollars)	843 ± 2352	892 ± 2455	0.00

ITE, individual treatment effect; COPD, chronic obstructive pulmonary disease; ED, emergency department.
High predicted benefit was defined as ITE greater than the median.
Categorical variables are expressed as proportions and continuous variables are expressed as mean ± standard deviation.

Table 3. The average treatment effect of targeted Medicaid expansion to individuals with high predicted benefit compared to average treatment effect of Medicaid expansion to all individuals in sample.

	High-benefit approach (expanding coverage to individuals with high predicted benefit) (n=5034)	Population approach (expanding coverage to all individuals in sample) (n=10068)	Difference	P-value of difference
ATE on depression (percentage point reduction)	21.5 (9.8, 33.2)	8.8 (0.8, 16.8)	+12.7 (+4.6, +20.8)	0.003
Total annual healthcare spending	$18,002,940	$31,980,514		
Annual healthcare spending per case of depression prevented	$16,627 ($10,775, $36,396)	$36,048 ($18,897, $390,077)	-$18,598 (-$156,953, -$3,120)	0.04

ATE, average treatment effect; ITE, individual treatment effect.
High predicted benefit was defined as ITE greater than the median.
Treatment effects were estimated with instrumental variable regressions, and are expressed as percentage point reduction in prevalence of depression. A positive treatment effect represents decreased risk of screening positive for depression.
The difference in ATEs, its 95% confidence interval, and P-value were obtained using percentile bootstrapping with 10,000 replications.

Figure 1. Distribution of individual treatment effects. Individual treatment effects showed a bimodal distribution, and the median individual treatment effect of 11.6 percentage points was used as the cutoff between low and high individual treatment effects. Treatment effects are expressed as percentage point reduction in prevalence of depression. A positive treatment effect represents decreased risk of screening positive for depression.

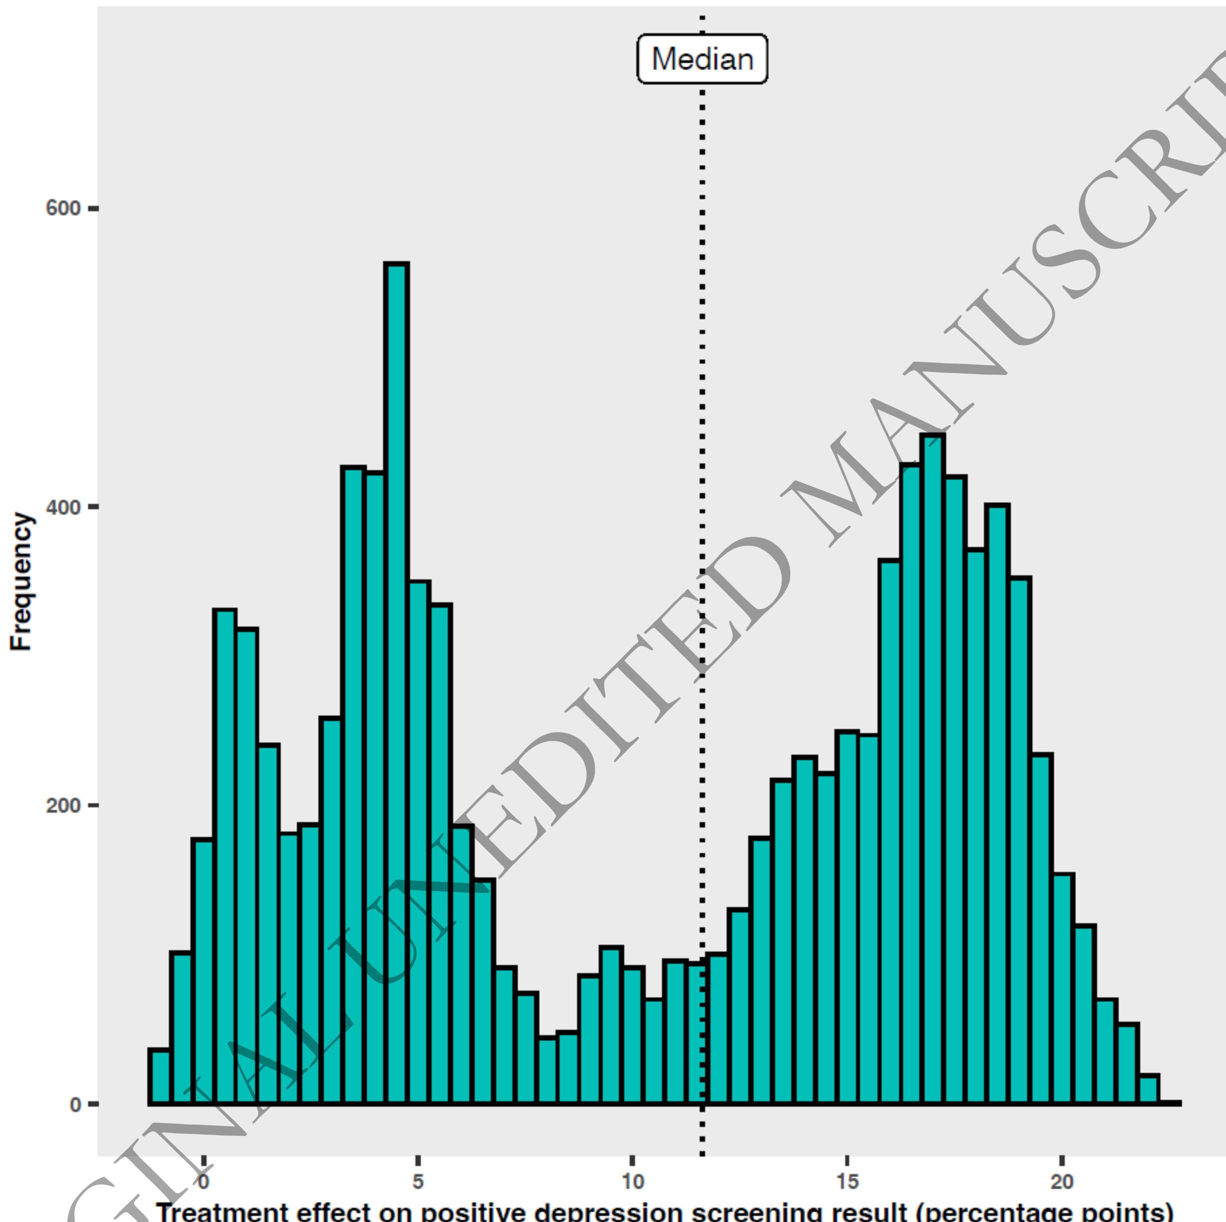

Figure 2. Individual treatment effects across age and number of comorbidities. ITEs are expressed as percentage point reduction. A positive ITE represents decreased risk of screening positive for depression. The x-axis represents the age of the individuals, and the y-axis represents the predicted ITEs. The number of comorbidities is color-coded, as represented in the legend. Two individuals aged 65 or above were excluded from the plot. ITE, individual treatment effect.