

## Response to Referee 1

We are grateful to the referee for his/her comments. Below, we repeat them and provide our response.

*(a) The literature goes back to 2014. In the meantime, further studies were published – e.g., for the UK and the US. The authors should make references to this literature.*

We have updated the literature to include more recent contributions.

*(b) The used covariates should be described in the data section. E.g., it would be important to consider the economic sectors, at least to distinguish between employment in the private and the public sector, mainly because the way how firms react on the MW could depend on their economic activities and the market where they are active.*

We added a table of descriptive statistics (Table 1 on p. 9).

*(c) In the analysis for the job exits as outcome, the authors tried to compare similar workers that are paid “around” the NMW with a treatment group of workers paid above the existing NMW but below the next adjusted level of the NMW with a control group of workers that are already paid at or above this latter level of the NMW. This implies that relevant but very specific groups are analysed. This also means that the computed probabilities refer to this specific groups and not to the employment population. This should be clarified. For the analysis for the job exits, the authors compare job entry probabilities of former unemployed now paid at or below the current NMW level with job entry probabilities of former unemployed that are paid above the current NMW level. The same argument applies: the results for the job entry probabilities are specific to the treatment group and this should be clarified. Maybe this is the reason why the authors hesitated to compare predicted employment exit and entry flows based on the whole employment and unemployment population, though this would presumably be possible.*

The interpretation of the results has been amended, to make it clear that the findings only apply to low-wage workers.

*(d) Referring to the specific treatment and control groups the authors should reconsider their concluding remarks (part. the last sentence on page 21): obviously the job entry and job exit probabilities do not refer to the whole working population and, thus, are hardly comparable.*

The conclusions have also been amended.

## Response to Referee 2

We are grateful to the referee for his/her comments. Below, we repeat them (slightly abridged) and provide our response.

*First, in order to learn more about the effects of the NMW, much more background information has to be provided. How did the NMW develop over time? How many and what type of workers were affected by each increase? How did wages of low-wage workers react to increases in the NMW? What has already been established in the existing (huge) literature on the NMW? What are the theoretical hypotheses and the existing empirical findings in terms of labour market transition and turnover at the firm level? The last point is crucial since the usual argument is that a minimum wage decreases turnover – implying lower job entry and job exit rates. Since the authors find increases in both rates, it should be explained in detail why the results deviate from other studies.*

We added information on the evolution of the NMW. We also added additional references to our discussion of the literature. However, the literature is indeed very voluminous and it is not the objective of this paper to offer a comprehensive review of the literature.

*The provision of such background information would also help to interpret the results in terms of economic significance. At the moment, the authors only provide information on “the causal impact of the NMW”. What is actually provided is the average marginal effect of all increases in the NMW. In order to interpret the magnitude of the effect, each increase in the NMW should be known. In this context it should also be discussed that the relationship between minimum wage increases and employment outcomes (including job entry and job exit) is most likely non-linear. Stated differently, modest increases in the minimum wage might have quite different effects compared to strong increases. This is especially important given the long observation period of 12 years, over which all increases are pooled. This is problematic to the extent that other policies might have been implemented that affect labour turnover of low-wage workers in addition to the minimum wage. This should be discussed explicitly.*

This is a good point. We have added a brief discussion of the effect of the size of NMW increase (see Figure 3), which shows that accounting for size does not change our main conclusions.

*Further, the authors should additionally include estimations with aggregate (un-)employment as an outcome variable to be able to relate their findings to the general discussion on employment effects of minimum wages. Alternatively, or additionally, aggregate statistics on (un-)employment could be used to make predictions on overall employment effects based on the estimated effects on labour market flows. In this context, the heterogeneity of the effect across the business cycle is an interesting result. Given that the heterogeneity in the employment effect across age groups is well established, it might be worthwhile to focus on this.*

Our analysis is carried out at the individual level and as such it explains individual variation in labor-market outcomes. We include year dummies in our analysis, which control for for all sources of macroeconomic heterogeneity: inflation, GDP growth, unemployment, etc.

*Second, and focusing on methodological aspects, the control group is problematic. The authors opt to use workers whose wages are up to 10 percent higher than the NMW in the next period. At the same time, spillover effects of the NMW on higher wage groups have been well documented in the literature. This implies that the control group itself is affected by the minimum wage leading to biased estimates. Indeed, a provided falsification test confirms this problem at least for the job entry rate: If the NMW is (artificially) set 2£ higher than the true NMW, the results remain significant. Though it remains unclear, if the true treatment group, i.e. those workers earning exactly the minimum wage, are excluded in this falsification tests. If not, the significant results is not really surprising. In addition, the treatment group for the job entry rate (those earning the NMW or less after moving from non-employment to employment) includes individuals who are exempted from the minimum wage. Since the*

*treatment and the control group are not well defined for the job entry rate, the corresponding results should be interpreted very carefully – if at all. Another question is how precisely hourly wages are measured in the LFS, which constitutes an important condition to correctly define the treatment and control group. In summary, BART is no improvement on existing studies if causality is not ensured.*

This is a very good suggestion. We have extended the analysis to use an alternative control group, one that is not immediately adjacent, in the wage distribution, to the treatment group. The results are summarized in Figure 8 and the accompanying discussion. We argue in the paper that using this alternative control group does not alter our findings in a significant way.

*Concerning BART, I'm not sure if the non-parametric estimation is really an important improvement or a simple technical exercise. I would be more convinced if the authors provided a more detailed discussion on existing techniques, their exact assumptions and plausible reasons for deviations from these assumptions. The authors actually show that BART does not make a big differences by providing results for a probit model with matching that delivers almost exactly the same estimates (compare Table 1 and Figure 1). Thus the distributional assumptions and specification assumptions needed for a probit model seem to be more or less correct. The authors additionally state that BART has the advantage of being able to study the interaction of covariates with the NMW. However, the provided heterogeneities (age, gender, qualification, business cycle) are standard and can be easily implemented with parametric techniques, such as RDD, differences-in-differences or matching.*

As we argue in the paper, the main advantage of BART is its flexibility in terms of covariates and interactive effects that can be considered. Therefore, even if some particular sets of results are similar to those obtained with other methods, the greater flexibility of the BART model is an important advantage. The fact that the results we report do not differ much from those obtained with other methods should serve as a reassurance rather than undermine the case for using this method.

*In either case, BART should be explained in more detail as the average reader interested in a paper on the NMW might not be too familiar with this method. For example, it remains unclear to me why the mean of the posterior distribution is interpreted in the baseline specification (Figure 1), but when looking at heterogeneities (Figure 2 – Figure 5), the spread of the posterior distribution is interpreted instead of the mean.*

We have elaborated on the explanation of the model and hope these issues are clearer now.