

Seeing beyond the Trees: Using Machine Learning to Estimate the Impact of Minimum Wages on Labor Market Outcomes

Doruk Cengiz, *OM Partners*

Arindrajit Dube, *University of Massachusetts Amherst, National
Bureau of Economic Research, and Institute of Labor Economics*

Attila Lindner, *University College London, Institute for Fiscal
Studies, Institute of Labor Economics, and Centre for Economic
and Regional Studies*

David Zentler-Munro, *University of Essex*

We assess the effect of the minimum wage on labor market outcomes. First, we apply modern machine learning tools to predict who is affected by the policy. Second, we implement an event study using 172 prominent minimum wage increases between 1979 and 2019. We find a clear increase in wages of affected workers and no change in employment. Furthermore, minimum wage increases have no effect on the unemployment rate, labor force participation, or labor market transitions. Overall, these findings provide little evidence of changing search effort in response to a minimum wage increase.

We thank Gábor Békés, Erin Conlon, Ezgi Cengiz, Ina Ganguli, Laura Giuliano, Carl Nadler, Hasan Tekgüç, Michael Reich, Jesse Rothstein, and participants at the

Submitted December 12, 2020; Accepted December 9, 2021.

Journal of Labor Economics, volume 40, number S1, April 2022.

© 2022 The University of Chicago. All rights reserved. Published by The University of Chicago Press in association with The Society of Labor Economists and The National Opinion Research Center. <https://doi.org/10.1086/718497>

I. Introduction

A long-standing question in economics centers around understanding how minimum wages affect low-wage labor markets. A key challenge to convincingly answer this question comes from the difficulty in successfully identifying most workers who are actually affected by the policy. While we can easily locate workers who are currently earning the minimum wage, it is difficult to identify all potential workers who also may have been working had the minimum wage been different. This difficulty has led many researchers to focus on specific industries or demographic groups, such as teens (Card 1992; Neumark and Wascher 1992; Giuliano 2013; Neumark, Salas, and Wascher 2014; Allegretto et al. 2017; Totty 2017), younger workers with lower educational credentials (Sabia, Burkhauser, and Hansen 2012; Manning 2016; Clemens and Strain 2017; Clemens and Wither 2019), and individuals without a high school degree (Addison and Blackburn 1999; Addison, Blackburn, and Cotti 2011). However, these groups constitute relatively small shares of all minimum wage workers. As a result, there is a tension between what is often analyzed (e.g., minimum wage effects on teens) and what is argued (effects of the policy on affected workers largely composed of adults; Belman, Wolfson, and Nawakitphaitoon 2015; Manning 2016).¹

In this paper, we use machine learning (ML) tools to predict which individuals were likely affected by minimum wage increases and then estimate the

Labor and Employment Relations Association 70th Annual Meeting, Institute for Research on Labor and Employment Research Presentation seminar, 44th Eastern Economic Association Conference, 2018 New School–University of Massachusetts Economics Graduate Student Workshop, and the authors' conference in honor of Alan Krueger for very helpful comments. We are also grateful to Jon Piqueras for outstanding research assistance. The previous version of the paper was circulated under the title "Seeing Beyond the Trees: Using Machine Learning to Estimate the Impact of Minimum Wages on Affected Individuals." Attila Lindner acknowledges financial support from the Economic and Social Research Council (new investigator grant, ES/T008474/1) and from the European Research Council (ERC) under the European Union's Horizon 2020 Research and Innovation Programme (grant agreement 949995). David Zentler-Munro acknowledges financial support from the Economic and Social Research Council (ESRC) under grant ES/R005745/1. Contact the corresponding author, Arindrajit Dube, at adube@umass.edu. Information concerning access to the data used in this paper is available as supplemental material online.

¹ The discrepancy is particularly relevant when the measured outcome is the teen employment rate, which has been the subject of extensive research in the United States. Belman and Wolfson (2014) consider 30 studies that examined the employment effects of the minimum wage on some demographic groups between 2001 and early 2013, and they find that 17 of them had teen employment as the dependent variable. Neumark (2017) shows that 12 of 13 studies that examined minimum wage effects on "lower-skilled" employment between 2010 and 2016 focused on teens (see his table 1). However, teens are less likely to be in the affected group than nonteen adult minimum wage workers (Lundstrom 2016). Compared with affected nonteens, only a relatively small share of teens live in poverty. According to the 2016 American Community Survey, 18.4% of teens were in households with incomes under the poverty level.

impact of minimum wages on the individuals predicted to be exposed to the minimum wage increase. Our prediction-based approach extends the classification of low-wage workers that was developed by Card and Krueger (1995) in *Myth and Measurement* (see p. 135) and has been undeservedly neglected in the literature ever since.²

We construct various groups according to the predicted probabilities to assess the impact of the policy on workers who are highly likely to be exposed to the minimum wage (whom we refer to as the “high-probability” group) and also a wider group where we can retrieve 75% of likely minimum wage workers (whom we refer to as the “high-recall” group). We then study the impact of the policy on various labor market outcomes, such as employment, unemployment, and labor force participation (LFP) of workers with different exposure to the minimum wage. The impact of the minimum wage on these latter outcomes has been extensively studied in the theoretical literature (e.g., Flinn 2011), despite the scant empirical evidence on them.

Our approach of using a prediction model to classify workers who are likely to be exposed to a minimum wage treatment has several advantages. First, we can assess the effect of the minimum wage on a large fraction of low-wage workers and not just on some specific subgroups with high exposure, such as teens or youth (see, e.g., Laws 2018). In that sense the spirit of our approach is close to Cengiz et al. (2019), who assess the overall employment impact of the policy using the frequency distribution of wages. Additionally, we can directly study the impact of the policy on the affected nonteen (20–64) and prime age (25–55) individuals who are more likely to live in low-income households than teens and tend to be the intended beneficiaries of the policy.

Second, we can also study the impact of the policy on individuals with low predicted probability of being minimum wage workers—that is, those who should not be affected by the policy. Specifications that show an unrealistically large impact on workers not participating in the low-wage labor market should be cautiously interpreted, as they raise questions about the credibility of the particular research design. Using the impact on the “low-probability” group as a falsification test is analogous to studying employment changes in the upper tail of the wage distribution—a fruitful approach that successfully resolved some of the discrepancies in the minimum wage literature (see Cengiz et al. 2019).

Finally, our prediction-based approach allows us to study the impact of the policy on various labor market outcomes that would not be feasible with the

² We are aware of only one previous publication that applied this method (Cengiz et al. 2019), and the list of coauthors includes three of the authors of this paper. That paper utilized the Card and Krueger prediction-based approach primarily to show the differences between that method and the bunching method developed in Cengiz et al. (2019).

distribution-based approach developed in Cengiz et al. (2019). This is a major advantage, as it allows us to provide the first comprehensive picture on how low-wage labor markets evolve in response to the minimum wage. The impact of the policy on unemployment and participation rates, as well as flows in and out of unemployment, are often mentioned in the policy discourse, but evidence on these topics is limited and mainly available for narrow subgroups. Furthermore, as we will demonstrate later, understanding the impact of the policy on outcomes besides employment is relevant, since it has welfare implications in various noncompetitive models (e.g., Flinn 2011).

We implement our approach by using ML methods along with demographic information to predict which individuals are likely to be minimum wage workers in the Current Population Survey (CPS) data between 1979 and 2019. In particular, we use individuals' demographic characteristics to predict their probability of having an hourly wage less than 125% of the statutory minimum wage.³ We consider three tree-based learning tools in the training data: decision trees, random forests, and gradient-boosting trees, as well as the elastic net regularization of a logistic regression. A key advantage of the ML tools over the Card and Krueger (1995) approach is that they do not require the researcher to prespecify the functional form of the prediction model, which is instead determined in a data-driven way. Then we compare the performance of various prediction models on the test data.

The best-performing prediction comes from the gradient-boosting tree model. The original linear prediction model proposed by Card and Krueger (1995; with a judiciously chosen set of interactions) also performs relatively well, although not quite as well as the state-of-the-art ML tools. When compared with commonly used demographic groups in the literature (such as teens or those under the age of 30 with high school or lower educational credentials), the boosting approach can form groups with a similar number of (correctly classified) minimum wage workers while substantially reducing the number of (misclassified) non-minimum-wage workers. The gains in precision (i.e., the share of predicted minimum wage workers who are classified correctly) for a given level of recall (i.e., the share of true minimum wage

³ In the prediction exercise, we restrict the sample to states instituting prominent minimum wage hikes and periods preceding those policy changes. The full set of predictors and how they are coded are reported in app. B (apps. A–E are available online). In the prediction model, we do not use variables related to past employment status or occupation/industry in the CPS Outgoing Rotation Group (CPS-ORG). They are sometimes missing even if the individual is currently in the labor force and looking for a job. We prefer to keep the observations with missing information on these variables in the sample, as they potentially carry information about labor supply effects of the policy. In our preferred prediction model, we do not use state of residence or year information either. This choice is primarily to be able to build samples that are comparable and consistent across time and space.

workers who are classified correctly) are sizeable when we limit attention to nonteen workers—a group that is of particular interest to policy makers.

Armed with the prediction model, we implement an event study analysis that exploits 172 prominent state-level minimum wage increases between 1979 and 2019. We assess the impact of the policy on various groups formed on the basis of the predicted exposure probability. The high-probability group comprises the 10% of the population with the highest likelihood of being affected by the policy. We also study the impact of the policy on the high-recall group, which captures 75% of all minimum wage workers.

For both groups, we find a considerable increase in wages after the policy change; as expected, the wage increase is somewhat lower for the high-recall group. At the same time, we detect a small, positive, statistically insignificant effect on employment for both groups. The implied employment elasticity with respect to own wage—the labor demand elasticity in the standard competitive model of the labor market—is 0.29 (SE, 0.32) for the high-probability group and 0.14 (SE, 0.25) for the high-recall group. The confidence bounds on both of these estimates can rule out anything more than modest negative disemployment effects at the conventional significance levels.

We find no evidence of substantial changes in the unemployment or participation rates in response to the policy. We are also not able to detect any economically meaningful (or statistically significant) changes in labor market transitions between employment, unemployment, and nonparticipation. This lack of response on the LFP margin provides new evidence that minimum wages have a limited impact on search effort when we focus on individuals who are most likely to be affected by the policy.

Our results are robust to controlling for time-varying heterogeneity in a wide variety of ways. Moreover, the increase in wages lines up well with the timing of the minimum wage increases, and the effects emerge only in the group of individuals likely to be exposed to the minimum wage increase. We find no significant differences in labor market outcomes for the low-probability group—suggesting that no unusual changes took place in the states' labor markets around the minimum wage increases we study here. All of these findings reinforce the credibility of our research design.

Furthermore, we also study whether the responses to the policy vary across demographic groups. Most importantly, we study whether differential responses can be detected on employment, unemployment, and participation margins for workers who are thought to have larger extensive margin labor supply elasticities—such as teens, older workers, and single mothers. In addition, we also assess the impact of the policy by the likelihood of moving into or out of the labor force. We use demographic information and apply ML tools again to classify workers as being more likely to move into and out of the labor force. Even when we focus on the group of workers with the highest predicted transition probabilities, we find no evidence of substantial change in the unemployment or participation rates.

This paper contributes to several strands of the minimum wage literature. Although this is a thick literature, there is only a handful of studies that examine broad segments of workers in the United States affected by the minimum wage policy.⁴ Linneman (1982), Currie and Fallick (1996), and Clemens and Wither (2019) study the short-term impact of minimum wages by examining the probability of remaining employed for workers earning below the new minimum wage before the policy change. Assigning workers to groups according to their baseline wages not only requires a richer panel data set than the one used here but also misses the large fraction of minimum wage workers who are entering the labor market in each year.⁵ This is especially problematic when we are interested in understanding changes at the participation margin and the impact on job flows. Moreover, it is difficult to study longer-term effects using this design, as the age composition of the cohort is changing along with time elapsed since the policy change, and the share of the cohort earning close to the minimum wage attenuates over time.

Only a few studies focus on the overall impact of the policy on low-wage jobs. Cengiz et al. (2019) use the frequency distribution of wages to focus on the number of low-wage jobs, while Meer and West (2015) simply consider total state-level employment to study the overall impact of the policy. Our estimates for the high-recall group complement the existing evidence by providing an alternative way of assessing the impact of the minimum wage on overall employment.

Our paper also fills an important gap in the literature by going beyond studying the wage and employment effects of the policy. While there has been a long-standing interest in understanding the impact of minimum wages on the unemployment and participation rates (see, e.g., Mincer 1976; Ragan 1977), there are only a handful of papers examining the impact on these outcomes while applying credible difference-in-differences style estimators. The few exceptions focus mainly on some specific subgroups, such as teens (see, e.g., Wessels 2005; Laws 2018; Marimpi and Koning 2018), parents (see Godoy, Reich, and Allegretto 2020), and workers close to retirement (see, e.g., Borgschulte and Cho 2019). Adams, Meer, and Sloan (2018) study the impact on aggregate job search. Similarly to this paper, they find no indication of significant changes in job search activity. However, it is unclear whether they have enough statistical power to detect significant changes by studying the impact on all US workers, including those with high wages. Since a relatively

⁴ See Belman and Wolfson (2014) for a thorough literature review on the subject.

⁵ Assigning workers according to their baseline wages requires panel data. Panel data sets, such as the National Longitudinal Survey of Youth (Currie and Fallick 1996) and Survey of Income and Program Participation (Clemens and Wither 2019), are often smaller than the CPS applied here and cover fewer years. The prediction probability approach can be applied to cross-sectional data, so we can include many more prominent minimum wage changes in our analysis.

small fraction of the US workforce is typically affected by minimum wage policies, positive wage effects cannot be detected without focusing on workers earning close to the minimum wage (Cengiz et al. 2019). As a result, it may not be surprising that Adams, Meer, and Sloan (2018) were not able to detect significant changes in aggregate job search.

Methodologically, our use of a demographics-based predictive model for minimum wage workers is inspired by Card and Krueger (1995), who examine the 1988 California minimum wage increase and use a linear prediction model to sort individuals living in the state in 1987 according to their likelihood of having hourly wages between the old minimum wage and the new minimum wage (\$3.35 and \$4.25). The Card and Krueger (1995) model is based on subjective judgments about predictors and the functional form, which includes complicated multiway interactions. Even though this subjective assessment turns out to have been implemented incredibly well in Card and Krueger (1995), the key advantage of the ML-based approach proposed here is that we do not need to rely on such judgments. Instead, the ML tools determine the prediction model in a data-driven way and can provide a guarantee against overfitting and specification hunting.

The remainder of the paper is organized as follows. In the next section we briefly discuss the benefits of examining outcomes other than employment. Section III describes the data sets used in our analysis. Section IV explains how we apply various learning tools to predict exposure to the minimum wage. Section V examines the empirical implementation and the key results. Section VI concludes.

II. Participation, Unemployment, and the Welfare Impacts of the Minimum Wage

The approach developed in this paper allows us to analyze the impact of the minimum wage on a wide range of outcomes that go beyond the traditional employment and wage impacts emphasized in the literature. Studying the impact on the unemployment and participation rates is particularly interesting, since a large class of theoretical models have direct predictions on these outcomes. For instance, efficiency wage models (Dräzen 1986), models with information asymmetry (Lang 1987), and search models (Swinnerton 1996) suggest that minimum wages will raise unemployment rates.

The participation margin is also of interest given the growing number of empirical studies that find close to zero effect on employment. Flinn (2006) and Ahn, Arcidiacono, and Wessels (2011) discuss in detail what mechanisms are needed to find positive employment effects in a search model. They point to endogenous participation rate or endogenous search effort as being essential. Flinn (2006) also points out that to simultaneously find an increase in the employment rate and a decrease in the unemployment rate, one needs to introduce two elements: an increase in both the participation rate and the search

effort. Therefore, by providing direct evidence on employment, unemployment, and participation rates simultaneously, we can assess the empirical validity of this search-based explanation.

In addition, the policy impact on labor market participation is potentially informative about the welfare effects on some groups of minimum wage workers. While the employment and unemployment responses often reflect both supply and demand sides of the market, whether someone is participating in the labor market depends solely on the worker's decision. As a result, if some workers choose to participate in the labor market as a result of the policy, they directly reveal that the payoff from searching and finding a job makes them better off in expectation. Assuming that minimum wages do not have an impact on the nonparticipation payoffs, the individuals who decide to participate expect to be made better off by the minimum wage. Conversely, workers who leave the labor force in response to the minimum wage expect to be made worse off.

Our emphasis on studying participation differs somewhat from the literature that interprets the presence of wage spillovers as an indicator of welfare impacts (e.g., Flinn 2002). A key requirement for wage spillovers to be sufficient for assessing welfare changes is the lack of nonwage amenities. If minimum wage policies alter job attributes, the positive (negative) wage spillovers are neither a necessary condition nor a sufficient condition for a welfare gain (loss). This is especially problematic, since some evidence suggests that minimum wages may affect job amenities. For instance, Dustmann et al. (2022) find that the introduction of the German minimum wage led to an increase in commuting times. Clemens, Kahn, and Meer (2018) find that minimum wages alter the provision of the employer-provided health insurance.

The advantage of considering participation decisions is that they are driven by workers' overall assessment of the quality of jobs available (and the probability of getting those jobs). As a result, examining the impact on participation can improve on a piecemeal approach of looking at wage spillovers or specific measures of amenities. Of course, the change in participation is a discrete decision, so the welfare impact of the minimum wage is revealed only for workers who are on the margin of the participation decision. As a result, we also study the impact of the minimum wage on unemployment rates and transition probabilities, which are affected by changes in search intensity. At the same time, we note that these outcomes also reflect changes in labor demand (and are not just workers' decisions), so their relationship to welfare is more complicated.

III. Data

The primary data we use throughout the analysis come from the CPS. We use the 1979–2019 CPS-ORG sample for the hourly wage and weekly earnings variables. This is a subset of the Basic Monthly CPS (CPS-Basic), a monthly survey of approximately 60,000 households in the United States. The

CPS-ORG includes only the fourth and eighth sample months, when usual hourly wages, weekly earnings, and weekly hours worked are asked. These variables are of primary importance for the prediction as well as for the estimation, and thus we rely on their accuracy. For this reason, we exclude observations with imputed hourly wages, imputed weekly earnings, or imputed hours worked. For hourly workers we use the reported hourly wage, and for other workers we define the hourly wage to be their usual weekly earnings divided by usual weekly hours. We also use a range of demographic variables in the data set when predicting an individual's likelihood of having a wage close to the minimum. These variables indicate an individual's age, race, Hispanic status, gender, education, veteran status, marital status, and rural status of the residency (for the exact definitions, see app. B).

We use the 1979–2019 CPS-Basic files for the employment, unemployment, and LFP variables as well as for a number of secondary variables describing the nature of employment (part-time, overtime, and self-employment). Unlike the CPS-ORG, CPS-Basic contains observations for every month that a respondent is surveyed. Therefore, using the CPS-Basic to estimate employment, unemployment, and LFP effects of the minimum wage results in greater precision than using the CPS-ORG to estimate these effects. It also allows us to estimate the impacts of the minimum wage on transitions between employment, unemployment, and inactivity. We obtain the minimum wage data from Vaghul and Zipperer (2016), which has been extended through 2019 by the authors.

IV. Predicting Who Is a Minimum Wage Worker

We build a prediction model to first explain the relationship between being a minimum wage worker (defined as having an hourly wage of less than 125% of the statutory minimum wage) and various demographic variables.⁶ Then we use the model to predict the likelihood of an individual being a minimum wage worker. As the model relies on demographic characteristics, we can ascertain the likelihood of an individual being affected by the policy even if that individual were currently not employed, or had no wage. As a result, we can examine the effects of the policy not only on incumbent workers but also on those who are currently nonemployed but are likely affected by the policy.

We create the following data set to build the prediction model. First, we select all workers in states and quarters that satisfy two criteria: (1) there had not had been any prominent minimum wage events in the past 20 quarters

⁶ Our results are not sensitive to the definition of minimum wage workers based on alternative cutoff values. Setting the thresholds to 3% above the minimum wage or 200% above the minimum wage produces virtually the same ordering of observations according to predicted probabilities, suggesting that the specific definition that we use has essentially no bearing on the conclusions.

and (2) there is a prominent minimum wage change in the next 12 quarters. The former criterion ensures that we are not training the model using workers in states/quarters where the wage distribution may not have stabilized following a minimum wage event. The latter criterion ensures that we are training the model on workers who will experience a minimum wage event in the near future and are therefore pertinent to our analysis. There are 469,174 worker-level observations in the CPS that satisfy this screen between 1979 and 2019.

Second, we divide the 469,174 observations into two mutually exclusive samples: a training sample and a test sample. To create the training sample we randomly draw 150,000 observations. We apply various learning tools, such as random forests, tree boosting, basic logistic, elastic net, and the linear probability model along the lines of Card and Krueger (1995), and fit each model on the training sample. In the next section we describe the key idea behind each prediction algorithm. For further details on these prediction models, we refer the reader to appendix C and Friedman, Hastie, and Tibshirani (2009).

The test sample is composed of the complement of the training sample. We use the test sample to compare the performance of the prediction models by plotting the precision-recall curves (explained below) along with other descriptive statistics.

Once we have optimized over the prediction models, we use the preferred (best-performing) model to calculate the predicted probability in the full data set that includes all time periods and states between 1979 and 2019. We use that full data set to investigate the causal effects of the minimum wage on various predicted probability groups.

A. Prediction Algorithms

Decision Trees

A single decision tree lies at the heart of many learning techniques, including random forests and gradient boosting. A decision tree recursively divides the feature (predictor) space into two in a way that reduces the prespecified loss function the most.⁷ More concretely, in the beginning the algorithm tries every possible split to divide the entire sample space into two and picks the one that diminishes the loss function the most. Subsequently, each subsample is treated as the new sample, and the first step is repeated. Once the splitting is over, it predicts the class of every observation according to the majority vote in the subspace (terminal node) to which the observation belongs.

This procedure requires a decision on when to stop the splitting. In principle, the splitting could continue until there is only one data point at each

⁷ The loss function is the deviance, defined as $-2\sum_m \sum_k n_{mk} \times \log(\hat{p}_{mk})$, where n_{mk} indicates the number of observations at terminal node m that belongs to class k and \hat{p}_{mk} is the share of observations at terminal node m that belongs to class k .

terminal node. Such a tree fits the training sample perfectly but would suffer from overfitting. To overcome the problem, it is common to use cross validation to determine the complexity of the tree. For a more accurate prediction, we collapse some internal nodes (“prune the tree”) and decrease the prediction variance at the expense of bias.

Decision trees are not among the most successful learners, yet they are relatively easy to interpret. In figure 1 we plot a pruned decision tree produced to predict whether a worker has an hourly wage of less than 125% of the statutory minimum wage using demographic and educational characteristics. The tree predicts that the only group in the training sample with hourly wages less than the threshold is the one with those who are 19 years old or younger. The majority vote in all of the other terminal nodes is “false,” indicating that nonteen observations are expected to work for hourly wages higher than the threshold.

It is noteworthy that the recommendation based on a simple decision tree is to proxy minimum wage workers with teens—which happens to be the most common approach taken in the literature. However, as we show below, it is possible to obtain much better predictions by combining multiple decision trees. The two most common ways to do so are the random forest by Breiman (2001) and the gradient-boosting trees by Friedman (2001).

Random Forest

The random forest is a tree-based ensemble learning technique. It provides a way to overcome the bias-variance trade-off of a single tree. It constructs a multitude of fully grown decision trees formed using different training bootstrap samples. Each tree produces unbiased predictions that have large variances. We calculate the average of the predictions, thereby reducing the variance. To further reduce the variance, we decrease the correlation among trees by employing a randomly selected portion of the predictors at each split. Although this results in the loss of the interpretability of individual trees, it has no impact on the bias, since individual trees are still fully grown.⁸ Our fivefold cross validation finds that the optimum random forest is achieved with 2,000 trees and only two predictors tried at each split.

Boosting

Boosting approaches the problem of how to combine multiple trees from a different angle. Instead of producing many fully grown trees and averaging them, the trees in this model are grown sequentially where subsequent trees attempt to fix the errors of the preceding ones. As a result, while the first tree

⁸ Note that if trees are perfectly correlated, the reduction of the variance would be nil. If they are independent, the variance of the final model would be σ^2/B , where σ^2 is the prediction variance of a single tree and B indicates the number of trees.




FIG. 1.—Minimum wage workers according to pruned trees. The figure plots a pruned decision tree produced to predict whether an individual is a minimum wage worker (has an hourly wage of less than 125% of the statutory minimum wage) using demographic and educational characteristics. In the beginning the tree tries every possible split to divide the entire sample space into two and picks the one that diminishes the loss function the most. Then each subspace is treated as the new feature space, and the first step is repeated. “True” indicates that the tree predicts that workers in the terminal node are minimum wage workers and “false” indicates otherwise. While the tree explores characteristics such as gender, marital status, veteran status, and rural residency status, it picks only age and education to make splits. LTHS = less than high school; HSG = high school graduate; SC = some college; CG = college graduate.

in the boosting is interpretable, the subsequent trees are not independently meaningful. Intuitively, with the boosted trees one starts with the lowest-hanging fruit and, say, classifies teens as minimum wage workers and others as not minimum wage workers. Then the second or subsequent tree builds a model that focuses more on correctly classifying nonteen minimum wage workers and teen non-minimum-wage workers, namely, the observations misclassified by the first tree. The change of focus is usually achieved by altering

the outcome variable (e.g., using the residual as the outcome variable) or slightly changing the loss function (weighting the misclassified observations more heavily). After building the subsequent tree, we combine the predictions of all trees through a weighted majority vote. On the basis of our fivefold cross validation, the optimum boosted tree model is obtained with the following parameters: number of trees = 4,000, shrinkage factor = 0.005, depth of tree = 6, and minimum observations in a node = 10.

Elastic Net

We use the elastic net regularization developed by Zou and Hastie (2005). The underlying model is very similar to the logistic regression except that the elastic net model penalizes model complexity. The penalty term is a linear combination of the lasso and ridge methods; lasso tends to drop poor predictors, while ridge tends to shrink their coefficients toward zero, so elastic net combines both. As opposed to tree-based models, the elastic net regression requires prespecification of the exact functional form for the predictors in the prediction equation. Therefore, we purposefully build a fairly complex model, where we include all of the predictors, their four-way interactions, and all of the interactions with the quadratic, cubic, and quartic terms of the age variable. We rely on the regularization to simplify the model and prevent overfitting.

Card and Krueger's Linear Probability Model

This is a trial and error method employed by Card and Krueger (1995). We follow the functional form proposed on page 135. They estimate a linear probability model using the following independent variables: a set of three-way interaction variables between teen, nonwhite, and gender indicators; three-way interaction variables between young adult (aged 20–25), nonwhite, and gender indicators; three-way interactions of age, categorical education, and gender variables; quadratic and cubic terms of the age variable; and indicator variables for Hispanic and nonwhite individuals. We can think of the Card and Krueger model as a sort of lasso approach to predictor selection but where the regularization is based on subjective judgment instead of a formal learning algorithm.⁹

B. Precision-Recall Curves and Predicted Probabilities

To compare the models with each other, we employ two concepts from the ML literature: precision and recall. “Precision” refers to the share of those we classify as being in the predicted group who are true minimum wage workers.

⁹ We also tried to implement neural networks and support vector machines. While the model constructed using the neural networks performs slightly worse than the boosted tree model, the models using the support vector machines fail to provide a well-performing prediction model.

“Recall” refers to the share of true minimum wage workers who we correctly classify as being in the predicted group. For instance, if a predicted group has only one observation and the observation is a true minimum wage worker, then the precision is 1; however, here the recall is very small, as the sample will cover only a minuscule fraction of the minimum wage workers in the population. On the other hand, if the predicted group contains every observation in the population, then the recall rate is 1, as the sample, by construction, includes all minimum wage workers in the population. However, here the precision is going to be small, since the predicted group also includes all the non-minimum-wage workers in the population. The ideal is to construct a predicted group that includes all the minimum wage workers and none of the non-minimum-wage workers so that both the precision and the recall are 1. Generally, the higher the precision for a given recall rate, the better the performance of the model.¹⁰

Figure 2A shows the precision-recall curves corresponding to the various prediction algorithms. We also estimate and report the performance of a basic logistic model with age and the categorical education variables as predictors for comparison. To plot the curve, we calculate the predicted probabilities for each individual in the test sample. We then define the predicted group for alternative probability thresholds, where all workers in the group have a predicted probability greater than the threshold. We calculate the precision and the recall for each of these groups and obtain the curve. In other words, each point on the curves corresponds to a separate predicted group. When we raise the threshold, we expect the precision to increase but at the cost of a reduced recall rate. How strong this trade-off is between the precision and recall rates for various prediction models is shown in the figure.

The figure shows that the boosted tree model (solid line) outperforms other prediction models, since it provides the highest precision at almost all recall levels. For comparison, in figure 2B we report the other prediction models relative to the boosted tree model. The boosted tree model (and also the other prediction algorithms) improves precision considerably relative to the basic logistic model. Nevertheless, the differences between the other prediction models and the boosted tree model are relatively small, especially at higher recall rate levels. The random forest model achieves almost the same result as the boosted tree model. It is also notable that the Card and Krueger subjective judgment approach does almost as well as the elastic regularization of the logistic model, and the performance of their model is not far behind the best-performing prediction model.¹¹

¹⁰ Another approach commonly used to compare models is to plot the receiver operating characteristic (ROC) curve. The ROC curve plots the recall against the false positive rate, the latter defined as the number of non-minimum-wage observations as a proportion of the number of non-minimum-wage workers in the population. In our case, we reach the same conclusion whether we use the ROC curve or the precision-recall curve.

¹¹ An alternative way to assess model performance is to compare the fraction of true minimum wage workers in each predicted probability decile. Table A.1 (tables A.1–A.5,

In figure 3 we compare the performance of the best prediction model, the boosted tree model, relative to the strategy of choosing specific subgroups to proxy minimum wage workers.¹² First, in terms of precision, the teen sample performs better than all of the other commonly used samples that we compare (workers younger than 30 with less than high school education [LTHS, age < 30], workers younger than 30 with high school or less education [HSL, age < 30], and LTHS workers). In fact, when compared with LTHS, age < 30, both recall value and precision values of the teen sample are higher. This indicates that the teen sample includes more workers that truly have hourly wages lower than 125% of the minimum, and it captures them more accurately than the former samples. Second, the commonly used samples that include nonteen workers tend to achieve a higher recall than the teen sample. However, the rise in the recall is expensive in terms of the lost precision. For instance, including all HSL workers younger than 30 increases the recall value by 0.124 compared with the teen sample, yet the precision decreases by 0.256. It implies that many nonteen observations in the HSL, age < 30 sample are actually non-minimum-wage workers.

Overall, it is clear that points on the curve are closer to the top-right corner than the points corresponding to the commonly used samples. Furthermore, the difference in the precision values between the commonly used samples and the samples recommended by the tools increases as the recall increases. For instance, the precision value that the learning tools achieve for the teen sample's recall value is only slightly greater than that of the teen sample (the vertical distance between the black triangle and the curve). For the recall value of the HSL, age < 30 sample, however, the learning tools achieve a substantially larger precision value than that of the former sample. Therefore, the figure highlights that the learning tools improve the precision-recall trade-off considerably, especially if the aim is to include nonteen observations.

C. Who Are the Minimum Wage Workers?

Before we study the impact of the minimum wage on labor market outcomes, we examine who are the most likely minimum wage workers according to our (best-performing) prediction model. To do this, we examine the

C.1, E.1 are available online) shows that the boosted tree model has a slightly higher fraction of true minimum wage workers in the most likely predicted deciles and a lower fraction in the least likely predicted deciles. This provides further support of the slightly better performance of the boosted tree model.

¹² Of course, it is possible that someone is directly interested in the impact of the policy on the labor market outcomes of certain demographic groups or industries. Nevertheless, in most cases researchers pick specific subgroups (e.g., teens) or sectors (e.g., restaurants) not because they are the main subjects of interest but because these are subgroups where the fraction of minimum wage workers is high. Furthermore, the prediction approach can also be applied if someone is specifically interested in the impact of the minimum wage on some subgroups (see table 5).

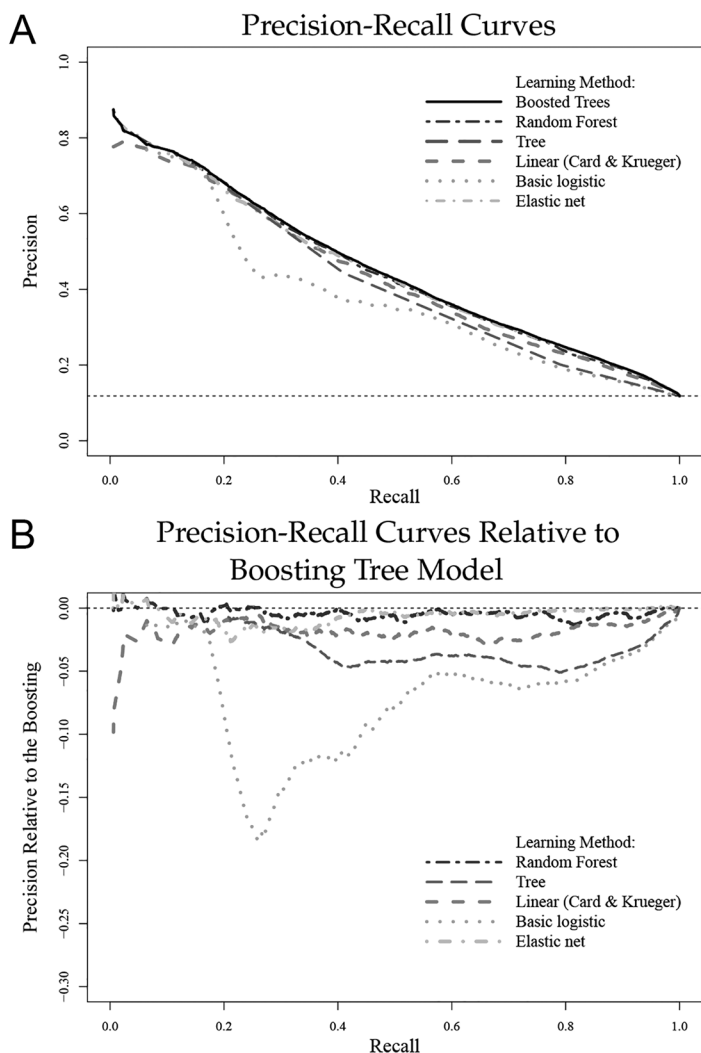


FIG. 2.—Precision-recall curves. *A* plots the precision-recall curves for various prediction models described in section IV.A and for a basic logistic model that we estimate using (linear) age and categorical education variables. We obtain the precision-recall curves in the following way: we use our prediction model to calculate the probability that someone is a minimum wage worker, and we assign all individuals to the predicted group if that probability is above a certain threshold. The figure shows the estimated precision and recall rates obtained when we vary the threshold value. The horizontal dotted line shows the average share of minimum wage workers in the sample. The areas below the precision-recall curves are 0.449 for boosted tree, 0.445 for random forest, 0.443 for elastic net, 0.435 for Card and Krueger’s linear probability model, 0.342 for basic logistic, and 0.269 for single tree. *B* shows the difference in precision rate between the best-performing model—the boosted tree model—and the other models at each recall rate. A color version of this figure is available online.




FIG. 3.—Performance of the boosted tree model relative to demographic subgroups. The figure compares the performance of the best-performing prediction model—the boosted tree model—relative to the strategy of choosing specific subgroups to proxy minimum wage workers. The solid line shows the precision-recall curve for the boosted tree model. The black triangle shows the precision and recall level for teens, the black circle shows the level for individuals with less than high school (LTHS) education, the gray circle shows the level for LTHS individuals younger than 30, and the gray triangle shows the level for individuals with high school or less (HSL) education younger than 30. The horizontal dashed line shows the average share of minimum wage workers in the sample.

main characteristics of individuals in various predicted probability deciles. Table 1 shows the share of workers with various demographic characteristics in each column, while predicted probabilities deciles are in descending order (starting with the highest-probability group).

Examining the table, we observe that age is a highly important predictor. Teens constitute 71.9% of the workers in the decile most likely to be exposed to the minimum wage. This suggests that an analysis using only the highest decile would be very similar to an analysis based on just the teen sample. However, the teen share drops to 4.7% in decile 9 and is virtually zero thereafter, and more than 34% of the observations in deciles 9 and 8 are adults aged 20–30. This suggests that the importance of the age variable is not limited to determining whether an observation belongs to the highest decile. The model also heavily employs educational attainment in determining the predicted probabilities. Observations with less than high school education are mostly in the highest deciles, and high school graduates are concentrated on the top half of the exposed individuals. There is almost no observation

Table 1
Demographic Characteristics for Each Predicted Probability Decile

	Teen (1)	20 ≤ Age < 30 (2)	LTHS (3)	HSG (4)	Female (5)	White (6)	Black or Hispanic (7)
Most likely decile	.719	.038	.752	.145	.592	.837	.244
Probability decile 9	.047	.405	.534	.238	.674	.847	.359
Probability decile 8	.004	.341	.344	.437	.594	.834	.243
Probability decile 7	.004	.298	.187	.575	.571	.833	.351
Probability decile 6	.000	.191	.085	.660	.673	.873	.150
Probability decile 5	.000	.187	.100	.475	.492	.784	.253
Probability decile 4	.000	.178	.067	.236	.512	.794	.237
Probability decile 3	.000	.162	.004	.297	.404	.865	.175
Probability decile 2	.000	.088	.000	.143	.385	.848	.122
Least likely decile	.000	.015	.000	.039	.314	.741	.134

NOTE.—The table shows some demographic characteristics for each predicted probability decile. The predicted probability refers to the probability that an individual has an hourly wage lower than 125% of the minimum wage preceding the minimum wage hike and is calculated according to the best-performing prediction model—the boosted tree model. Each row shows the average characteristics of individuals in the particular predicted probability decile. The top (bottom) row shows the characteristics at the top (bottom) decile, which consists of individuals that are most (least) likely exposed to the minimum wage according to our prediction model. Each cell shows the share of the selected demographic group: col. 1, the share of teens (i.e., those younger than 20); col. 2, the share of individuals who are between 20 and 30 years of age; col. 3, the share of individuals with less than high school (LTHS) education; col. 4, the share of high school graduates (HSGs) with no college education; col. 5, the share of females; col. 6, the share of white individuals; and col. 7, the share of Black or Hispanic individuals.

without some college education in the lowest deciles. Another finding worth noting is that the share of women workers is high in the top deciles (e.g., 67.4% in decile 9) and is lower in the bottom deciles (31.4% in the least likely decile), indicating that an individual’s gender also plays an important role. Last, Black/Hispanic individuals’ share in the least likely decile is 13.4%, whereas they make up at least 24% of the top two deciles.

An alternative way to examine who are the minimum wage workers is to consider the relative importance of each predictor. In figure 4 we plot the “relative influences” of the variables calculated following Friedman (2001).¹³ The

¹³ The figure shows the reduction in the loss function caused by each variable used in the nonterminal nodes. We normalize the relative influences so that they sum up to 100. The average importance of each variable is

$$\mathcal{I}_l^2 = \frac{1}{M} \sum_{m=1}^M \mathcal{I}_l^2(T_m),$$

where $\mathcal{I}_l(T_m)$ is the reduction in the loss function due to the use of variable l in the nonterminal nodes of tree m . However, we wish to caution against interpreting them directly. First, the importance is in terms of prediction, not explanation. Second, there are cases where one variable needs to be interacted with another one for high predictive power. In those cases, only one of them is deemed to have a strong influence, whereas both are essential.




FIG. 4.—Relative influences of the predictors in the boosted tree prediction model. We plot relative influences of the variables in the best-performing prediction model—the boosted tree model—calculated as in Friedman (2001; for details, see n. 13). The bars, which indicate the decline in the loss function associated with the corresponding variable, are normalized so that they sum up to 100.

figure largely confirms our previous observations. It finds age as the most important predictor in the sample with a very large margin. The variable for educational credentials comes after age.¹⁴ Gender variables are also relatively important in the prediction. The indicator variables for Hispanic, rural, race, and veteran status appear to have less influence on the prediction.

D. High-Probability and High-Recall Groups

While we show responses to the minimum wage change throughout the whole predicted probability spectrum, in most of the paper we will focus on two key subgroups in the baseline specifications. For the first group, we follow Card and Krueger (1995) and define the high-probability group as consisting of the top 10% of individuals in terms of predicted probability of being minimum wage workers. When we use the boosted tree prediction model, the threshold probability that we need to apply to get this group is 35%—so all individuals with predicted probabilities above this value are assigned to the high-probability group. At this threshold, the precision rate is around 60%, which means that around 60% of the individuals in the high-probability group

¹⁴ In fact, dropping teen observations from the sample decreases the relative importance of the age variable substantially. It renders age to be the close second-most important variable in the prediction. The educational credentials variable of the observation becomes the most important predictor.

are indeed minimum wage workers. The associated recall rate is 36%, which means that this group covers around 36% of all minimum wage workers.

Since the high-probability group covers just over a third of all minimum wage workers, we also study the impact of the minimum wage on a more broadly defined group. In the high-recall group, we set a threshold probability such that 75% of all minimum wage workers are captured. In practice, this leads us to set the threshold at 12%; at that level we achieve a 35% precision rate. This high-recall group covers just more than 40% of all workers in the data. To study the impact of the policy on workers unlikely to be affected by the policy, we also define a group for whom the predicted probability is less than 12%. Throughout the paper, we refer to that group as the low-probability group.

V. Impact of the Minimum Wage on Labor Market Outcomes

Identification Strategy

We estimate the effect of the minimum wage by implementing an event study strategy similar to the baseline specification in Cengiz et al. (2019). We exploit state-level variation in the minimum wage and compare labor market outcomes in the treated states with the labor market outcomes in states without any minimum wage hike. The event-based approach that we apply here is similar to that of Autor, Donohue, and Schwab (2006).¹⁵ We focus on changes within an 8-year window around 172 prominent state-level minimum wage events instituted between 1979 and 2019. Prominent minimum wage changes are those where the (real) minimum wages increased by more than \$0.25 and where at least 2% of the workforce earned between the new minimum wage and the old minimum wage.¹⁶

Our basic regression specification is as follows:

$$Y_{st}^g = \sum_{\tau=-3}^4 \beta_{\tau} treat_{st}^{\tau} + \Omega_{st} + \mu_s + \rho_t + u_{st}, \quad (1)$$

where Y_{st}^g is the labor market outcome (e.g., employment rate, unemployment rate, participation rate) in state s and at quarter t for group g . Unlike when estimating the prediction model, here we use all states and quarters available in the CPS data. As we discussed above, we study the impact of the minimum wage on various groups defined by the prediction model, such

¹⁵ Appendix G in Cengiz et al. (2019) shows how this event study approach is related to alternative methods applied in the literature like the two-way fixed effects estimator with log minimum wage.

¹⁶ We show the graphical distribution of prominent minimum wage changes and the number of such changes in each year in fig. A.1 (figs. A.1–A.6, C.1, C.2 are available online). In fig. A.2 we plot the change in state-level log (real) minimum wage following a prominent minimum wage increase.

as the high-probability group and the high-recall group. Here, $treat_{st}^{\tau}$ is a binary variable that takes on the value of 1 if the minimum wage was increased τ years from date t in state s . This definition implies that $\tau = 0$ represents the first year following the minimum wage increase (i.e., the quarter of treatment and the subsequent three quarters), and $\tau = -1$ is the year (four quarters) before treatment. Our benchmark specification controls for state and period fixed effects, μ_s and ρ_t , and we also include controls for small or federal increases, Ω_{st} .¹⁷ We cluster our standard errors by state, which is the level at which policy is assigned.

Our baseline approach uses staggered variation of minimum wage increase. As shown in many recent papers, this can lead to negative weighting bias in the presence of heterogeneous treatment effects (see, e.g., Sun and Abraham 2020). To alleviate these concerns, in table E.1 we present estimates with a stacked regression approach following Cengiz et al. (2019). In this approach, we align events by event time (and not calendar time) and use only within-event variation (between the treated unit and clean control units), which is equivalent to a setting where all of the events happened all at once and were not staggered. Gardner (2021) derives the implicit weights for such a stacked regression approach and shows that they do not suffer from negative weighting.

Main Results

Table 2 shows the estimated effects on the high-probability, the high-recall, and the low-probability groups. We report 5-year averaged posttreatment estimates for the key labor market outcomes (wages, employment, unemployment, and LFP) relative to the $\tau = -1$, formally $(1/5)\sum_{\tau=0}^4(\beta_{\tau} - \beta_{-1})$. Columns 1 and 3 establish that the minimum wage has a significant positive impact on wages for groups of workers predicted to be exposed to the minimum wage. In the high-probability group, wages increased by around 2.3% (SE, 0.3%), while in the high-recall group—which captures 75% of the minimum wage workers—wage increase was a little smaller but still significant (1.6%; SE, 0.3). In contrast, column 5 shows no indication of any significant wage effects for the low-probability group. This confirms that the wage growth occurred only for workers exposed to the minimum wages and not for individuals unlikely to be directly exposed to the minimum wage shock.

Columns 2, 4, and 6 of table 2 show analogous estimates but classifying workers according to the Card and Krueger linear probability model.¹⁸ It is

¹⁷ The variables we use to control for federal and small events are the same as the ones employed in Cengiz et al. (2019). We collapse the windows for small and federal events into three periods: EARLY, PRE, and POST. EARLY is for 3 and 2 years before, and PRE is for 1 year before the event. POST is for 0–4 years after the event.

¹⁸ Table A.2 shows that the boosted tree prediction model picks a sample that is a bit older, more educated, more female, and less white than the sample selected by the Card and Krueger model.

Table 2
Impact of the Minimum Wage on Labor Market Outcomes

	(1)	(2)	(3)	(4)	(5)	(6)
Δ wage (%)	.023*** (.003)	.025*** (.004)	.016*** (.003)	.014*** (.003)	-.001 (.003)	-.001 (.003)
Δ employment (pp)	.002 (.002)	.001 (.002)	.001 (.001)	.001 (.001)	.001 (.001)	.000 (.001)
Δ unemployment (pp)	-.001 (.001)	-.001 (.001)	-.000 (.000)	-.001* (.000)	-.000 (.000)	.000 (.000)
Δ participation (pp)	.002 (.002)	.001 (.002)	.000 (.001)	.001 (.001)	.001 (.001)	.001 (.001)
Employment elasticity with respect to minimum wage	.071 (.076)	.036 (.066)	.020 (.038)	.028 (.032)	.012 (.011)	.006 (.013)
Employment elasticity with respect to wage	.286 (.316)	.138 (.245)	.114 (.216)	.191 (.211)	NA	NA

Number of events	172	172	172	172	172	172	172	172
Number of observations	7,854	7,854	7,854	7,854	7,854	7,854	7,854	7,854
Number of individuals in sample	6,639,492	5,812,367	20,917,455	24,737,455	29,370,470	25,511,485		
Mean employment	.338	.399	.415	.439	.741	.771		
Mean unemployment	.058	.064	.046	.043	.030	.030		
Mean participation	.395	.462	.460	.481	.771	.801		
Group	High probability	High probability	High recall	High recall	Low probability	Low probability		
Prediction model	Boosted tree	CK linear	Boosted tree	CK linear	Boosted tree	CK linear		

NOTE.—The table reports the effects of the minimum wage on labor market outcomes according to the event study analysis (see eq. [1]) using 172 state-level minimum wage changes between 1979 and 2019. The table reports 5-year averaged posttreatment estimates for each key labor market outcome: percent change in wages and the change in employment to population, unemployment to population, and labor force participation rate. We also report the employment elasticity with respect to the minimum wage and the employment elasticity with respect to the wage, which is the ratio of the percent change in employment and wage. To calculate the percent change in employment, we divide the change in employment to population by the mean employment to population rate preceding the minimum wage hikes (reported at the bottom of the table). “Number of observations” refers to the number of quarter-state cells used for estimation, while “Number of individuals” refers to the underlying CPS sample used to calculate labor market outcomes in these cells. Columns 1 and 2 show estimates for the high-probability group, which captures the 10% of the population with highest predicted probability. Columns 3 and 4 show estimates for the high-recall group, which consists of individuals whose predicted probability is above 12%—a threshold that leads to a 75% recall rate of minimum wage workers. Columns 5 and 6 show the estimates for workers whose predicted probability is below 12%. Columns 1, 3, and 5 use the best-performing prediction model—the boosted tree model. Columns 2, 4, and 6 use the Card and Krueger (CK) linear prediction model (here, the high-recall group is defined by having a predicted probability, using the linear prediction model, above 12%, which again generates a 75% recall rate). All regressions are weighted by state-quarter population. Robust standard errors in parentheses are clustered by state. pp = percentage point.

* $p < .10$.
 *** $p < .01$.

worth noting that the wage effects are almost the same for the best-performing prediction model and for the Card and Krueger approach and suggests that the Card and Krueger model performs quite well in this setting.¹⁹

Table 2 also reports the effect of the minimum wage on employment. Considering the high-probability and high-recall groups, we find a small and statistically insignificant positive effect on employment. The employment elasticity with respect to minimum wage is around 0.07 (SE, 0.08) and 0.02 (SE, 0.04), respectively. The 90% confidence intervals around these estimates can rule out an elasticity of -0.1 , the lower bound (in magnitude) of the range suggested by Neumark and Wascher (2008).

Table 2 also shows that the employment effects are somewhat smaller for the high-recall group than for the high-probability group, which is in line with the wage effects. This leads to a similar elasticity of employment with respect to own wage, which would be the labor demand elasticity in the neoclassical model, in the two groups. When we calculate the employment elasticity with respect to own wage, we obtain an elasticity of 0.29 (SE, 0.32) for the high-probability group and 0.11 (SE, 0.22) for the high-recall group. The estimates are quite precise, especially those from the high-recall group, which can rule out all but a modest negative impact of the policy on employment.

A key advantage of the probability-based approach is that we can study outcomes other than employment and wages. In table 2 we also report the effect of the minimum wage on unemployment rate and on participation rate. For the high-probability group (col. 1), we find a slight decrease in unemployment and a slight increase in the participation rate. Importantly, however, none of these changes in unemployment and participation rates are statistically significant. The estimates for the high-recall group (col. 3) show no change in either the unemployment rate or the participation rate.

The estimated slight decrease in unemployment or the slight increase in participation (or just unchanged values of both) suggests that search effort is unlikely to fall in response to the policy. This set of empirical findings is difficult to reconcile with a Flinn-type search-and-matching model (Flinn 2006), where adjustments on the participation margin play a vital role, or with models predicting a considerable increase in the unemployment rate in response to the policy (Drazen 1986; Lang 1987; Swinnerton 1996). Furthermore, as we discussed in section II, the lack of a visible drop in participation implies that the policy did not have a significant negative effect on the welfare of workers at the margin of LFP, even as it raised wages for inframarginal workers.

¹⁹ However, the key advantage of applying ML tools is that someone can select the predictors in a data-driven way without knowing much about the context. Even if the functional form chosen by Card and Krueger (1995) performs very well, it is unclear how someone with less knowledge about US labor markets could come up with that functional form. Moreover, there is no guarantee that it would perform well in all contexts.

Nonlinearity in the Extent of Exposure

Table 2 shows that the high-probability group (which captures the 10% most exposed individuals) and the high-recall group (which covers a broader group comprising 75% of the minimum wage workers) provide strikingly similar estimates across the subgroups studied here. This suggests that the impact of the policy does not seem to depend on nonlinearity in the extent of exposure. Figure 5 explores this nonlinearity in more detail. We show the estimated impact on the key labor market outcomes by applying various cut-offs in the predicted probability. The solid lines show the 5-year averaged posttreatment estimates for individuals whose predicted probability is above the particular predicted probability threshold. As this threshold decreases, we also naturally decrease the precision, as our group will include more non-minimum-wage workers in the sample. At the same time, we also capture a larger fraction of minimum wage workers, so we attain a higher recall rate.

Figure 5A shows that the wage effects tend to decline as we lower the minimum predicted probability. This is what we would expect given the declining precision rate: as we lower the threshold, our sample will cover more and more non-minimum-wage workers whose wages are unaffected (directly) by the policy. Figure 5B shows that employment responses, which start from a small positive effect, decline as we lower the precision. As a result, when we divide the employment effects by the wage effects, we obtain a stable employment elasticity with respect to own wage.

Furthermore, figure 5C shows that the unemployment rate declines at the high-probability group but that the decline shrinks as we increase the recall rate. The estimated change in participation rate, which is shown in figure 5D, is close to zero and unrelated to the recall rate. Overall, the graphical evidence in figure 5 finds no clear indication for nonlinearities in response to the policy.

Timing of the Impact

Figures 6 and 7 show the impact of the minimum wage on various labor market outcomes over time for the high-recall and high-probability groups, respectively. In the figures we plot β_τ , expressed relative to the event date $\tau = -1$, or the year just before treatment. Panel A shows the evolution of wages around the minimum wage increase. There is a clear increase in wages in line with the timing of the minimum wage increase, and this is not driven by preexisting trends. Over time the wage effects are somewhat attenuated, which reflects that the most recent minimum wage changes tend to be larger.²⁰

²⁰ Since we do not fully observe the impact of the policy 5 years after the minimum wage increase for the most recent events, those events impact the estimates only in earlier posttreatment years. In figs. A.4 and A.5 we assess the timing of the policy when we focus only on minimum wage changes where we see responses for the entire event window. For these events with a balanced sample, which are not subject to the composition effect, we find no decline in wage effects over time.

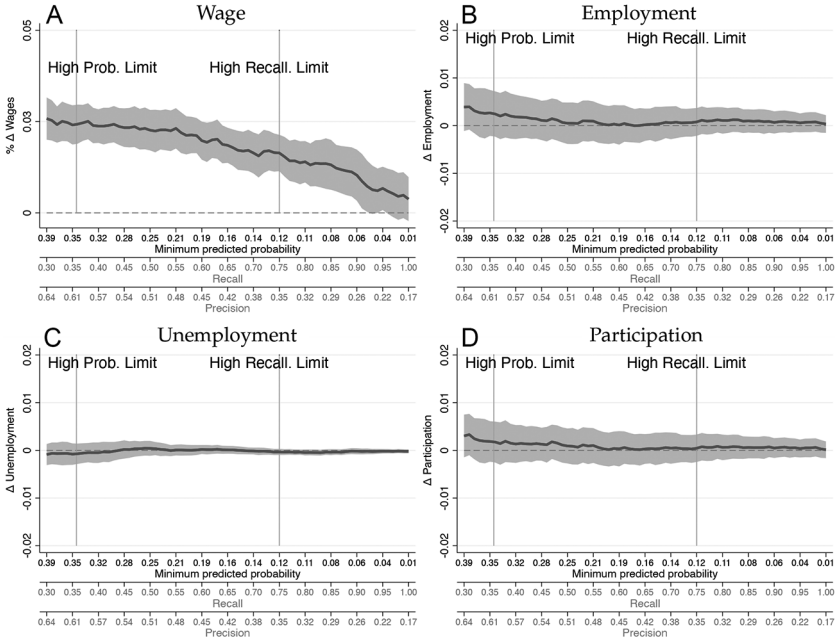


FIG. 5.—Impact of the minimum wage for alternative predicted probability threshold values. The figure shows the main results from our event study analysis (see eq. [1]) using alternative predicted probability threshold values. We exploit 172 state-level minimum wage changes between 1979 and 2019. The figure shows the effect of a minimum wage increase on wages (A), on employment to population (B), on unemployment to population (C), and on labor force participation rate (D). In each panel the solid line shows the 5-year averaged posttreatment estimates for individuals whose predicted probability is above the minimum predicted probability threshold. On the x-axis we also report the corresponding recall rate (the fraction of minimum wage workers retrieved by the prediction model if the particular minimum predicted probability threshold is applied) and the precision rate (the fraction of minimum wage workers in the predicted group if the particular minimum predicted probability threshold is applied). We also plot the thresholds corresponding to the high-probability group capturing the 10% of the population with the highest predicted probability and to the high-recall group capturing 75% of all minimum wage workers. To calculate the predicted probabilities, we use the best-performing prediction model—the boosted tree model. The shaded areas show the 95% confidence interval based on standard errors that are clustered at the state level. A color version of this figure is available online.

Panel B shows the impact of the policy on employment around the minimum wage hike. For both the high-recall group and the high-probability group, we see a similar pattern: there is no clear evidence of preexisting trends in employment, although there is a slight dip in employment 2 years before the minimum wage increase when we look at the high-recall group (fig. 6). Nevertheless, there is no unusual employment change if we look at the longer

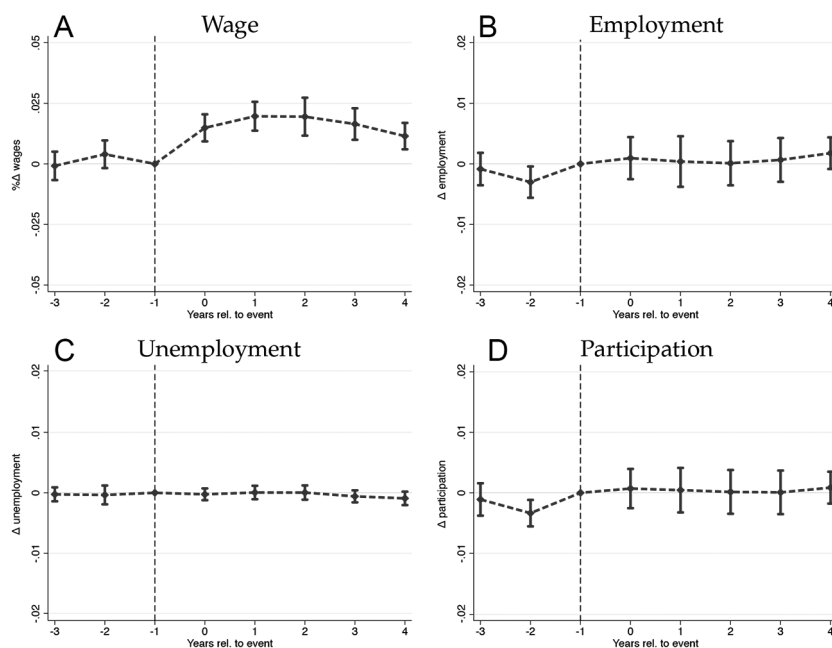


FIG. 6.—Impact of the minimum wage over time, high-recall group. The figure shows the main results from our event study analysis (see eq. [1]) using 172 state-level minimum wage changes between 1979 and 2019. The figure shows the effect of a minimum wage increase on wages (A), on employment to population (B), on unemployment to population (C), and on labor force participation rate (D) for the high-recall group. The high-recall group consists of all workers whose predicted probability is above 12%—a threshold that corresponds to a 75% recall rate. To calculate the predicted probabilities, we use the best-performing prediction model—the boosted tree model. We also show the 95% confidence interval based on standard errors that are clustered at the state level. A color version of this figure is available online.

horizon between 1 and 3 years preceding the minimum wage increase. Furthermore, the small drop in employment between 1 and 2 years preceding the minimum wage increase would imply that the economy slightly deteriorates before an average minimum wage hike, so we would expect a decrease in employment rate after the policy change. In contrast, we see no clear break in employment after the policy change—if anything, there is a small, statistically insignificant increase.

Panel C shows the impact on the unemployment rate. There is neither any preexisting trend nor any break after the minimum wage increase. Panel D shows the impact of the policy on participation rate, which also shows a flat response following the policy change. Overall, the preexisting trends are reassuring, and we find no indication of any break in the participation rate at the time of the minimum wage increases. Moreover, we do not find evidence

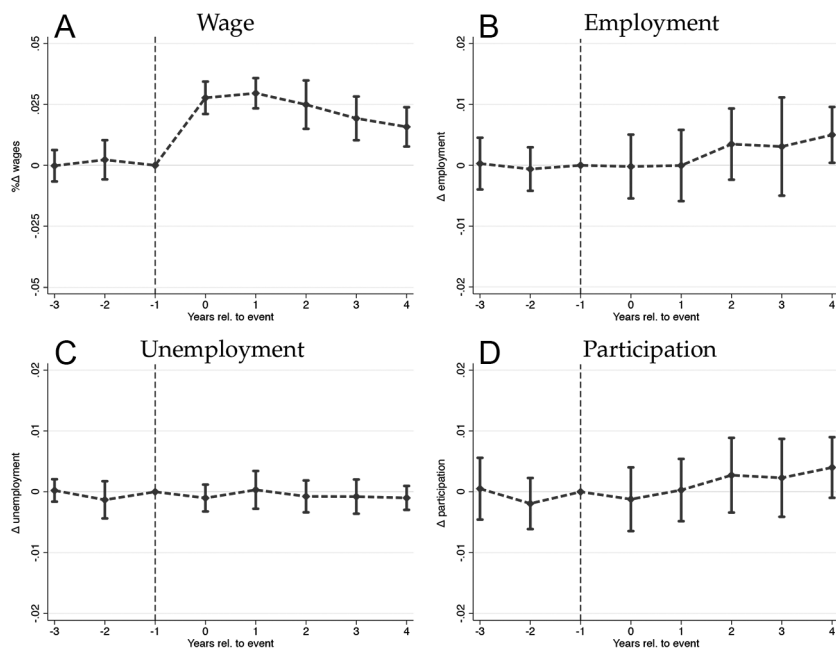


FIG. 7.—Impact of the minimum wage over time, high-probability group. The figure shows the main results from our event study analysis (see eq. [1]) using 172 state-level minimum wage changes between 1979 and 2019. The figure shows the effect of a minimum wage increase on wages (A), on employment to population (B), on unemployment to population (C), and on labor force participation rate (D) for the high-probability group. The high-probability group consists of the 10% of the population with the highest likelihood of being affected by the policy. To calculate the predicted probabilities, we use the boosted tree model. We also show the 95% confidence interval based on standard errors that are clustered at the state level. A color version of this figure is available online.

of lower employment (or higher unemployment) rates when we look up to the fifth year following the policy change.

Impact by Predicted Probability Quintiles

Figure 8 shows the effect of the minimum wage separately for each predicted probability quintile. Figure 8A shows that there is a clear wage effect among the individuals most exposed to the minimum wage. Wages also increase for the second-highest quintile of predicted probabilities, but this increase is statistically insignificant. Figure 8B–8D shows the effect of the policy on labor market outcomes, such as employment, unemployment, and participation. None of these outcomes shows any substantial change even at the highest predicted probability quintile. Reassuringly, none of the outcomes are

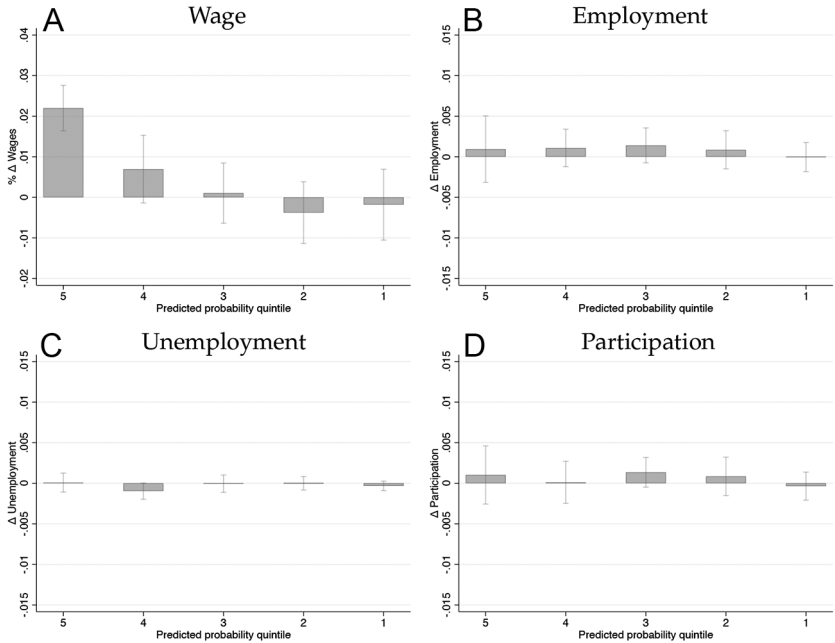


FIG. 8.—Impact of the minimum wage by predicted probability quintiles. The figure shows the effect of the minimum wage separately for each predicted probability quintile. The highest quintile comprises individuals with predicted probabilities (of being minimum wage workers) in the top 20%. We estimate equation (1) for each quintile separately and report the 5-year averaged posttreatment estimates. We use 172 state-level minimum wage changes between 1979 and 2019. The figure shows the effect of a minimum wage increase on wages (A), on employment to population ratio (B), on unemployment rate (C), and on labor force participation rate (D). We also show the 95% confidence intervals based on standard errors that are clustered at the state level. A color version of this figure is available online.

statistically distinguishable from zero in the bottom quintiles of predicted probability, which provides support for the validity of the research design.

Robustness

In table 3 we assess the robustness of the main results shown in table 2 to the inclusion of various versions of time-varying heterogeneity for the high-recall group, while we report the same robustness checks for the high-probability group in table A.3. In column 1 we report the estimates for the baseline specification shown in table 2. Column 2 allows the period effects to vary by the nine census divisions. The results are similar to the baseline specification: we find a positive and statistically significant wage effect and a positive employment effect—which comes from an increase in the participation rate.

Table 3
Impact of the Minimum Wage on Labor Market Outcomes, Robustness to Alternative Specifications (High-Recall Group)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Δ wage (%)	.016*** (.003)	.016*** (.003)	.018*** (.003)	.012*** (.003)	.016*** (.005)	.016*** (.002)	.016*** (.002)
Δ employment (pp)	.001 (.001)	.002 (.002)	.001 (.002)	-.001 (.001)	.002 (.002)	.000 (.001)	.000 (.001)
Δ unemployment (pp)	-.000 (.000)	.000 (.000)	-.001 (.000)	.000 (.001)	-.001* (.000)	-.000 (.000)	-.000 (.000)
Δ participation (pp)	.000 (.001)	.002 (.002)	.000 (.002)	-.000 (.001)	.001 (.002)	.000 (.001)	.000 (.001)
Employment elasticity with respect to minimum wage	.020 (.038)	.047 (.038)	.018 (.040)	-.014 (.037)	.033 (.052)	.010 (.021)	.007 (.035)
Employment elasticity with respect to wage	.114 (.216)	.284 (.242)	.096 (.208)	-.109 (.285)	.233 (.330)	.071 (.141)	.048 (.205)
Number of events	172	172	406	172	99	172	172
Number of observations	7,854	7,854	7,854	7,854	7,854	7,854	7,854
Number of individuals in sample	20,917,455	20,917,455	20,917,455	20,917,455	20,917,455	20,917,455	20,917,455
Mean employment	.415	.415	.425	.419	.426	.415	.415
Mean unemployment	.046	.046	.049	.045	.050	.046	.046
Mean participation	.460	.460	.473	.464	.476	.460	.460

So far, we have focused only on prominent state-level minimum wage changes. In column 3 of table 3 we expand the event definition to include (nontrivial) federal minimum wage increases. This leads us to 406 minimum wage changes. Inclusion of these federal-level events produces similar results, although the (small) reduction in the unemployment rate is now marginally statistically significant. In column 4 we provide estimates without using population weights. The estimates are similar, albeit the employment and participation estimates without population weight are somewhat closer to zero and more noisily estimated. The implied employment elasticity with respect to wage without population weights is -0.109 (SE, 0.29), which can rule out only sizeable negative responses to the policy at the conventional significance levels.

In column 5 we restrict our analysis to prominent state-level minimum wage changes where we observe all of the 5 years after reform. In this analysis we include only events that took place on or before the first quarter of 2014. Restricting the sample to these events does not affect the key findings: we obtain a slightly larger participation estimate (0.1 percentage points instead of zero), but the estimates also become a little noisier. Finally, columns 6 and 7 control for the state unemployment and employment rates (as a fraction of population), using all workers in the case of column 6 and workers in the low-probability group (i.e., probability less than 12%—the high-recall cutoff) in the case of column 7. Controlling for the state unemployment and employment rates—either using all workers or just workers with a low probability of being a minimum wage worker—produces results that are very close the baseline estimation. The ability to control for state-level labor market conditions of the workers unlikely to be treated by the policy is another advantage of our prediction-based approach. The findings confirm that our results are not driven by unusual labor market conditions in treated states around the time of the policy changes.

Effects on Labor Market Transitions

Table 4 shows the estimated impact on monthly labor market transition rates between employment, unemployment, and participation. We find no statistically or economically significant effects on the transition probabilities in response to the policy change for the high-probability group (col. 1) or the high-recall group (col. 2). This clarifies that the minimum wage increase did not lengthen unemployment durations; if anything, the policy accelerated monthly transitions from unemployment to employment by 0.006 (SE, 0.005) percentage points for the high-probability group and by 0.004 (SE, 0.004) percentage points for the high-recall group. Therefore, we find no indication that some individuals were pushed permanently into long-term unemployment or that the labor markets became more sclerotic (with low job-finding rates) in response to increases in the minimum wage.

Table 4
Impact of the Minimum Wage on Labor Market Transitions

	(1)	(2)	(3)
Δ E-U flow as a share of employment (pp)	.000 (.001)	.000 (.000)	.000 (.000)
Δ E-I flow as a share of employment (pp)	−.000 (.001)	.001 (.001)	.000 (.000)
Δ U-E flow as a share of unemployment (pp)	.006 (.005)	.004 (.004)	.004 (.003)
Δ U-I flow as a share of unemployment (pp)	.002 (.005)	.003 (.004)	.000 (.003)
Δ I-E flow as a share of inactivity (pp)	.001 (.001)	.000 (.000)	.001 (.001)
Δ I-U flow as a share of inactivity (pp)	.000 (.001)	.000 (.000)	−.000 (.000)
Number of events	172	172	172
Number of observations	7,854	7,854	7,854
Number of individuals in sample	4,701,665	14,856,017	21,156,039
Mean E-U flow as a share of employment (%)	.028	.022	.009
Mean E-I flow as a share of employment (%)	.107	.061	.020
Mean U-E flow as a share of unemployment (%)	.230	.246	.258
Mean U-I flow as a share of unemployment (%)	.369	.295	.190
Mean I-E as a share of inactivity (%)	.057	.042	.053
Mean I-U as a share of inactivity (%)	.036	.024	.023
Group	High probability	High recall	Low probability
Prediction model	Boosted tree	Boosted tree	Boosted tree

NOTE.—The table reports the effects of the minimum wage on labor market transition rates according to the event study analysis (see eq. [1]) using 172 state-level minimum wage changes between 1979 and 2019. The table reports 5-year averaged posttreatment estimates for percentage point changes in the employment-to-unemployment (E-U) transition rate as a share of employment (row 1), the employment-to-inactivity (E-I) transition rate as a share of employment (row 2), the unemployment-to-employment (U-E) transition rate as a share of unemployment (row 3), the unemployment-to-inactivity (U-I) transition rate as a share of unemployment (row 4), the inactivity-to-employment (I-E) transition rate as a share of inactivity (row 5), and the inactivity-to-unemployment (I-U) transition rate as a share of inactivity (row 6). We also report the mean levels of each of these variables for the period preceding the minimum wage hikes. “Number of observations” refers to the number of quarter-state cells used for estimation, while “Number of individuals” refers to the underlying CPS sample used to calculate labor market outcomes in these cells. Column 1 shows estimates for the high-probability group, which captures the 10% of the population with highest predicted probability. Column 2 shows estimates for the high-recall group, which consists of individuals whose predicted probability is above 12%—a threshold that leads to a 75% recall rate of minimum wage workers. Column 3 shows the estimates for workers whose predicted probability is below 12%. Columns 1–3 use the best-performing prediction model—the boosted tree model. All regressions are weighted by state-quarter population. Robust standard errors in parentheses are clustered by state. pp = percentage point.

Effects by Demographic Subgroups

In some cases, policy makers may be interested in the impact of minimum wages on specific demographic groups. Moreover, individuals are more likely to be on the participation margin in some demographic groups than in others. As a result, the impact of the policy on participation decisions might vary considerably across demographic groups. We study the heterogeneity in response to the minimum wage in table 5. To make sure that changes in labor market outcomes are driven by the minimum wage itself and not something

Table 5
Impact of the Minimum Wage on Labor Market Outcomes by Demographic Group I

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ wage (%)	.016*** (.003)	.012*** (.003)	.015*** (.004)	.001 (.005)	.027*** (.003)	.016*** (.007)	.015*** (.004)	.014*** (.003)
Δ employment (pp)	.001 (.001)	-.003 (.003)	.001 (.001)	-.001 (.002)	.004 (.003)	-.002 (.004)	.001 (.002)	.000 (.001)
Δ unemployment (pp)	-.000 (.000)	-.001 (.001)	-.001* (.000)	-.000 (.000)	-.001 (.002)	.000 (.001)	-.000 (.000)	-.000 (.000)
Δ participation (pp)	.000 (.001)	-.004 (.002)	.000 (.001)	-.002 (.002)	.003 (.002)	-.001 (.004)	.000 (.002)	.000 (.001)
Employment elasticity with respect to minimum wage	.020 (.038)	-.069 (.056)	.020 (.039)	-.044 (.050)	.112 (.088)	-.072 (.147)	.019 (.058)	.013 (.038)
Employment elasticity with respect to wage	.114 (.216)	-.526 (.431)	.128 (.248)	NA	.396 (.318)	-.430 (.927)	.120 (.356)	.090 (.258)

Number of events	172	172	172	172	172	172	172	172	172	172	172	172	172
Number of observations	7,854	7,841	7,854	7,854	7,854	7,854	7,854	7,854	7,854	7,854	7,854	7,854	7,854
Number of individuals in sample	20,917,455	5,680,302	13,357,475	4,883,626	3,763,811	2,163,408	9,118,096	16,755,128					
Mean employment	.415	.499	.385	.347	.333	.264	.361	.405					
Mean unemployment	.046	.061	.038	.024	.068	.015	.050	.049					
Mean participation	.460	.560	.423	.371	.401	.279	.412	.453					
Group	High recall	High recall	High recall	High recall	High recall	High recall	High recall	High recall	High recall	High recall	High recall	High recall	High recall
Demographic group	All	Black or Hispanic	Female	Married female	Teen	Aged 60–70	LTHS	HSL					

NOTE.—The table reports the effects of the minimum wage on labor market outcomes by demographic group according to the event study analysis (see eq. [1]). We exploit 172 state-level minimum wage changes between 1979 and 2019. We assess the impact of the minimum wage on the high-recall group including individuals whose predicted probability is above 12%—a threshold that leads to a 75% recall rate of minimum wage workers. The table reports 5-year averaged posttreatment estimates for each key labor market outcome: percent change in wages and the change in employment to population, unemployment to population, and labor force participation rate. We also report the employment elasticity with respect to the minimum wage and the employment elasticity with respect to the wage, which is the ratio of the percent change in employment and wage. To calculate the percent change in employment, we divide the change in employment to population by the mean employment to population rate preceding the minimum wage hikes (reported at the bottom of the table). “Number of observations” refers to the number of quarter-state cells used for estimation, while “Number of individuals” refers to the underlying CPS sample used to calculate labor market outcomes in these cells. In all of the regressions, we use the best-performing prediction model—the boosted tree model. The demographic subgroups are Black or Hispanic, woman, married woman, teen, aged 60 and older and less than 70, less than high school (LTHS) education, and high school or less (HSL) education. All estimates are weighted by the corresponding subgroup’s population in the state. Robust standard errors in parentheses are clustered by state. pp = percentage point.

* $p < .10$.
 ** $p < .05$.
 *** $p < .01$.

else, in each subgroup, we focus on workers who are in the high-recall group.²¹ Restricting the sample to workers who are likely to be minimum wage workers is also necessary for getting first-stage wage effects, which would not be possible if we had all workers in the sample (see table A.1 in the online appendix of Cengiz et al. 2019).

For most subgroups in table 5, we find a clear and significant impact on wages, which confirms a key advantage of restricting the sample to groups that are likely to be exposed to the minimum wage. In column 1 we simply reproduce the benchmark estimates for the overall high-recall group for comparability. In column 2 we report the estimates for workers who are Black or Hispanic.²² While the estimates are somewhat noisy, the point estimates indicate a nontrivial drop in employment and participation. This suggests that in the Black or Hispanic group, some workers (who were at the participation margin) may have been made worse off by the policy. However, the standard errors are too large to draw a clear conclusion on this question.

In columns 3 and 4 we examine the impact of the policy on all women and on married women, respectively. Since the extensive margin labor supply elasticities are often found to be larger for women than for men, it might be the case that minimum wages lead to a greater increase in women's participation. Furthermore, married women are also typically thought to have larger responsiveness on the extensive margin. However, our estimated effects for employment and participation outcomes for women and married women are very similar to the benchmark specification; however, we note that there is no statistically (or economically) significant wage effect for married women in our sample, which makes it difficult to draw any strong conclusions. (We find more informative evidence when we consider married mothers with younger kids below.)

Columns 5 and 6 report the impact on teens and on older individuals (aged 60–70), respectively. Both younger individuals and older individuals have lower participation rates than prime-age individuals, and they are also thought to have more elastic labor supply (Blundell, Bozio, and Laroque 2011). For teens, we find larger wage effects and slightly larger employment increase

²¹ We use the predicted probabilities estimated on all workers. One could estimate a separate prediction model for each subgroup and then use those predicted probabilities. However, it is unlikely that there are substantial gains from estimating a separate prediction model for each subgroup. If there were substantial gains from using some predictors differently within a particular subgroup, then the boosting algorithm would tend to detect and incorporate this fact into the prediction model—even if it were estimated using the full sample. We also explored separately estimating prediction models for some specific demographic groups but found negligible changes in the precision of our estimates.

²² We attempted to estimate the effect of the policy on Black and Hispanic individuals separately, but the estimates were too imprecise to be informative.

than for the overall sample. The increase in employment comes from the changes in participation, which is in line with the idea that more teens are on the participation margin and also with the evidence presented by Laws (2018). For older individuals, the wage effects are imprecise (although the point estimates are close to the overall sample). This makes interpreting the results for older individuals difficult. Still, we document a slight positive (statistically insignificant) employment and participation effect, which is in line with Borgschulte and Cho (2019), who document a similar-sized response in employment for those aged between 62 and 70.

Columns 7 and 8 show the impact on those with lower educational credentials. The labor market impact of the policy on these education groups is very similar to the impact of the policy on the overall sample. These findings suggest that workers with lower education credentials seem to benefit from the minimum wage policies.

Table 6 presents additional heterogeneity analysis. Columns 2–4 show the estimates by the predicted probability of moving into and out of the labor force. The goal of this exercise is to help us assess any heterogeneity in the effects by how likely workers are to be at the margins of LFP. We predict the probability that an individual changes their LFP status—either from nonparticipation to participation or from participation to nonparticipation—by applying the boosting tree ML method. For estimating this prediction model, we use all of the demographic variables that we had used before to predict minimum wage exposure but also add the number of children under the age of 5, since it is likely to be an important predictor of LFP. The relative importance of the predictors is shown in figure A.6. Similar to the prediction model on minimum wage exposure, age, education, and gender are the three most important predictors of changes in participation. In addition, race and number of children under the age of 5 also substantially influence the prediction model.

Since the number of children is coded consistently only since 1986 in the CPS, we report estimates using the 1986–2018 period in table 6. Column 1 reports estimates for all workers using only that period, and they are very similar to the estimates using the 1979–2019 sample. Column 2 shows the estimates for the group of workers that has a high predicted probability of changing LFP status. The estimated impacts for this group are very similar to those from the overall sample. Columns 3 and 4 show the estimates with lower probability of switching. Again, we find responses similar to those from the overall sample. The key takeaway is that even when we look at individuals who are more likely to switch LFP, we find no meaningful difference in the causal effects of minimum wage policies. These findings cast doubt that there was much of an impact from minimum wage changes in our sample on job search and participation behavior, even among groups that are likely to be at the margin of participation. In columns 5 and 6 we also report estimates on single and married mothers, focusing on those with kids

Table 6
Impact of the Minimum Wage on Labor Market Outcomes by Demographic Group II

	(1)	(2)	(3)	(4)	(5)	(6)
Δ wage (%)	.015*** (.003)	.017*** (.003)	.012*** (.003)	.011 (.008)	.025** (.012)	.006 (.009)
Δ employment (pp)	.001 (.001)	.000 (.002)	.002 (.002)	-.001 (.001)	-.000 (.006)	-.005 (.004)
Δ unemployment (pp)	-.001 (.000)	-.001 (.001)	-.001 (.001)	.001* (.000)	-.001 (.003)	-.001 (.001)
Δ participation (pp)	.000 (.001)	-.001 (.002)	.001 (.002)	-.001 (.001)	-.001 (.007)	-.006 (.004)
Employment elasticity with respect to minimum wage	.017 (.035)	.006 (.035)	.035 (.038)	-.084 (.071)	-.005 (.113)	-.115 (.097)
Employment elasticity with respect to wage	.107 (.217)	.033 (.193)	.276 (.339)	NA	-.021 (.442)	NA

Number of events	156	156	156	156	156	156	156
Number of observations	6,222	6,222	6,222	6,222	6,222	6,222	6,222
Number of individuals in sample	15,760,550	7,883,899	2,915,439	4,961,212	525,347	618,103	
Mean employment	.417	.512	.561	.148	.520	.417	
Mean unemployment	.047	.068	.042	.008	.091	.040	
Mean participation	.463	.581	.604	.155	.612	.457	
Probability group	High recall	High recall	High recall	High recall	High recall	High recall	
Demographic group	All	High LFP switch	Medium LFP switch	Low LFP switch	Single mother kids under 5	Married mother kids under 5	

NOTE.—The table reports the effects of the minimum wage on labor market outcomes by demographic group according to the event study analysis (see eq. [1]). We exploit 156 state-level minimum wage changes between 1986 and 2018. We assess the impact of the minimum wage on the high-recall group including individuals whose predicted probability is above 12%—a threshold that leads to a 75% recall rate of minimum wage workers. The table reports 5-year averaged post-treatment estimates for each key labor market outcome: the change in employment to population, unemployment to population, and labor force participation rate. “Number of observations” refers to the number of quarter-state cells used for estimation, while “Number of individuals” refers to the underlying CPS sample used to calculate labor market outcomes in these cells. In all of the regressions, we use the best-performing prediction model—the boosted tree model. Column 1 shows the estimates on overall employment using the 1986–2018 periods. Columns 2–4 show the estimates for individuals with the different predicted probability of moving into or out of the labor force. Column 5 shows the estimates for single mothers with children under the age of 5, while col. 6 shows the estimates for married mothers with children under the age of 5. All estimates are weighted by the corresponding subgroup’s population in the state. Robust standard errors in parentheses are clustered by state. pp = percentage point.

* $p < .10$.

** $p < .05$.

*** $p < .01$.

under 5. For single mothers, we find a very small increase in employment and participation; for married mothers, we find a small reduction in employment and participation. Neither set of estimates are economically or statistically significant.²³

Additional Labor Market Outcomes

Finally, in table 7 we study the impact on other labor market outcomes, such as self-employment, part-time (working fewer than 30 hours per week), and overtime (working more than 40 hours per week) status. Changing working hours and pushing workers to self-employment are often argued to be important margins of adjustment to the minimum wage. We find little indication that the high-probability or the high-recall group experienced any changes in self-employment. We do find a decline in the share of employees in part-time jobs for the high-probability group, with close to no change in the share of employees in overtime jobs. The former change is statistically significant at the conventional levels. This implies that the minimum wages increase the share of full-time jobs (without any overtime) in the low-wage workforce.

VI. Conclusion

In this paper we estimate the impact of the minimum wage on various labor market outcomes using 172 prominent minimum wage changes between 1979 and 2019. To capture the impact of the policy on a broad group of affected workers, we utilize modern ML techniques to estimate the likelihood that someone is a minimum wage worker. While the best-performing prediction model does better than the linear prediction model of Card and Krueger (1995), the gap is not large. One implication of these findings is that minimum wage researchers who are not interested in investing in a ML approach may do fairly well by simply applying the Card and Krueger linear probability prediction model. Of course, the advantage of the ML approach is that

²³ We find an effect of minimum wages on employment for single mothers with young children that is close to zero and statistically insignificant. This differs from Godoy, Reich, and Allegretto (2020), who find a large, statistically significant increase in employment for single mothers with kids under 5 in response to a minimum wage increase. Table A.5 provides a reconciliation of the discrepancy between these findings in greater detail. There are a number of differences between the two papers, but the primary reason why our estimates differ from Godoy, Reich, and Allegretto (2020) is that we use the CPS-Basic files while they use the (smaller) CPS-ORG data. While use of the CPS-Basic versus CPS-ORG data by itself does not make a difference in the overall sample (see cols. 1–5 in table A.5), the results are more sensitive to the data sets for the single mother sample, which is much smaller (see cols. 6–10 in table A.5).

Table 7
Impact of the Minimum Wage on Alternative Labor Market Outcomes

	(1)	(2)	(3)
Δ self-employment as share of employment (pp)	.001 (.001)	−.000 (.001)	−.001 (.001)
Δ part-time as share of employment (pp)	−.005** (.002)	−.001 (.001)	.000 (.000)
Δ overtime as share of employment (pp)	.001 (.001)	.000 (.001)	−.001 (.001)
Number of events	172	172	172
Number of observations	7,854	7,854	7,854
Number of individuals in sample	6,639,492	20,917,455	29,370,470
Mean self-employment	.029	.066	.129
Mean part-time	.414	.216	.080
Mean overtime	.029	.067	.196
Group	High probability	High recall	Low probability
Prediction model	Boosted tree	Boosted tree	Boosted tree

NOTE.—The table reports the effects of the minimum wage on alternative labor market outcomes according to the event study analysis (see eq. [1]) using 172 state-level minimum wage changes between 1979 and 2019. The table reports 5-year averaged posttreatment estimates for the following labor market outcomes: self-employment, part-time (working fewer than 30 hours per week), and overtime (working more than 40 hours per week). Each of these variables is expressed as a share of total employment. Column 1 shows estimates for the high-probability group, which captures the 10% of the population with the highest predicted probability. Column 2 shows estimates for the high-recall group, which consists of individuals whose predicted probability is above 12%—a threshold that leads to a 75% recall rate of minimum wage workers. Column 3 shows the estimates for workers whose predicted probability is below 12%. All columns use the best-performing prediction model—the boosted tree model. All of the regressions are weighted by state-quarter population. Robust standard errors in parentheses are clustered by state. pp = percentage point.

** $p < .05$.

it will successfully locate minimum wage workers in environments different from the one we study here, including in other data sets or time periods.

The prediction-based approach allows us to study the impact of the minimum wage not just on wages and employment but also on unemployment and participation rates for groups that cover 75% of all minimum wage workers. These groups include substantially more minimum wage workers than the demographic-based and industry-based subsamples commonly used in the literature. In line with much of the existing evidence in the literature, we find that the minimum wage has a positive and significant impact on wages, while employment effects are modest in the US context. We also show that the slight (statistically insignificant) employment increase comes from a slight drop in unemployment and a slight increase in the participation rate. These responses indicate that the minimum wage is unlikely to have a negative impact on workers by discouraging them to search for new jobs.

We also find no significant heterogeneities in the responses to the minimum wage. The most likely exposed group and a much boarder group that covers 75% of all minimum wage workers responded very similarly to the policy change. The only indication for a potential negative impact of the policy was on the Black or Hispanic subgroup, but those estimates are too noisy

to draw a definitive conclusion. Overall, our results underscore the positive impact of the policy on key labor market outcomes.

Our findings suggest that the prediction-based approach of defining a likely treatment group could be naturally applied to study the impact of the policy on various other important outcomes measured in different data sources. Since the most important predictors (age, education, gender) are available in most data sources, it is straightforward to assess the impact of the minimum wage hikes along the predicted probability spectrum (and show estimates like those in fig. 5). A key advantage of doing so is to study directly whether there is some nonlinearity in response to the minimum wage by the extent of exposure. Furthermore, studying the impact of the policy on the low-probability group can serve as an additional falsification test and can provide further support for the credibility of the research design. Given these advantages, the prediction-based approach should be a part of the standard toolkit for analyzing the impact of minimum wages.

References

- Adams, Camilla, Jonathan Meer, and Carly Will Sloan. 2018. The minimum wage and search effort. NBER Working Paper no. 25128, National Bureau of Economic Research, Cambridge, MA.
- Addison, John T., and McKinley L. Blackburn. 1999. Minimum wages and poverty. *Industrial and Labor Relations Review* 52, no. 3:393–409.
- Addison, John T., McKinley L. Blackburn, and Chad Cotti. 2011. Minimum wage increases under straightened circumstances. IZA Discussion Paper no. 6036, Institute of Labor Economics, Bonn.
- Ahn, Tom, Peter Arcidiacono, and Walter Wessels. 2011. The distributional impacts of minimum wage increases when both labor supply and labor demand are endogenous. *Journal of Business and Economic Statistics* 29, no. 1:12–23.
- Allegretto, Sylvia, Arindrajit Dube, Michael Reich, and Ben Zipperer. 2017. Credible research designs for minimum wage studies: A response to Neumark, Salas, and Wascher. *ILR Review* 70, no. 3:559–92.
- Autor, David H., John J. Donohue III, and Stewart J. Schwab. 2006. The costs of wrongful-discharge laws. *Review of Economics and Statistics* 88, no. 2:211–31.
- Belman, Dale, and Paul J. Wolfson. 2014. *What does the minimum wage do?* Kalamazoo, MI: W.E. Upjohn Institute.
- Belman, Dale, Paul Wolfson, and Kritkorn Nawakitphaitoon. 2015. Who is affected by the minimum wage? *Industrial Relations: A Journal of Economy and Society* 54, no. 4:582–621.
- Blundell, Richard, Antoine Bozio, and Guy Laroque. 2011. Labor supply and the extensive margin. *American Economic Review* 101, no. 3:482–86. <http://www.jstor.org/stable/29783793>.

- Borgschulte, Mark, and HeePyung Cho. 2019. Minimum wages and retirement. *ILR Review* 73, no. 1:153–77.
- Breiman, Leo. 2001. Random forests. *Machine Learning* 45, no. 1:5–32.
- Card, David. 1992. Do minimum wages reduce employment? A case study of California, 1987–89. *ILR Review* 46, no. 1:38–54.
- Card, David, and Alan B. Krueger. 1995. *Myth and measurement: The new economics of the minimum wage*. Princeton, NJ: Princeton University Press.
- Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer. 2019. The effect of minimum wages on low-wage jobs. *Quarterly Journal of Economics* 134, no. 3:1405–54.
- Clemens, Jeffrey, Lisa B. Kahn, and Jonathan Meer. 2018. The minimum wage, fringe benefits, and worker welfare. NBER Working Paper no. 24635, National Bureau of Economic Research, Cambridge, MA.
- Clemens, Jeffrey, and Michael R. Strain. 2017. Estimating the employment effects of recent minimum wage changes: Early evidence, an interpretative framework, and a pre-commitment to future analysis. NBER Working Paper no. 23084, National Bureau of Economic Research, Cambridge, MA.
- Clemens, Jeffrey, and Michael Wither. 2019. The minimum wage and the Great Recession: Evidence of effects on the employment and income trajectories of low-skilled workers. *Journal of Public Economics* 170:53–67.
- Currie, Janet, and Bruce C. Fallick. 1996. The minimum wage and the employment of youth evidence from the NLSY. *Journal of Human Resources* 31, no. 2:404–28.
- Drazen, Allan. 1986. Optimal minimum wage legislation. *Economic Journal* 96, no. 383:774–84. <https://doi.org/10.2307/2232990>.
- Dustmann, Christian, Attila Lindner, Uta Schönberg, Matthias Umkehrer, and Philipp vom Berge. 2022. Reallocation effects of the minimum wage. *Quarterly Journal of Economics* 137, no. 1:267–328.
- Flinn, Christopher J. 2002. Interpreting minimum wage effects on wage distributions: A cautionary tale. *Annales d'Economie et de Statistique* 67/68:309–55. <http://www.jstor.org/stable/20076351>.
- . 2006. Minimum wage effects on labor market outcomes under search, matching, and endogenous contact rates. *Econometrica* 74, no. 4:1013–62. <http://www.jstor.org/stable/3805915>.
- . 2011. *The minimum wage and labor market outcomes*. Cambridge, MA: MIT Press.
- Friedman, Jerome H. 2001. Greedy function approximation: A gradient boosting machine. *Annals of Statistics* 29, no. 5:1189–232.
- Friedman, Jerome, Trevor Hastie, and Robert Tibshirani. 2009. *The elements of statistical learning: Data mining, inference, and prediction*. New York: Springer.
- Gardner, John. 2021. Two-stage differences in differences. Working paper, University of Mississippi.

- Giuliano, Laura. 2013. Minimum wage effects on employment, substitution, and the teenage labor supply: Evidence from personnel data. *Journal of Labor Economics* 31, no. 1:155–94.
- Godoy, Anna, Michael Reich, and Sylvia A. Allegretto. 2020. Parental labor supply: Evidence from minimum wage changes. IRLE Working Paper no. 103-19, Institute for Research on Labor and Employment, University of California, Berkeley.
- Lang, Kevin. 1987. Pareto improving minimum wage laws. *Economic Inquiry* 25, no. 1:145–58. <https://doi.org/10.1111/j.1465-7295.1987.tb00729.x>.
- Laws, Athene. 2018. Do minimum wages increase search effort? CWPE Working Paper no. 1857. Cambridge Working Papers in Economics, University of Cambridge.
- Linneman, Peter. 1982. The economic impacts of minimum wage laws: A new look at an old question. *Journal of Political Economy* 90, no. 3:443–69.
- Lundstrom, Samuel M. 2016. When is a good time to raise the minimum wage? *Contemporary Economic Policy* 35, no. 1:29–52.
- Manning, Alan. 2016. The elusive employment effect of the minimum wage. Working paper, Washington Center for Equitable Growth, Washington, DC.
- Marimpi, Maria, and Pierre Koning. 2018. Youth minimum wages and youth employment. *IZA Journal of Labor Policy* 7:1–18.
- Meer, Jonathan, and Jeremy West. 2015. Effects of the minimum wage on employment dynamics. *Journal of Human Resources* 51, no. 2:500–522.
- Mincer, Jacob. 1976. Unemployment effects of minimum wages. *Journal of Political Economy* 84, no. 4:S87–S104.
- Neumark, David. 2017. The employment effects of minimum wages: Some questions we need to answer. NBER Working Paper no. 23584, National Bureau of Economic Research, Cambridge, MA.
- Neumark, David, J. M. Ian Salas, and William Wascher. 2014. Revisiting the minimum wage-employment debate: Throwing out the baby with the bathwater? *ILR Review* 67, suppl.: 608–48.
- Neumark, David, and William Wascher. 1992. Employment effects of minimum and subminimum wages: Panel data on state minimum wage laws. *ILR Review* 46, no. 1:55–81.
- . 2008. *Minimum wages*. Cambridge, MA: MIT Press. <https://mitpress.mit.edu/books/minimum-wages>.
- Ragan, James F. 1977. Minimum wages and the youth labor market. *Review of Economics and Statistics* 59, no. 2:129–36.
- Sabia, Joseph J., Richard V. Burkhauser, and Benjamin Hansen. 2012. Are the effects of minimum wage increases always small? New evidence from a case study of New York State. *ILR Review* 65, no. 2:350–76.
- Sun, Liyang, and Sarah Abraham. 2020. Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics* 225, no. 2:175–99.

- Swinerton, Kenneth. 1996. Minimum wages in an equilibrium search model with diminishing returns to labor in production. *Journal of Labor Economics* 14, no. 2:340–55.
- Totty, Evan. 2017. The effect of minimum wages on employment: A factor model approach. *Economic Inquiry* 55, no. 4:1712–37.
- Vaghul, Kavya, and Ben Zipperer. 2016. Historical state and sub-state minimum wage data. Working paper, Washington Center for Equitable Growth, Washington, DC. <http://cdn.equitablegrowth.org/wp-content/uploads/2016/09/02153029/090716-WP-Historical-min-wage-data.pdf>.
- Wessels, Walter J. 2005. Does the minimum wage drive teenagers out of the labor force? *Journal of Labor Research* 26:169–76.
- Zou, Hui, and Trevor Hastie. 2005. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society B* 67, no. 2:301–20.