

Discussion Paper

No. 2016-17 | April 26, 2016 | <http://www.economics-ejournal.org/economics/discussionpapers/2016-17>

Minimum Wage and Employment: Escaping the Parametric Straitjacket

Stefano Cabras, Jan Fidrmuc, and Juan de Dios Tena

Abstract

Parametric regression models are often not flexible enough to capture the true relationships as they tend to rely on arbitrary identification assumptions. Using the UK Labor Force Survey, the authors estimate the causal effect of national minimum wage (NMW) increases on the probability of job entry and job exit by means of a non-parametric Bayesian modelling approach known as Bayesian Additive Regression Trees (BART). The application of this methodology has the important advantage that it does not require ad-hoc assumptions about model fitting, number of covariates and how they interact. They find that the NMW exerts a positive and significant impact on both the probability of job entry and job exit. Although the magnitude of the effect on job entry is higher, the overall effect of NMW is ambiguous as there are many more employed workers. The causal effect of NMW is higher for young workers and in periods of high unemployment and they have a stronger impact on job entry decisions. No significant interactions were found with gender and qualifications.

JEL C23 C11 C14 J3 J4

Keywords BART; causal inference; regression approach; matching regression

Authors

Stefano Cabras, Universidad Carlos III and Università di Cagliari

Jan Fidrmuc, ✉ Brunel University; Charles University; and CESifo Munich,
jan.fidrmuc@gmail.com

Juan de Dios Tena, University of Liverpool; Università di Sassari; CRENoS; and Instituto Flores de Lemus

We received helpful comments and suggestions from Brendan McCabe, Andy Tremayne and David Forrest, as well as seminar/conference participants at the University of Liverpool and “The German Minimum Wage – First Evidence and Experiences from Other Countries” conference at the Institute for Employment Research (IAB) in Nuremberg.

Citation Stefano Cabras, Jan Fidrmuc, and Juan de Dios Tena (2016). Minimum Wage and Employment: Escaping the Parametric Straitjacket. Economics Discussion Papers, No 2016-17, Kiel Institute for the World Economy. <http://www.economics-ejournal.org/economics/discussionpapers/2016-17>

1 Introduction

The most characteristic feature of the literature on the causal impact of the minimum wage on employment is the general lack of consensus. Neumark and Washer (2007) compile an extensive survey of previous research and conclude that the minimum wage exerts an adverse impact on employment of low-skilled workers and a non significant impact on total employment. However, other surveys on this issue, a meta-analysis by Card and Krueger (1995) and the subsequent extension by Doucouliagos and Stanley (2008), find that there is a wide range of results in the previous research, and that once the publication selection bias is accounted for the mean estimate is consistent with a non-significant impact of the minimum wage on employment.

A possible reason for the wide range of findings is the fact that the results hinge dramatically on ad hoc assumptions about the parametric specification of the empirical model and on the definition of the control group in the analysis. This is corroborated in the highly insightful and interesting discussion in a series of papers by Allegretto et al. (2011, 2013), Dube et al. (2010), and Neumark et al. (2014) in a state-level panel analysis for the US. Dube et al. (2010) and Allegretto et al. (2011) suggest that it is essential to control for spatial heterogeneity in order to estimate the impact of the minimum wage in a panel data setting. In particular, they propose to include two types of local controls consisting of: (1) jurisdiction-specific linear time trends; and (2) interactions between time dummy variables for sets of neighboring states or neighboring counties so they could be used as controls to determine the impact of the minimum wages. Subsequently, Neumark et al. (2014) and Sabia et al. (2013) criticize these measures on the grounds that there are other non-linear ways of controlling for unobserved trends and that this approach excludes other potential controls apart from those for the neighboring regions. Crucially, the parametric form of the model appears to be the critical determinant of whether a significant or insignificant impact of minimum wage on state employment is obtained.

Another potential problem of the minimum-wage literature mentioned above is the aggregation bias. This is because aggregate data might mask the real effect of minimum wage at the

individual level. Moreover, the analyses based on aggregate data could be affected by endogeneity as minimum wage movements could be caused by regional or national macroeconomic variables (Baskaya and Rubinstein, 2011; Sabia, 2014).

In this paper, we use the UK Labor Force Survey to estimate the causal impact of the UK national minimum wage (NMW) on employment using a non-parametric Bayesian modelling approach known as the Bayesian Additive Regression Trees (BART henceforth) that was originally developed by Chipman et al. (2010) and applied to the analysis of causal inference by Hill (2011). This procedure shares some similarities with standard matching estimation strategies, see for example Abadie and Imbens (2006), as it compares unemployment-to-employment and employment-to-unemployment transitions of individuals affected by the NMW with a similar individuals whose salary are not affected but are sufficiently close to those in the treatment group. However, this procedure has important advantages over other more traditional parametric specifications. Among them, it does not require any type of hypotheses or priors over the covariates to be included in the model, it can consider a large number of regressors, and it can estimate any type of interactive effects between the treatment variable and any other variable. Thus, under the BART model, the definition of the closest untreated individual for each treated individual and the interactions between the different clusters of individuals and time or and any other relevant covariate is not constrained to follow any specific (and potentially ad-hoc) parametric function. Furthermore, and more importantly, it need not be specified a priori.

Our paper is closely related to at least two previous works that estimate the impact of the UK NMW on employment at the individual level using micro-data: Stewart (2004) and Dickens and Draca (2005). Stewart (2004) analyzes how the introduction of the UK NMW in 1999 and its subsequent changes in 2000 and 2001 affected the employment-to-employment transition. Dickens and Draca (2005) follow a similar approach for the NMW increases in October 2003 but they extend the analysis to consider the separate effect of the NMW on job entry and job exit decisions. Both study the impact of the NMW by applying the difference-in-difference

technique to the UK Labor Force Survey data, and find both that the NMW does not have a significant adverse effect on employment. Unlike these papers, we do not consider a specific year's increase in the minimum wage but take into account all NMW changes since its introduction in 1999. Finally, our approach allows us to identify the interactions of the NMW effect with other relevant variables such as gender, age, qualifications and business cycle without the necessity of proposing a parametric specification.

The contribution of our paper is twofold. First, we shed some new light on the relationship between the minimum wage and employment. In particular, we find that the NMW exerts a positive and significant impact on both the probability job entry and job exit. Although the magnitude of the effect on job entry is larger, the overall effect of NMW is ambiguous as there are a lot of more employed workers than the number of the unemployed. This could explain the insignificant effect found in the previous work based on aggregate macroeconomic estimations. We find also that the effect is stronger for younger workers and in high unemployment periods. On the other hand, gender and qualifications play little role in shaping the minimum wage effect.

Second, we demonstrate the applicability of the BART approach to analyses of economic outcomes without imposing a specific parametric form a priori. While we chose the minimum wage effect on employment, this method could be applied to a broad range of other contexts equally well.

In the next section, we present the data used. Sections 3 and 4 discuss methodological approaches used for analyzing the labor-market impact of the minimum wage and explain the main features of the BART model, respectively. Empirical results are shown and discussed in Section 5. The final section summarizes our findings and offers some conclusions.

2 Data

Our analysis is based on the UK Labor Force Survey (LFS). The LFS is a quarterly nationally-representative survey of households across the UK. Each quarter, approximately 60 thousand

households and over 100 thousand individuals aged 16 and above are surveyed. Each household is retained in the survey for five consecutive quarters, with one-fifth of households replaced in each wave. The survey contains detailed demographic and socio-economic information on the respondents, including, importantly, their labor-market outcomes. Since the NMW was introduced in April 1999, we use all quarterly datasets available from April-June 1999 to October-December 2011, pooling all available LFS waves during this period. In order to have a sufficient number of observations, we include all individuals aged between 16 and 40.

The UK NMW features three different age-dependent rates: the 16-17 years old rate, the youth rate (applying to those aged 18-21¹), and the adult rate.² Historically, the youth rate has remained some 35% higher than the 16-17 rate while the adult rate has exceeded the youth rate by around 20%. The LFS reports the date of birth of every respondent and also the date the survey was carried out. By comparing these two dates, we can determine the precise age of each respondent on the day of the survey.³ We therefore know whether a particular individual is below or above the age threshold at which they become eligible for a different (higher) NMW rate.

3 Methodological Considerations

We analyze the effect of the NMW increases on employment by going beyond standard regression and matching estimation methodologies traditionally used for this purpose. Regardless of the methodology, the analysis involves comparing the changes in labor-market outcomes (such as employment) after a NMW change for the treatment and control groups.⁴

¹ The upper limit for the youth rate has been lowered to 20 from October 2010. Where relevant, our analysis takes this change into account.

² A fourth rate, for apprentice workers, was introduced in October 2010 (we do not consider those subject to this rate in our analysis). No minimum wage applies to those who belong to one of the few exemptions such as members of the armed forces, volunteers, students on work placements, workers living in the employers' households, and (until 2010) apprentices.

³ The precise date of birth is not available in the publicly released LFS datasets. We are grateful to the Office for National Statistics for making the restricted release of the LFS available to us.

⁴ Note that our approach is similar in spirit to the difference-in-difference approach in Stewart (2004) and Dickens and Draca (2005) who compare the average change in the employment status before and after the introduction of a very specific minimum wage policy. Given that we consider a whole sample of minimum wage changes through thirteen years, our analysis is based on the estimation of the effect on the

Consider the impact of NMW on the probability of job loss. The treatment group comprises workers whose wages have to go up in the wake of an annual NMW increase because the new NMW rate is higher than their current wage. The wages of those in the control group should be close to but just above the new rate so as not to have to change.

More specifically, the treatment group can be defined as the individuals whose wages meet the following condition:

$$nmw_t < w_{it} < nmw_{t+1} \quad (1)$$

where nmw_t is the (age-dependent) NMW rate in effect at time t while w_{it} is the worker i 's wage. The control group is defined as the workers whose wage before the increase is greater than the new NMW rate but lower than some upper bound to ensure that we only consider workers earning just above the minimum wage, so that are likely to possess similar characteristics as those earning the minimum wage. If we set the upper bound as c above the new rate, the control group comprises workers meeting the following condition:

$$nmw_{t+1} \leq w_{it} < nmw_{t+1} * (1 + c) \quad (2)$$

We can then estimate the following equation

$$P(e_{t+1} = 0 | e_t = 1) = \Phi(\alpha * D_i + \gamma * X_{it}) \quad (3)$$

where the dependent variable is the probability that individual i is unemployed conditional on being employed in the preceding quarter, $\Phi(.)$ is the standard normal cumulative distribution function, D_i is a dummy variable denoting individuals belonging to the treatment group, included on its own and in interaction with the gap between individual i 's wage and the new NMW rate, and X_{it} collects all remaining covariates (individual socio-economic characteristics and time effects). An analogous equation can be estimated for the probability of remaining employed conditional on employment in the previous quarter. In line with the standard practice,

change in employment status, before and after the policy application, for the treatment and control groups.

equation (3), and in particular the coefficient estimate of the first term, is interpreted as capturing the differentiated effect of the minimum-wage increase on the probability of becoming unemployed for the treated individuals relative to those in the control group.

A similar approach can be used to estimate the impact of NMW on the probability of job entry. In this case the equation to estimate is

$$P(e_{t+1} = 1 | e_t = 0) = \Phi(\alpha * D_i + \gamma * X) \quad (4)$$

A particular problem presents itself here in the fact that we do not have any previous wage information for those who enter employment only after the NMW increase. In other words, we do not know whether those entering into employment after the increase would have earned more or less than the minimum wage before the increase. Dickens and Draca (2005) resolve this by defining the treatment group as those whose earnings are less than or equal to the (age-relevant) new NMW rate and the control group as those who earn up to c percent above the NMW:

$$\text{Treatment group: } w_{t+1} \leq nmw_{t+1} \quad (5)$$

$$\text{Control group: } nmw_{t+1} < w_{t+1} < nmw_{t+1} * (1 + c) \quad (6)$$

A somewhat uncomfortable implication of this specification is that the treatment group now includes also those who earn less than the NMW (there are specific cases when this is allowed, for example for apprentices or for those who receive employer-provided accommodation or other in-kind payments). An alternative specification would entail constructing the treatment group as including only those who earn the minimum wage after the NMW increase. Using that specification yields very similar results.

Note that our analysis could suffer from a potential endogeneity problem as in a non-experimental sample, such as the Labor Force Survey, workers earning less than the new NMW rate are more likely to lose their jobs even if NMW does not change because they are likely to

be less productive than workers earning higher wages.⁵ Therefore, it is the characteristics associated with their lower wages (and not the minimum wage itself) that determine their higher probability of job loss compared to other individuals with above-NMW wages. In other words, if wages are not allocated randomly, the allocation of individuals into treatment and control groups is not random either but depends on their characteristics. In order to assume that the outcome is independent of the treatment, it is necessary to account for all possible conditioning factors by including a broad range of covariates, X . More specifically, the strong ignorability hypothesis with respect to the allocation of treatment states that Y is conditionally independent of D given X and that the probability of treatment allocation is always positive regardless of the specific value of X . Under this hypothesis, the estimation of the marginal effects associated to the treatment variable can be considered in general as a consistent and unbiased estimation of the causal effect of NMW on the probability of job exit and job entry: including a relevant set of covariates in equations (3) and (4) is a sufficient condition to ensure an unbiased estimation. However, as argued by Morgan and Winship (2007), the regression approach can be subject to two important drawbacks. The first relates to the fact that the causal effect of NMW is not constant across individuals. In this case, the estimated causal effect represent a conditional variance weighted estimate of causal effects of individuals and the causal estimation is only unbiased and consistent for this particularly weighted average that is not usually the parameter of interest. The second problem relates to the fact that the strong ignorability condition does not necessarily imply that treatment is uncorrelated with the error term net of adjustment for X as this error term depends on the specification of covariates, X . Therefore, in order to interpret the estimation of a regression strategy as a reliable causal effect, we require a fully flexible parameterization of X .

An alternative approach that overcomes the drawbacks mentioned above is matching: comparing the labor outcomes of the treated individuals with those of similar individuals, with similarity determined based on the set of variables X . Let Y_i be the response, that is $Y_i = 1$ if job

⁵ Note however that in expressions (2) and (6), small values of c would imply that the salary of the treated and control groups could be deemed to be very similar.

exit or job entry is observed and $Y_i = 0$ otherwise; and D_i be an indicator of whether the individual belongs to the treated or control group. In order to compute the causal effect of D_i on the response variable Y , we would need to know the outcome of interest for the same individual if treated, $Y_i(1)$, and if not treated, $Y_i(0)$. However, this is impossible because only one of them can be observed at any given point in time. The counterfactual result, therefore, has to be estimated with a regression model. In this case, we estimate the response Y to a “hypothetical treatment” D_i .

There are two standard approaches to estimate this causal impact. One is to compare the outcome variable of a treated individual with that of one or several nonparticipants as similar as possible to the participant considering the values of covariates X_{it} . A second approach matches participants and nonparticipants based on their estimated propensity scores. However, the application of these methodologies is only possible if there is a region of common support between the treatment and control groups.

Regardless of the approach used, the average treatment effect is defined as $ATE = E[Y(1) - Y(0)]$, where the expected value is computed with respect to the probability distribution of Y for all individuals. We focus on the causal effect for a given set of individuals, for example those who have received the treatment, $E[Y(1) - Y(0)|D = 1]$, that is, individuals affected by NMW increases. In this case, the expected value is estimated with respect to the conditional distribution of $(Y|D = 1)$. Even more generally, if we have a set of covariates X we can estimate the causal effect conditional on them, that is, conditional on $X = x$.

However, this is not always possible because matrix X typically has a very high dimensionality and comprises a wide range of covariates, including qualitative and quantitative variables, and some standard approaches such as, for example, the propensity score, cannot be applied if the number of covariates is too high. This forces the analyst to consider a set of variables of lower dimension, putting the strong ignorability assumption in doubt.⁶ Besides, the specification of

⁶ See Caliendo and Kopeining (2008) and references therein for a discussion on this issue.

regression models with many variables makes it not practical to consider all possible interactions among the variables. Again, this forces the analyst to consider only interactive effects among first or second order covariates or to use algorithms such as the forward or backward variable selection that may provide locally optimal models. Unfortunately, there is no theoretical justification, only empirical results, to guide us in assessing the scope of a local instead of a global optimum.

Due to these drawbacks, we make use of a particular type of matching estimation based on the BART model for the estimation of causal impact of NMW increases. Being a non-parametric model, this frees us from being restricted by a given model specification. Furthermore, it allows us to estimate with a satisfactory precision the response of the variable of interest to NMW increases, and with that, the counterfactual result even for a high dimensional X . An additional important advantage of this approach is that it allows for identification of the most significant interactive effects between the treatment variable and any of the covariates without being constrained to include these interactions in any parametric form.

4 BART Model

In the explanation of the model, we mainly follow the notation of Hill (2011) and references therein (see Chipman et al., 2010, for details of the statistical model, and Leonti et al., 2011, for an application of this model to the estimation of a causal effect of the use of medical plants). Let $Data$ be the available data, that is the set Y, X, D observed for N individuals and $\pi(\cdot | \cdot)$ the probability distribution of the left argument conditional to the right argument. The aim of the analysis is to estimate the *posterior* probability distribution of the causal effect, that is $\pi(ATE|Data)$, or the posterior distribution conditional on some covariates, $\pi(ATE|Data, X = x)$. In order to do this we use a non parametric regression model. The novelty in these types of causal inference analyses is the use of a Bayesian regression model known as BART. As in all Bayesian models, we need a likelihood function defined for a set of parameters, $\theta \in \Theta \notin \mathbb{R}$, and a prior distribution $\pi(\theta)$, $\theta \in \Theta$. The likelihood function, $L(Y|X, D, \theta)$, is obtained from the

following additive regression model, where the mean of Y , $E(Y) = P_i(Y = 1) = P$, is determined from the sum of estimated models for the response variable:

$$P = \Phi\left[\sum_{j=1}^m g(X, D; T_j, M_j)\right] \quad (7)$$

where $g(X, D; T_j, M_j)$ is a classification tree with the variables and split points represented by T_j and the terminal nodes denoted by M_j and computed with respect to the values x, D that belong to the individual whose response is Y . Essentially, g is a function that gives to each individual i their expected value in the j^{th} tree, $\mu_{ij} \in M_j$. The final score estimated for the i^{th} individual would correspond to the average of the m scores. It is well known that, in order to minimize the forecast error, classification trees tend to grow disproportionately until generating *overfitting* in the response and that in general an estimator obtained from many simple trees is more efficient than another one obtained from a single complex tree. Examples of these types of models are Boosting (Shapire and Singer, 1999) and Random Forest (Breiman, 2001).

In order to achieve this, it is necessary to use a regularization prior on the size of the tree $\pi(T, M)$ as specified in Chipman et al (2010). This regularization prior precludes the tree from growing too much and makes sure that each of the μ_{ij} contributes in a marginal way to the estimation of the response function. As Chipman et al (2010) show, the hyper parameters of all prior distributions are specified in relation to the observed sample. It produces priors that are dependent on the sample. This procedure, which is not very orthodox from a Bayesian point of view, is part of the approaches known as empirical Bayes that are very popular and have been enhanced from a theoretical point of view by Petrone et al. (2013). As explained by Hill (2011), the results of this type of analysis are robust with respect to prior modifications.

Using the priors specified above it is possible to simulate samples of the posterior distribution with a non-excessive computational effort using Markov Chain Monte Carlo (MCMC), more specifically using Metropolis Hastings. In particular, the distribution used to update the values of T_j and M_j consists of adding/dropping a terminal node and changing a split variable or a split point with the probabilities specified in Chipman et al. (2010). Once the posterior distribution of

$\theta = (T_1, \dots, T_m, M_1, \dots, M_m)$ has been obtained, the predictive distribution for the probability of job exit is:

$$m(Y_i|x_i, z_i) = \int_{\theta \in \Theta} L(Y_i; \theta) d\pi(\theta | \pi(\theta)) \quad (8)$$

where L is the likelihood function for $\theta \in \Theta$ and $\pi(\theta)$ is the posterior. Integral (8) is practically estimated by generating values of $P_i = P_i(Y_i = 1)$, using the normal distribution with the mean and variance for each value θ in the chain MCMC and the regression tress computed in x_i and z_i . In particular we use $m=500$ trees and 5000 MCMC steps after an initial burn-in of 1000 steps.

In this way, the distribution for each individual and the corresponding counterfactual response can be estimated simply by estimating the response in $D_i = 1$ if the worker is affected by NMW and in $D_i = 0$ otherwise. Once these predictive posterior distributions have been obtained, the difference between the factual and counterfactual responses are considered to obtain the distribution of the individual causal effect. Finally, $\pi(\text{ATE} \setminus \Omega)$ is estimated from the set of the differences for all the individuals. Then, the estimation of the conditional causal effect is required, this is obtained simply by considering the difference for the individuals that fulfill the condition $X = x$.

5 Results

As a first step, and to establish a benchmark to compare our results against, we report the results of a probit model as specified in Equations (3) and (4), where job entry and job exit are functions of the dummy variable for the treatment along with and a set of covariates (Table 1).⁷ In this regression, the parameter c defined in the previous section is set to be 0.1 to ensure that treatment and control individuals are comparable in terms of wages but the results are qualitatively similar when we consider $c = 0.3$, $c = 0.5$ and $c = 1$. The last row of Table 1

⁷ Besides standard socio-economic characteristics, we also include an indicator to account for the fact that the age limit for the adult rate was lowered from 22 to 21 from October 2010.

indicates that the probabilities of job entry and job exit are both positively correlated with being in the treatment group.

It is interesting to compare these results with those obtained with a standard matching procedure such as the propensity score. The estimation results are qualitatively, and even quantitatively, similar to those obtained from a regression probit model. More specifically, the estimated causal impact for job entry is 0.051 with standard deviation 0.013 while that for job exit is 0.03 with standard deviation 0.011.⁸ The fact that the two sets of results are very similar is not surprising as the matching estimation can be interpreted as being similar to a regression that puts more weight on the observations in the treatment and control groups that are very similar to each other.

Table 1. Probability of job entry and job exit as a function of treatment and change.

	Probit regression	
	Job entry	Job exit
LR Chi2	445.13	400.17
Prob > chi2	0.0000	0.0000
Number of obs	7792	6746
Treatment	.0517764** (.00961)	.024584** (.0086)

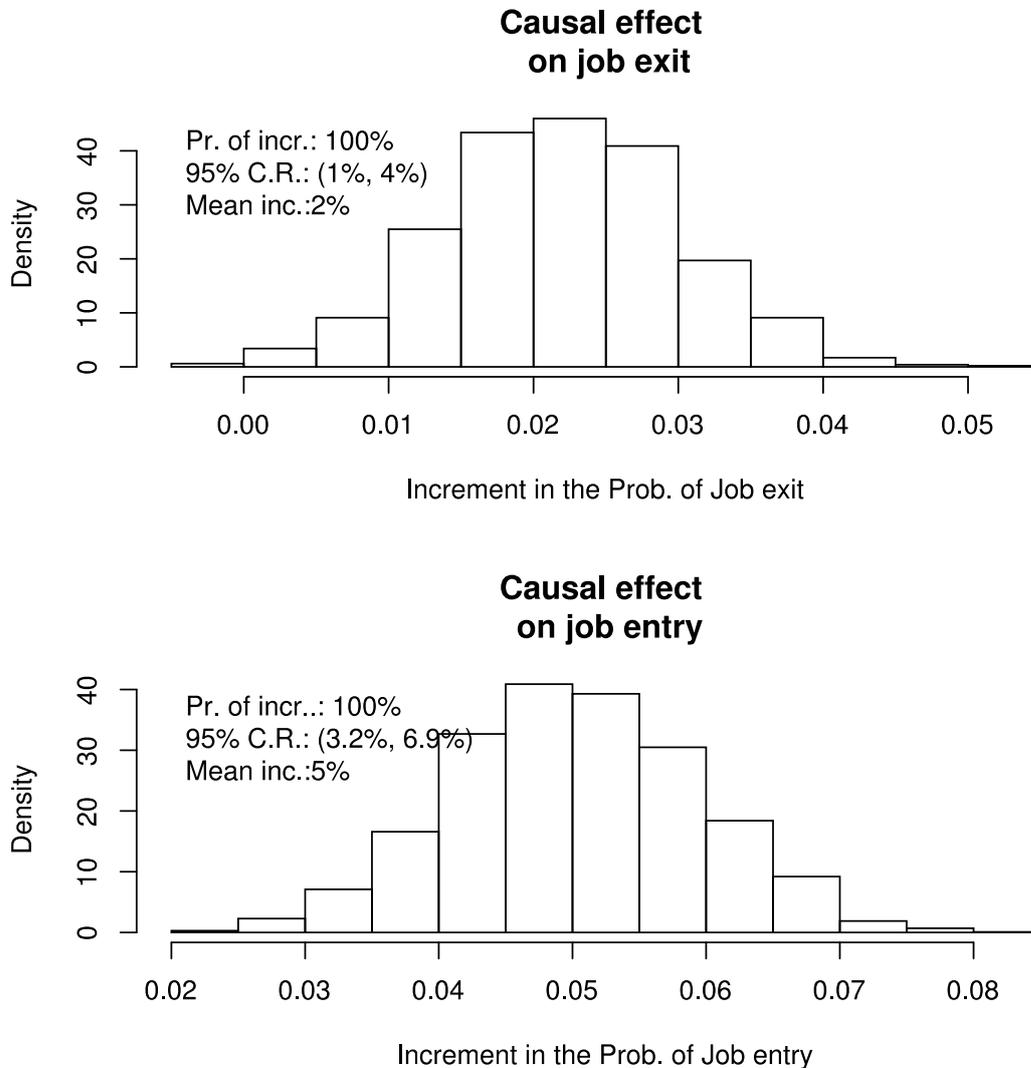
Notes: Marginal effects evaluated at mean values.

The Bayesian approach considered here is instead based on the estimation of the expected value of the treatment and control groups using the same explanatory variables in both cases. Figures 1 and 2 report the estimated distribution of the total causal impact of increases in the minimum wage rate on job exit and job entry using the BART model with all workers aged 18-40. The results indicate that treatment has positive effects both on job entry and job exit, in a manner similar to the probit results reported above. More specifically, the NMW exerts a positive impact on job entry, and the mean value of this causal impact is 5% with a 95% confidence interval equal to [3.2%, 6.9%]. For job exit, the effect is positive with the mean value equal to 2% and with a 95% confidence interval equal to [1%, 4%]. Note that although the estimated effect on job entry is larger than that for job exit, the overall effect of NMW is ambiguous as

⁸ These results are available upon request.

there are many more employed workers (who are candidates for job exit) than unemployed individuals (candidates for job entry).

Figure 1 Posterior distribution of the causal effect of job entry and job exit



As discussed above, one of the most important advantages of the BART approach is that it allows for the simultaneous estimation of any kind of interaction between the treatment variable and any of the covariates. This is possible either at model estimation and at description level of obtained results. Here we consider the result description level by inspecting the interaction between covariates and the estimated causal effect. In particular, the interaction with continuous variables is evaluated through boxplots, which include 95% percentile bootstrap confidence

intervals for median, while that with continuous covariates by local polynomial regression smoother (loess) along with their 95% confidence intervals (Cleveland et al., 1992, Chp. 8).

In Figure 2, we interact gender with the effect of NMW increases. Again, the previous finding of a greater effect of NMW increases on job exit than on job entry is reproduced. Although for job entry it is clear that the median values are significantly different, the whole distribution of the two effects are very similar which suggests that gender plays little role. In Figure 3, in turn, we consider the interaction with age (expressed in months rather than years). Here, the pattern is different for job exit and entry. While the causal impact of NMW is decreasing with age in both cases, that decline is much steeper for job entry. This is not surprising, given that young workers are more vulnerable to NMW increases. Besides, the interactive effect is clearly stronger for job entry. In Figure 4, we consider the interaction with the highest attained qualification. Again, although it is possible to observe significantly different mean values associated to the different qualifications, the whole distribution of the estimated causal effect indicates that this variable is not a relevant factor to explain differences in the causal impact of NMW either for job entry or job exit. Finally, Figure 5 presents the interaction with the regional business cycle – measured using the unemployment rate. Interestingly, this interaction effect is very different for the two labor-market flows: the minimum-wage effect on job exit is relatively low and depends little on the regional unemployment rate, whereas that for job entry is higher and positively related to regional unemployment. This implies that the effect of the minimum wage on job entry differs considerably between recessions and booms, whereas the business cycle has little bearing on how the minimum wage shapes job exits.

Finally, so far we have been considering the effects of NMW changes for two similar groups, those affected by the change and those who are unaffected but are otherwise similar to the affected individuals both in terms of their wage and in terms of the other covariates used in the analysis. However, to test for the robustness of our results even further, we carry out a falsification experiment whereby we define the treatment and control groups as if the NMW were equal to the actual NMW plus 2£. The results of this experiment, shown in Figure 6,

indicate that the causal impact of the (false) NMW increase is significant at the 5% level for job entry but not for job exit. We find similar conclusions for other artificial NMW (results are available from the authors upon request).

The fact that the falsification test is significant for job entry decisions could be due to spillover effects of the actual NMW increase: the NMW change can lead to ripple effects even for wage rates above the minimum wage. Importantly, the insignificant falsification test results for job exit give strong support to the finding that the employment of workers earning the minimum wage is adversely affected by NMW increases.

Figure 2 NMW Increases, Job Exit/Entry and Gender

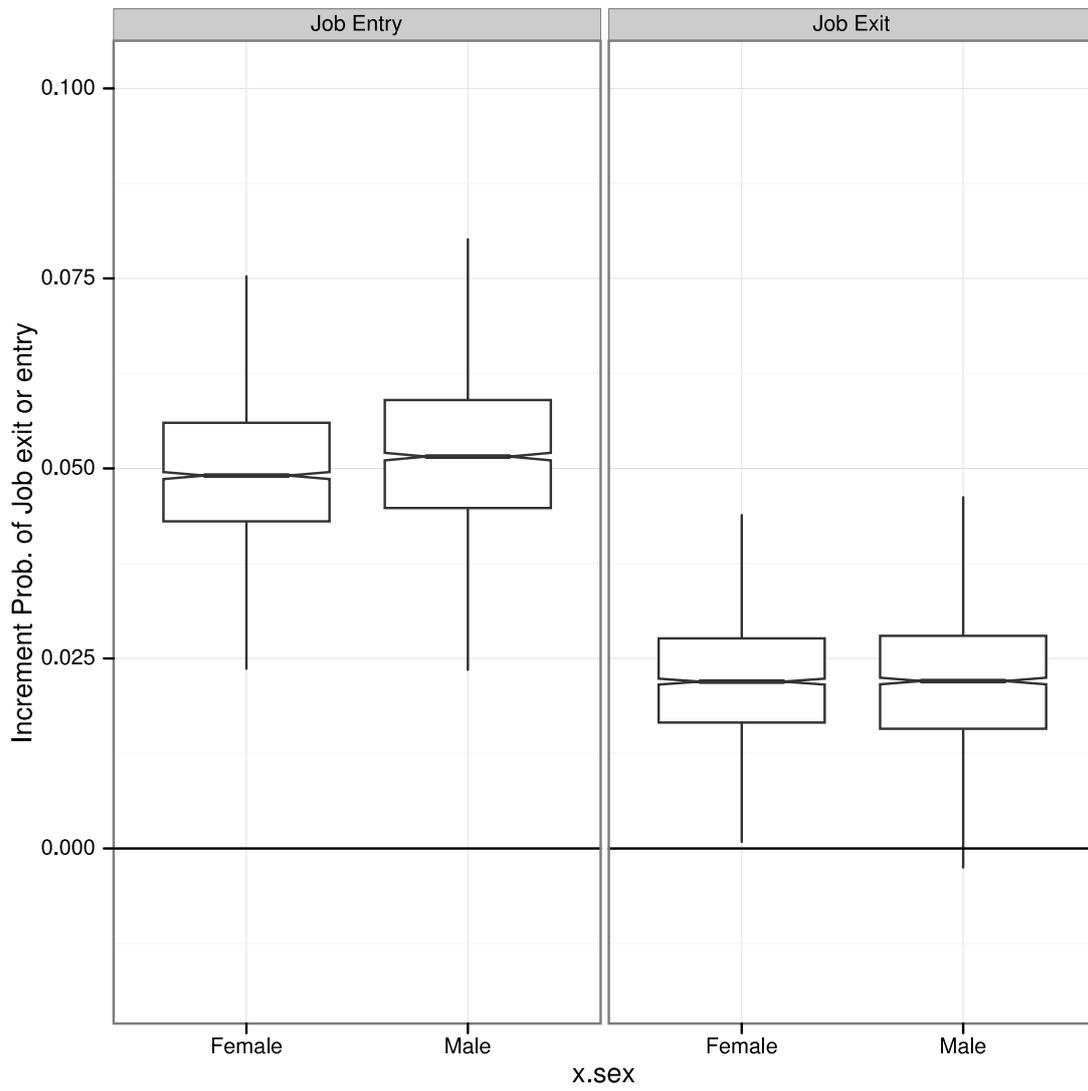
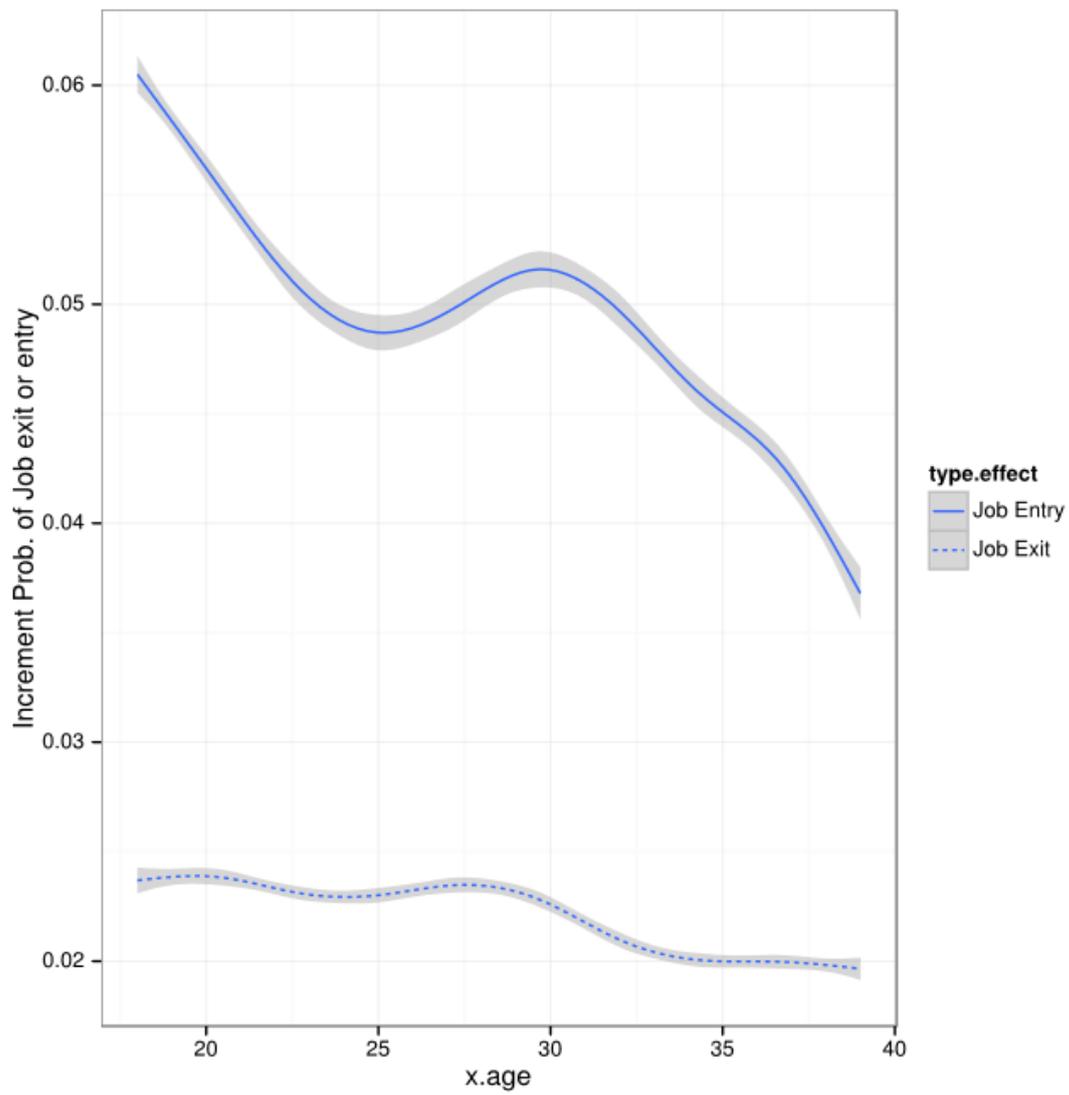
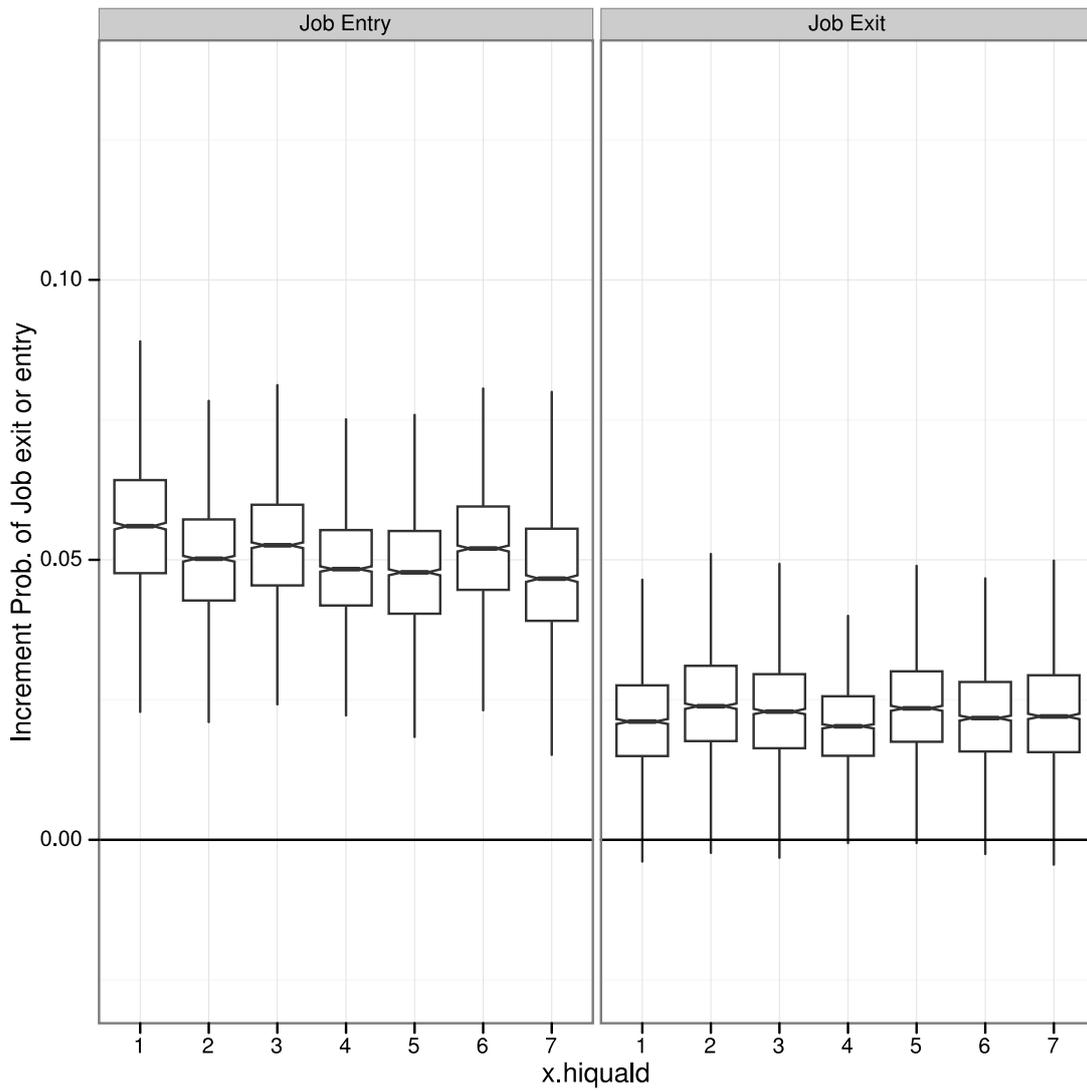


Figure 3 NMW Increases, Job Exit/Entry and Age



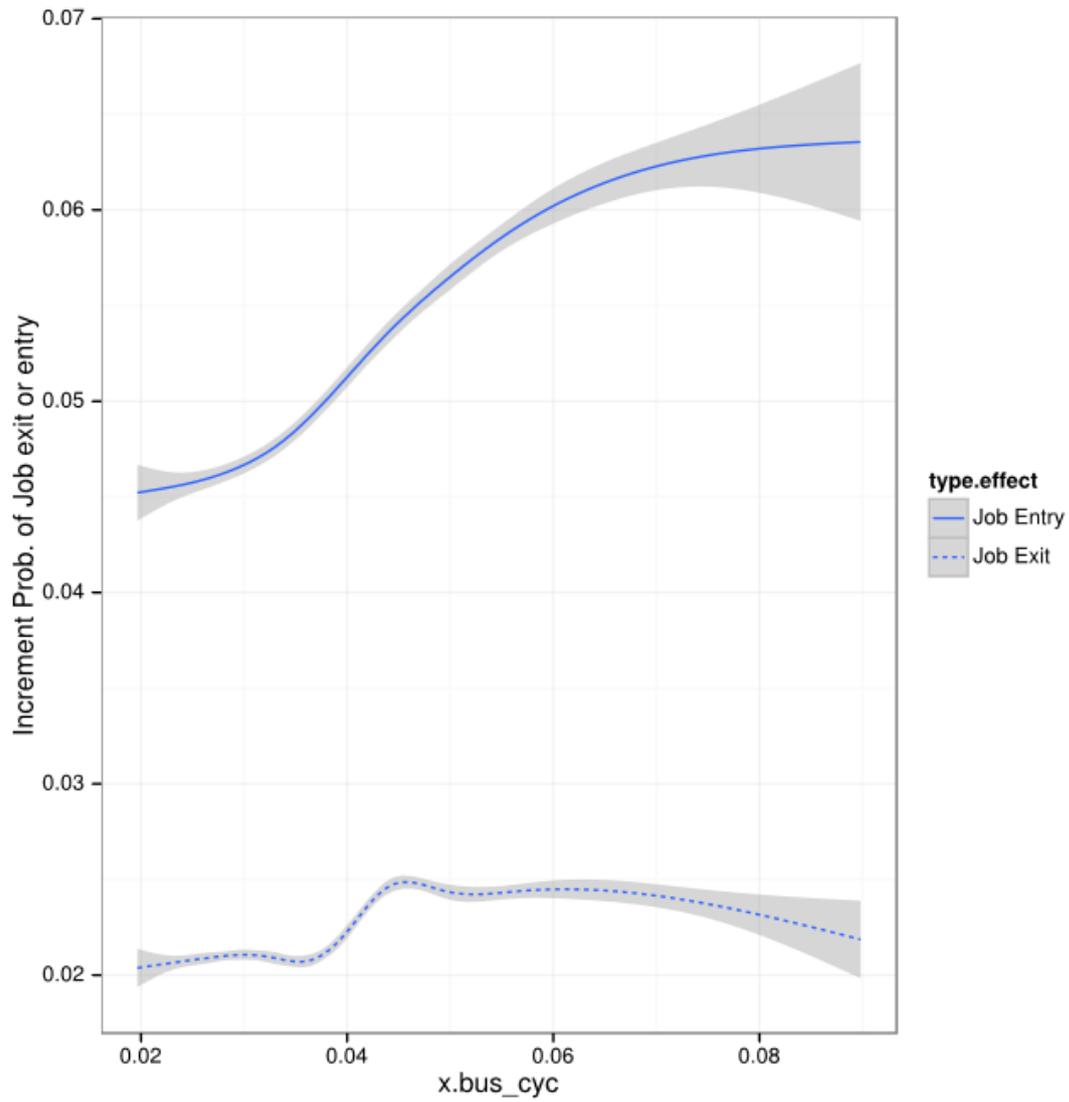
Note: Shadow area indicates the 95% confidence interval of the local polynomial regression estimator (loess).

Figure 4 NMW Increases, Job Exit/Entry and Qualifications



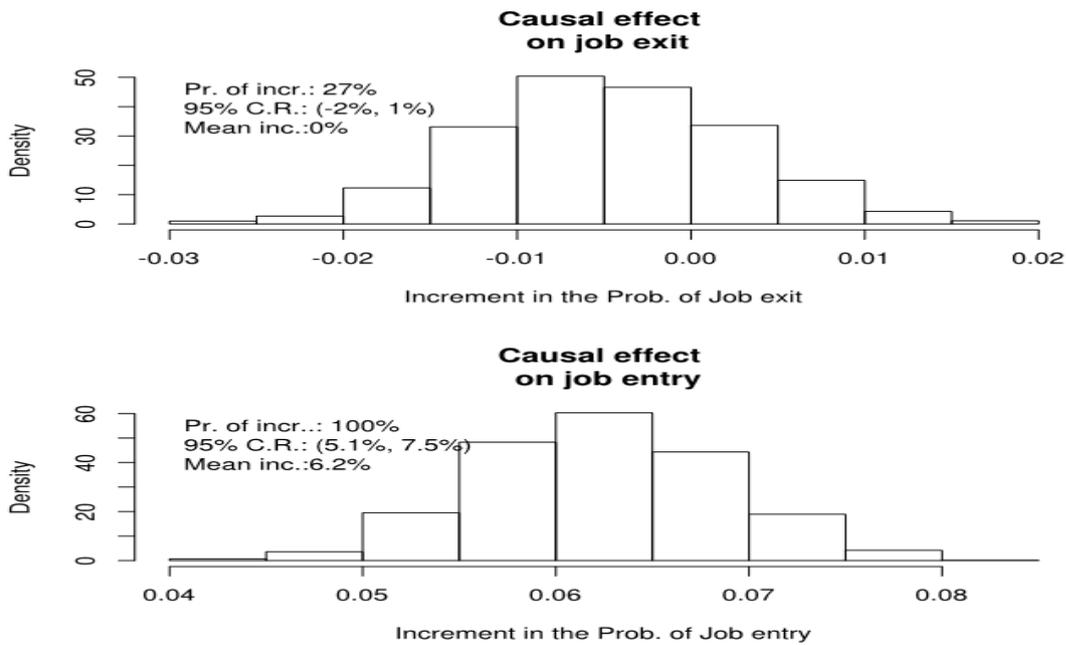
Notes: 1 Degree or equivalent, 2 Higher education, 3 GCE A Level or equivalent, 4 GCSE grades A-C or equivalent, 5 Other qualifications, 6 No qualification, 7 Don't know

Figure 5 NMW Increases, Job Exit/Entry and Business Cycle



Notes: The horizontal axis measures the regional unemployment rate. Shadow area indicates the 95% confidence interval of the local polynomial regression estimator (loess).

Figure 6 Falsification Test: NMW + £2



Notes: The falsification test simulates the NMW being £2 higher than the actual value.

6 Concluding remarks

We estimate the causal impact of the NMW on the probability of job entry and job exit in the UK, applying a novel methodology to this context, the Bayesian Additive Regression Trees (BART). An important advantage of this procedure is that it allows the identification of the most important interactions between the treatment variable and other covariates in the model. We find that the NMW exerts a significantly positive effect both on job entry and job exit, with the impact on job entry being relatively stronger (given that there are fewer unemployed than employed workers, the absolute size of the flows cannot be readily compared). The causal effect of NMW is found to be higher for young workers and in periods of high unemployment; both of these interactions are more prominent for job entry than for job exit. However, no significant interactions were found with gender and worker qualification. Overall, the effect of NMW is stronger for job entry than for job exit.

Most importantly, our paper opens new lines of research that can be explored in subsequent work. For example, this fully flexible approach could be adapted to deal with some recent issues in the literature about the importance of the econometric specification on estimating the effect of minimum wage using panel data models in US states. Also, it could be used to estimate the possible interactions between the federal minimum wage and the state minimum wages, as done by Baskaya and Rubinstein (2012), without the necessity of estimating two different models.

References

- Abadie, A., and G.W. Imbens (2006). "Large Sample Properties of Matching Estimators for Average Treatment Effects," *Econometrica* 74(1), 235-267.
- Allegretto, S., Dube, A., Reich, M., (2011), "Do minimum wages really reduce teen employment? accounting for heterogeneity and selectivity in state panel data." *Industrial Relations* 50, 205-240.
- Allegretto, S., Dube, A., Reich, M. and Zipperer, B., (2013), "Credible research design for minimum wage studies." Working paper, Institute for Research on Labor and Employment, UC Berkeley.
- Baskaya, Y., Rubinstein, Y., (2012), "Using federal minimum wages to identify the impact of minimum wages on employment and earnings across the U.S. States." Unpublished paper, Central Bank of Turkey.
- Breiman, L. (2001). "Random forests". *Machine Learning* 45, 5-32.
- Card, D. and Krueger, A.B. (1995), "Time-series minimum wage studies: a meta-analysis." *American Economic Review* 85, 238-243.
- Caliendo, M. and Kopeining, S. (2008), "Some practical guidance for the implementation of propensity score matching." *Journal of Economic Surveys* 22, 31-72.
- Chipman H.A. , George E.I., and Mculloch R. (2010). "BART: Bayesian Additive Regression Trees", *Annals of Applied Statistics* 4(1), 266–298.
- Cleveland W. S., Grosse E. and Shyu W. M. (1992) Local regression models. Chapter 8 of *Statistical Models in S* eds J.M. Chambers and T.J. Hastie, Wadsworth & Brooks/Cole.
- Dickens, R. and M. Draca (2005), "The Employment Effects of the October 2003 Increase in the National Minimum Wage," CEP Discussion Paper No 693, Centre for Economic Performance, London School of Economics.
- Doucouliagos, H. and Stanley, T.D. (2009), "Publication selection bias in minimum-wage research? a meta-regression analysis." *British Journal of Industrial Relations* 47, 406-428.
- Dube, A., Lester, T.W., and Reich, M. (2010), "Minimum wage across state borders: estimates using contiguous counties." *Review of Economics and Statistics* 92, 945-64.
- Hill, L.H. (2011). "Bayesian Nonparametric Modeling for Causal Inference", *Journal of Computational and Graphical Statistics* 20(1), 217-240.
- Leonti, M., S. Cabras, C. Weckerle, M. Solinas, and C. L. (2010). "The causal dependence of present plant knowledge on herbals—contemporary medicinal plant use in campania (italy) compared to matthioli (1568)", *Journal of Ethnopharmacology* 130 (2), 379–391.
- Morgan, S.L., and C. Winship (2007), *Counterfactuals and Causal Inference: Methods and Principles for Social Research*. Cambridge University Press.
- Neumark, D. and Wascher, W. (2007), "Minimum wages on employment." *Foundations and Trends in Macroeconomics* 3, 1-182.
- Neumark, D., Salas, I.M.I, and Wascher, W. (2014), "Revisiting the minimum wage employment debate: throwing out the baby with the bathwater?" *Industrial and Labor Relations Review* 67, 608-648.
- Petrone S, Rousseau J, and Scricciolo S. (2014). "Bayes and empirical Bayes: Do they merge?" *Biometrika* (to appear).
- Sabia, J.J. (2014), "Minimum Wages and the Business Cycle: Does a Wage Hike Hurt More in a Weak Economy?" Employment Policies Institute: Washington, DC.
- Schapire, R.E. and Singer, Y. (1999). "Improved boosting algorithms using confidence-rated predictions", *Machine Learning* 37, 297-336.
- Stewart, M.B. (2004). "The employment effects of the national minimum wage", *Economic Journal* 114, 110-116.

Please note:

You are most sincerely encouraged to participate in the open assessment of this discussion paper. You can do so by either recommending the paper or by posting your comments.

Please go to:

<http://www.economics-ejournal.org/economics/discussionpapers/2016-17>

The Editor