

Referee Report for

Minimum Wage and Employment: Escaping the Parametric Straightjacket

This paper analyses the effect of all increases in the UK national minimum wage (NMW) on the job entry and job exit rate over the observation period 1999-2011. The estimations are based on the UK Labour Force Survey (LFS) and make use of Bayesian Additive Regression Trees (BART). This method is novel to the minimum wage literature and has the advantage that fewer specification assumptions are needed compared to standard parametric techniques. The authors find positive effects on both transition rates, but do not make any statements on the overall employment effect. Further, heterogeneous effects are found for different age groups (younger workers are affected more strongly) and across the business cycle.

The contribution of this paper can be seen along two dimensions. First, the authors focus on labour market flows (job entry and job exit) instead of employment levels. While this is not the first paper to focus on these outcomes, the results may help to improve our understanding of the mechanics of the minimum wage leading to aggregate results (increasing or decreasing employment). Second, the paper could contribute to the minimum wage literature by explaining biases in previous studies due to specification assumptions, thereby resolving conflicting existing results. Unfortunately, the current version of the paper has to be revised considerably in order to offer one (or both) of these contributions. In what follows, I will therefore focus on the two possible contributions and only provide major comments.

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.

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.

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.

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.

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.

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.