NBER WORKING PAPER SERIES

THE MINIMUM WAGE, TURNOVER, AND THE SHAPE OF THE WAGE DISTRIBUTION

Pierre R. Brochu
David A. Green
Thomas Lemieux
James H. Townsend

Working Paper 33479
http://www.nber.org/papers/w33479

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2025

We would like to thank Iain Snoddy, Daniel Gao and Bo-Yu Chen for expert research assistance, and SSHRC and the Bank of Canada Fellowship Award for financial support. We are grateful for comments from Attila Lindner, Denis Dutz, Magne Mogstad, Bruce Meyer, and seminar participants at the University of Chicago, University College London, the Vancouver School of Economics, and the University of Waterloo. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

The Minimum Wage, Turnover, and the Shape of the Wage Distribution
Pierre R. Brochu, David A. Green, Thomas Lemieux, and James H. Townsend
NBER Working Paper No. 33479
February 2025
JEL No. J31, J38

ABSTRACT

This paper proposes an empirical approach to decompose the distributional effects of minimum wages into effects for workers moving out of employment, workers moving into employment, and workers continuing in employment. We estimate the effects of the minimum wage on the hazard rate for wages, which provides a convenient way of re-scaling the wage distribution to control for possible employment effects. We find that minimum wage increases do not result in an abnormal concentration of Job Leavers below the new minimum wage, which is inconsistent with employment effects predicted by a neoclassical model. We also find that, for Job Stayers, the spike and spillover effects of the minimum wage are simply shifted right to the new minimum wage. Our findings are consistent with a model where entry wages are set according to a job ladder, and where firms anchor their internal wage structure on the minimum wage due to fairness or internal incentives issues.

Pierre R. Brochu
University of Ottawa
pbrochu@uottawa.ca

David A. Green
Vancouver School of Economics
University of British Columbia
6000 Iona Drive
Vancouver, BC V6T1L4
Canada
green@econ.ubc.ca

Thomas Lemieux
Vancouver School of Economics
University of British Columbia
6000 Iona Drive
Vancouver, BC V6T 1L4
and NBER
thomas.lemieux@ubc.ca

James H. Townsend
University of Winnipeg
j.townsend@uwinnipeg.ca

1 Introduction

Minimum wages alter the wage distribution through truncating it, inducing a spike at the minimum wage, and generating spillover effects on wages just above the minimum wage. This has now been well established in a set of recent, innovative papers (Cengiz et al. (2019), Gopalan et al. (2021), Jales (2017)). For the United States, where much of the evidence is focused, the spillovers reach to approximately \$2 to \$2.50 above the minimum wage. These effects are important, in part, because of their direct implications for inequality (e.g., Fortin et al. (2021)). Beginning with Meyer and Wise (1983) and more recently (and more completely) Cengiz et al. (2019) and Gopalan et al. (2021), the patterns have also been used as a way to study the employment impacts of minimum wages. But the changes in the shape of the wage distribution also have the potential to provide insight into the relevance of different labour market models. Indeed, both Cengiz et al. (2019) and Gopalan et al. (2021) end their discussions with general statements that their results showing the extent and location of spillovers could be used to discern among labour market models.

Previous research on the minimum wage has provided fruitful ground for understanding which models best describe the functioning of the low-wage labour market. While a standard neoclassical model predicts that higher minimum wages should reduce employment, Card and Krueger (1995) conclude that a model in which employers have monopsony power and set their own wages best accounts for the lack of minimum wage employment effects they document in their research (see also Manning (2003)). Likewise, the presence of a spike at the minimum wage is incompatible with a standard neoclassical model where workers are paid their marginal product, which is smoothly distributed over the workforce. Again, monopsony power –or other sources of labour market imperfections– can readily account for the existence of the minimum wage spike.(e.g., Dickens et al. (2012), Haanwinckel (2024) and Engbom and Moser (2022)). Continuing in this vein, our goal in this paper is to expand the set of data patterns that can be used for examining the implications of key labour market models. More specifically, using rich Canadian data consisting of a large set of short panels (and covering a period with over 150 changes in nominal minimum wages), we are able to decompose the effects of a change in the minimum wage on the wage distribution into elements related to job turnover and wage changes for continuing workers. We argue that these elements can be related directly to implications from the main models of the labour market and, thus, can be used to give more weight to some over others.

Our contribution stems from four novel features of the paper. The first is our use of a hazard based estimator presented in Donald et al. (2000) that allows us to examine the impact of the minimum wage on the shape of the wage distribution while controlling for potentially confounding factors such as education and age. Employing the hazard rates corresponding to a wage distribution (e.g., the probability that an observed wage is \$10 conditional on it being at least \$10) as the basis of our estimation has advantages in terms of flexibility, allowing us to tap into the rich specifications that have been developed in the hazard literature (e.g., Meyer (1990)). We employ a specification with a non-parametric estimator for the underlying shape of the hazard function and a very flexible mechanism for introducing covariate effects that allows them to have differing effects across the range of wages. [1]

Most importantly, though, using hazard rates improves our ability to discern the impact of minimum wages on the shape of the wage distribution. In particular, if a minimum wage increase leads to reduced employment then the wage density will have to be inflated by 1 minus the reduction in the probability of employment in order to make it continue to integrate to 1. The results of a density based estimator will then conflate this mechanical effect (that extends across the entire wage range) with the economic effects (selection and changes in firm wage setting policies) that we are interested in (a point raised by Flinn (2006) and Stewart (2012)). The hazard does not face this difficulty since the re-scaling factor affects both the numerator (the density) and the denominator (1 minus the CDF) and cancels out. In addition, the hazard estimator, by definition, focuses attention on movements in the CDF, which we will argue is particularly useful for assessing models in which the ranking of wages on a job relative to other wages is important.

The second novel feature of the paper is built on the flexibility of our hazard estimator, which allows us to control for completely flexible time by province effects on the shape of the wage distribution. We combine those effects with an assumption that minimum wages do not affect the wage distribution above some cut-off wage to define a triple difference estimator of the effect of minimum wages on the shape of the distribution below that cut-off. This novel identification approach identifies the difference between the shape of the distribution with and without a minimum wage (i.e., the total effect of the minimum wage). This is different from the difference-in-difference estimators used in the key Cengiz

[1]The estimator is also straightforward to implement since we do not need to impose integration constraints that would be necessary for consistency in directly estimating densities, and we provide a new, Generalized Linear Estimator approach. As long as the hazard rates are non-negative, we can integrate them to form consistent, associated density and CDF estimates.

et al. (2019) and Gopalan et al. (2021) papers, which identify changes in the height of the density at the points where it changes. Thus, for example, if a minimum wage reduces the mass in the density below, say, \$5 but a *change* in the minimum wage from \$10 to \$12 does not affect the extent of that reduction then the difference-in-difference estimator will record no minimum wage effect below \$5 while our estimator will show the reduction. We argue the latter, overall shape effects are useful for discerning between models of the labour market. The cost of our approach is the assumption that there is no effect beyond the cut-off wage, and we show that our results are unaffected by changes in that cut-off as long as it is set high enough. We compare our triple difference hazard estimator to the difference-in-difference estimator in a Monte Carlo exercise in section 4 of the paper.

The third feature is a decomposition of changes in the overall distribution into groups defined by job change status. In particular, we estimate our model separately for: Leavers (workers who separate from a job - estimating the distribution of their wages just before they separate); Stayers (workers who remain on the same job between two periods); and Joiners (new hires - estimating their wage distribution just after they are hired). The wage distribution for Leavers reveals from where in the overall wage distribution separators are selected. The change in that distribution if the separation coincides with a minimum wage increase then tells us whether those increases affect the overall wage distribution by altering selection. In a standard neoclassical model, for example, we would expect there to be considerable mass near the minimum wage as low productivity workers are laid off. In contrast, changes in the Stayers' wage distribution at the time of a minimum wage increase has no selection element since we are following the same workers over time. Instead, changes in their distribution at the time of a minimum wage increase will reveal any changes in the wage setting policies of firms linked to the increase. We argue that separating selection from wage policy effects in this way is useful for delineating among economic models.

Finally, we also examine the dynamics of the adjustments to a minimum wage change for the overall sample and all our subgroups. This allows us to determine whether firms and workers take time to fully adjust to a change.

We implement our estimator, first, for the overall wage distribution. Mirroring the results in papers such as Cengiz et al. (2019), Gopalan et al. (2021), and Jales (2017), our estimates show a substantial spike at the minimum wage and spillover effects of minimum wage increases reaching up to approximately \$2.50 above the minimum wage. As we discuss in our Monte Carlo exercise, we expect different estimation approaches to yield similar estimates near the minimum wage, though our estimates show different patterns elsewhere in the distribution and bear a different interpretation. For both the overall wage distribution and our component groups, we find that the wage structure responds very rapidly to a minimum wage change. We find no evidence of dynamics beyond about 2 months after a minimum wage increase.

When we implement our estimator for Job Leavers, we find no effect of upcoming minimum wage changes. The wage distribution of Leavers is different from other workers (with more mass at lower wages) but it does not change at the time of a minimum wage increase.[2] This result clearly does not fit well with neoclassical models in which workers with abilities below the new minimum wage are simply laid off. It also does not fit with versions of monopsony or bargaining models in which firms with productivity below the new minimum wage shut down. Thus, this result (plus the existence of a spike at the minimum wage) rejects standard neoclassical models and is potentially problematic for other models as well.

For Job Stayers, we find that the spike and spillover effects of a minimum wage are simply shifted right to the new minimum wage at the time of a minimum wage increase. That is, the effects of a minimum wage for continuing workers maintain a fixed pattern that is anchored on the minimum wage but is independent of its level. This finding is difficult to reconcile with most models typically invoked to understand minimum wage effects. Monopsony and search and bargaining models predict that workers whose wages were between the old and new minimum wage should all be moved to the new minimum wage. This generates a spike but, typically, no spillover effects among Job Stayers. Burdett and Mortensen (BM) type job ladder models provide a natural explanation for spillover effects as firm set wages to maintain their relative position in the wage distribution (Burdett and Mortensen (1998)) but do not generate a spike at the minimum wage. The findings for Job Stayers are instead consistent with a class of models based on internal incentives (promotions or tournaments), efficiency wages, or fairness issues where it is important for firms to preserve their internal wage structure. Our results imply that the lower part of the internal wage structure is anchored on the minimum wage. These models, however, do not fit well with our finding of spillovers above the minimum wage for Joiners, since the internal considerations would not have emerged for them. Thus, on balance, we conclude that the evidence is consistent with a more general model where entry wages are set according to a job ladder, while wage growth for Job Stayers is driven by internal considerations. We offer these conclusions as an indication

[2] This potentially fits with results implying that minimum wages have small impacts on employment (e.g., Cengiz et al. (2019), Brochu and Green (2013)).

of how our results can be used for assessing models of the labour market, though our main contribution is to present the various patterns so that other researchers can use them in their own assessments and model building.

Our paper contributes to the growing literature on the impact of minimum wages on the wage distribution. The literature starts with Grossman (1983) and Meyer and Wise (1983), with important early contributions in DiNardo et al. (1996) and Card and Krueger (1995). These papers use US data and much of the existing literature is focused on either the United States or the United Kingdom.[3] For Canada, the only existing paper of which we are aware is Fortin and Lemieux (2015), who show important effects of minimum wage increases in reducing inequality. There is some disagreement in the literature on the extent of spillovers onto above-minimum-wage wages, ranging from large spillovers (Lee (1999)) to quite small ones (Autor et al. (2016)), though our reading of the most credible recent evidence for the United States is for spillovers up to about $2 above the minimum wage. We find similar results for Canada. To reiterate our point from earlier, much of the existing literature estimates the impact of minimum wage changes on different quantiles of the wage distribution. Estimates of effects on quantiles are useful for studying overall inequality effects but offer only a rough picture of the impact on the shape of the wage distribution. Our hazard based estimator is better suited to the latter task.

Both Cengiz et al. (2019) and Gopalan et al. (2021) share similarities with our paper in broad methodology, though not in focus. Both focus on the impacts of minimum wage changes on changes in employment at different points of the wage distribution, finding reductions below a new minimum wage that are exactly offset by increased employment at and above the minimum wage. Gopalan et al. (2021) show that layoff and quit rate effects are zero in all wage bins, which fits with our first result that the CDF of wages for Job Leavers does not change with a minimum wage increase. Neither paper provides the equivalent of our Stayers results. Our main contribution is to show the effects of minimum wages on the overall shape of the wage distribution and to use that to differentiate between different models of the labour market.

The paper proceeds as follows. We start by explaining, in Section 2, how the overall

[3]For instance, see Lee (1999), Teulings (2000), Neumark et al. (2004), Flinn (2006), Autor et al. (2016), Cengiz et al. (2019)) for the United States and Manning (2003), Dickens and Manning (2004a), Dickens and Manning (2004b), Dickens et al. (2012), Stewart (2012) for the United Kingdom. There is also an emerging literature examining Brazil, where the minimum wage appears to have played a role in substantially reducing inequality (Jales (2017), Engbom and Moser (2022), Alvarez et al. (2018), Haanwinckel (2024)).

wage distribution is connected to those for Leavers, Joiners, and Stayers. In Section 3, we describe our estimator and its advantages. In Section 4, we present the Monte Carlo exercise comparing estimators. In Section 5, we describe our data and the minimum wage variation we use. In Section 6, we discuss implementation details for our estimator and present the results.We then discuss the implications of our key findings for models of the labour market in Section 7. Section 8 contains conclusions.

2 Decomposing Changes in the Wage Distribution

We begin with a decomposition to show how the overall wage distribution depends on the wage distribution for people who separate from a job ("Leavers"), people who are hired into a job ("Joiners") and people who remain in the same job ("Stayers"). We will argue that different classes of models mostly differ in terms of the implications for the Leavers and Stayers distributions.

We start with a benchmark case where the value of the minimum wage remains constant at a level, m_0, over time. Consider a definitional expression for the CDF of wages in a period, t, as a weighted average of the CDF for Leavers (who are about to leave their current job) and Stayers (people who will still be in the same job in period t+1):

$$F_t(y|m_0) = \frac{N_S}{E_t} \cdot F_{S,t}(y|m_0) + \frac{N_{L,t}}{E_t} \cdot F_{L,t}(y|m_0), \tag{1}$$

where $F_t(y|m_0)$ is the CDF of wages for all workers in period t when the minimum wage equals m_0; $F_{S,t}(y|m_0)$ is the CDF of wages for Stayers; $F_{L,t}(y|m_0)$ is the CDF of wages for Leavers; N_S is the number of Stayers, which does not have a time subscript since it is the same in the two periods; $N_{L,t}$ is the number of Leavers; and E_t is the total number of workers in period t ($E_t = N_S + N_{L,t}$). Notice that the CDFs are all allowed to vary over time while holding the minimum wage constant, allowing for impacts from other policy changes and macro conditions. Note, also, that there is no restriction on why a person leaves a job or where they go after leaving it. They could be laid off, quit, or have their firm shut down, and they could move immediately into a new job or be non-employed in t+1.

A similar equation can be formed for period t+1, but focusing on job entrants:

$$F_{t+1}(y|m_0) = \frac{N_S}{E_{t+1}} \cdot F_{S,t+1}(y|m_0) + \frac{N_{J,t+1}}{E_{t+1}} \cdot F_{J,t+1}(y|m_0), \tag{2}$$

where $N_{J,t+1}$ is the number of Joiners in t+1 who were either in a different job or were not working in period t, and $F_{J,t+1}(y|m_0)$ is their CDF of wages. As with the Leavers, there are no restrictions on where the Joiners came from.

Combining equations (1) and (2) highlights three channels through which a change in the CDF relative to the previous period can occur:

$$F_{t+1}(y|m_0) - F_t(y|m_0) = \frac{N_{J,t+1}}{E_{t+1}} \cdot [F_{J,t+1}(y|m_0) - F_{L,t}(y|m_0)] \tag{3}$$

$$+\frac{N_S}{E_{t+1}} \cdot \Delta F_{S,t}(y|m_0)$$

$$+\left[\frac{N_S}{E_{t+1}} - \frac{N_S}{E_t}\right] \cdot F_{S,t}(y|m_0) + \left[\frac{N_{J,t+1}}{E_{t+1}} - \frac{N_{L,t}}{E_{t+1}}\right] \cdot F_{L,t}(y|m_0)$$

where $\Delta F_{S,t}(y|m_0) = F_{S,t+1}(y|m_0) - F_{S,t}(y|m_0)$. The first is differences between the CDFs for Leavers in period t and Joiners in period t+1 (the first line). This is a standard selection effect, capturing forces such as a tendency for firms to lay off less skilled workers and replace them with more skilled workers. The second channel is changes in the shape of the wage distribution for Stayers (the second line). This would arise either from policy changes other than the minimum wage, from changes in firm wage practices, or from macroeconomic changes. The third line captures compositional effects linked to changes in employment and worker flows in and out of employment.

Our main interest, of course, is in the impacts of changes in the minimum wage, and those changes can impact the wage CDF through any or all of the three components of equation (3). Previous papers mainly focus on the overall change in the CDF, showing that it involves both shifts in the location of a spike at the minimum wage and spillover effects above the minimum wage. But we see equation (3) as revealing that relating those overall shifts to economic forces is complicated. Policy and economic changes can operate through any of the three channels in ways that could be either reinforcing or offsetting.

As we will discuss in Section 7, the implications of minimum wage changes in different models are most easily seen in their impacts through the selection and shape channels. To understand how we estimate effects through those channels, it is helpful to see how we define the various types of spells. In figure 1, we show the workers in May of a year in a province where the minimum wage is initially \$10. We define Leavers as workers in March who are either not employed or on a different job in May.[4] Suppose (as we

[4] We use job status two months later rather than one month later because of the timing of some minimum wage changes, as described in the data section.

depict in the figure) that the minimum wage in the province increases to \$11 in April. We would then expect that some of those who separate from the job they held in March did so because of the minimum wage increase. We could, then, estimate the minimum wage effect on the wage distribution of Leavers as:

$$\Delta F^{k}_{L,t}(y|m_0 \to m_1) - \Delta F^{j}_{L,t}(y|m_0) \tag{4}$$

where, $F^{k}_{L,t}(y|m_0 \to m_1)$ is the CDF for job Leavers in province, k, where the minimum wage increases between t and t+1 (March and May, in our example). Our estimator actually allows for substantial flexibility in differences in general shifts in the wage distributions between the provinces by adopting a triple difference specification that further compares changes in a lower part of the distribution to changes in an upper part (assumed not to be affected by the minimum wage) within each province. The CDF for Leavers represents the wages of workers just before they separate from a firm and, hence, the selectivity of separations in terms of where they come from in the wage distribution (not surprisingly, this distribution tends to have more mass at lower wages compared to $F_{S,t}(y|m)$). Our estimated minimum wage effects then reveal whether that selectivity differs when the separations coincide with a minimum wage increase.[5] One might expect, for example, that workers earning the old minimum wage, m_0, would be laid off because their productivity might not justify being paid the new minimum wage.

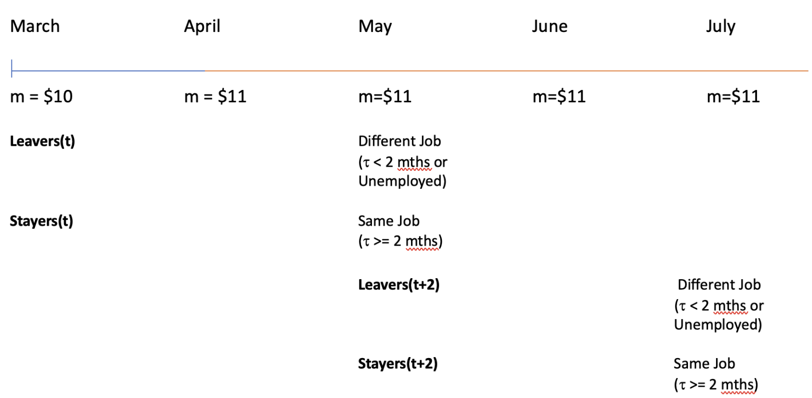

Figure 1: Definitions of Types of Job Spells

The remaining workers in March, in our example, are Stayers - workers who will still be on the same job in May. The change in the CDF between May and March for the

[5]This difference corresponds solely to selection as long as the firm does not change its payment policies in period t, in anticipation of the minimum wage change. In Appendix A, we show that this is true for Stayers in period t and take that as evidence that it is also unlikely for the workers about to separate.

Stayers (relative to the change for Stayers in provinces where there is no minimum wage change in April and changes in the upper part of the distribution) captures the effect of the minimum wage change on how firms pay their workers. Because it is the same set of workers through time, there is no selection effect by definition. Thus, we are able to separate selection and shape effects by examining the CDFs for Leavers and Stayers separately. The third type of workers are new job Joiners. Movements in their CDFs will reflect both selection (if firms decide to hire different workers when the minimum wage is higher) and shape effects (if firms change how they pay various workers in response to the minimum wage increase.) We present results for the Joiners but don't focus on them because they don't provide a clean representation of either channel.

3 Hazard Based Estimator

Our goal is to estimate the impact of the minimum wage on the hourly wage distribution. We want an estimator that allows for flexible conditioning of the distribution on covariates in order to control for potential differences in characteristics of the workforce such as education level across places and time and, most importantly, to allow for a rich conditioning on time by location effects. For that purpose, we use the hazard based estimator set out in Donald et al. (2000) to estimate $F(y|x)$: the conditional (on employment) cumulative distribution of wages, y, given a vector of covariates, x.

Defining $f(y|x)$ as the conditional density of wages, we can write

$$f(y|x) = h(y|x)S(y|x), \tag{5}$$

where $h(y|x)$ is the conditional hazard function (the probability that the wage equals y conditional on the wage being at least as large as y) and $S(y|x)$ is the survivor function (the probability that wages are at least as large as y). It is well known that,

$$S(y|x) = exp(-\int_{y_0}^{y} h(u|x)du) \tag{6}$$

where y_0 is the minimum level of wages defined in the data. In the Donald et al. (2000) approach, the hazard function, $h(y|x)$, is estimated directly and then, based on equation (6), one can retrieve estimates of $F(y|x) = 1 - S(y|x)$. This approach allows us to make use of hazard function estimators developed in the duration model literature that: i) permit covariates to be introduced easily; and ii) are flexible, in the sense that they impose a minimum of restrictions on the shape of the hazard function for any value of the covariate vector. In contrast to standard quantile estimators, the resulting estimates of $F(y|x)$ are also guaranteed to be consistent in the sense that the distribution function always lies between 0 and 1.

One key advantage of focusing on the hazard rate is that it provides a convenient way of re-scaling the observed wage density to control for possible employment effects of the minimum wage. That is, it eliminates the effects of the third (mechanical) channel, allowing us to concentrate on the effects of the selection and shape channels. To see this point, consider an example in which the minimum wage simply truncates the density without otherwise changing its shape. In this case, if the true wage density is given by $f(y)$ then the observed wage density is given by $f^*(y) = \frac{f(y)}{1-F(m)}$, where $F(.)$ is the CDF for the wage and m is the minimum wage. The denominator in this expression is needed to ensure that the observed density integrates to 1, and it implies an inflating of the underlying density that will increase as m rises and the employment rate declines. Stewart (2012) calls this effect a truncation spillover, distinguishing it from a shape spillover in which the actual shape of the density changes above m. As we will see, the various models we are interested in have differing implications for shape spillovers that we would like to distinguish from truncation spillovers. Unlike the density, the observed hazard rate at y given by $h^*(y) = \frac{f^*(y)}{1-F^*(y)}$ is unaffected by the truncation of the wage distribution since the numerator, $\frac{f(y)}{1-F(m)}$, and the denominator, $1 - F^*(y) = \frac{1-F(y)}{1-F(m)}$ (the observed survivor function), are both divided by $1 - F(m)$. Thus, in our example in which minimum wages only serve to truncate the wage distribution, the hazard rate would remain unaffected. The same logic holds in our more general situation described in the previous section, as we show in Appendix B.[6]

We adopt a proportional hazard specification in which:

$$h(y|x) = exp(q(x, y))h_0(y) \tag{7}$$

where $h_0(y)$ is the baseline hazard that is common for everyone regardless of their values of x and $exp(q(x, y))$ is a proportionality factor that alters the height of the hazard in a manner that depends on the covariates, x, and, potentially, with the value of y. Following Meyer (1990), we adopt a flexible form for the baseline hazard, implementing it as a step

[6]Of course, this result is not relevant if the minimum wage has no impact on the employment rate. Given that there is still some debate on the existence and size of employment effects, we believe it is better to use an estimator that is robust to their existence.

function in y. To do this, we first divide the wage range into P bins defined by thresholds:

$$y_0 < y_1 < ... < y_P. \tag{8}$$

All observations with $y > y_P$, including top-coded observations, are part of an upper-wage bin that is treated as right-censoring in the estimation. We then introduce a parameter, γ_p, corresponding to the height of the hazard in each bin.[7] We use a large number of baseline bins (P = 164), implying substantial flexibility in the baseline hazard function.

We specify the proportionality factor as:

$$q(x, y) = exp(x\alpha_{s(p(y))})exp(D_{mp(y)}\beta) \tag{9}$$

where, $p(y)$ is the baseline bin in which y falls. The first term in $q(x, y)$, $exp(x\alpha_{s(p(y))})$, captures the effects of a vector of covariates consisting of personal characteristics and time effects, described in greater detail below. Importantly, the effects of the covariates are allowed to vary across sets of the baseline bins that we call 'covariate segments' indexed by $s(p)$. This approach allows for the possibility of completely different shapes for the hazard functions for different values of the covariate vector. For example, it would allow for a bimodal wage distribution for the young, a left skewed distribution for the middle aged, and a right skewed distribution for older workers.

We have investigated a variety of specifications for province and year effect elements of the x vector. In our preferred specification, we treat the first covariate segment differently from the other four segments. Although all of the minimum wages lie in the first segment, spillovers effects range up to the fourth segment in some specifications. In the first segment, we include a complete set of province by year effects. In the other segments where most spillover effects take place, we include separate sets of province effects, year effects, and a quadratic trend in time in order to control for the large differences in real wage growth across provinces over time (Fortin and Lemieux (2015), Green et al. (2019)).[8] Recall that the effects of these variables are allowed to vary by covariate segment. To further our approximation, we also include interactions between the average wage in a

[7]Formally, with a continuous underlying baseline hazard, γ_p is the integration of the baseline hazard over the interval (y_{p-1}, y_p) (Meyer (1990)). Using this flexible parametrization, equation (7) becomes $h(y|x) = exp(q(x, y))\gamma_p$.

[8]Our estimates change little when we include the complete set of province by year effects in the second covariate segment as well as the first. Allowing for province specific time trends is sufficient to capture the important variation within province over time, with results changing little when we replace those trends with the complete set of province by year interactions in the first or first and second covariate segments.

bin and each of the full set of province and year effects. Because the bins are uniform in width (10 cents), this amounts to an interaction of the bin number with the province and year effects. We include these interactions in all the covariate segments, including the first, and allow their associated coefficients to differ by covariate segment. Thus, we employ a specification with considerable flexibility, with the shape of the distribution being allowed to vary by province and year but also, to a substantial extent, by province-year cells. This is important because Stewart (2012) showed that using high percentiles (e.g., the 50th percentile) as controls in an identification strategy for spillover effects of this form could be problematic because the higher percentiles appear to be following different trends from the lower percentiles in his UK data. Cengiz et al. (2019), similarly, show that employment shocks can affect different parts of the wage distribution differently and in ways that confound estimation of minimum wage effects in their US data. Our specification addresses these issues by allowing for different trends in different parts of the wage range corresponding to the covariate segments. Moreover, as we will see, our identification is based just on movements in the wage range directly above the top of our minimum wage effect range ($5 above the minimum wage in our main specification).

The second term in $q(x, y)$, $exp(D_{mp(y)}\beta)$, captures the effects of the minimum wage. Within the context of this estimator, minimum wages are analogous to time varying covariates in the standard duration literature. For expositional purposes, we break the minimum wage terms into two subcomponents and leave aside the wage y from the subscript $mp(y)$:

$$exp(D_{mp}\beta) = exp(D^b_{mp}\beta^b + D^c_{mp}\beta^c), \tag{10}$$

where, D^b_{mp} is a matrix of dummy variables corresponding to ranges that are anchored on the minimum wage in force in the given province and quarter (which is similar in nature to the approach taken in Cengiz et al. (2019)); D^c_{mp} is a matrix containing similar wage range dummy variables interacted with dummies corresponding to different time lags; and β^b and β^c are associated vectors of coefficients. $D^b_{mp}\beta^b$ correspond to long term minimum wage effects that are in place when there are no recent changes in the minimum wage, which we refer to as the base effects, and $D^c_{mp}\beta^c$ captures short and medium term dynamics. We will begin with a description of content and identification for the base effects then return to a discussion of the dynamic effects in the next subsection.

The columns of the D^b_{mp} matrix correspond to dummy variables that each equal 1 in a different range of wages as follows: 50 cents or more below the minimum wage (m); 30 to 49 cents below m; 10 to 29 cents below m; 10 to 29 cents above m; 30 cents to 49 cents above m; and 50 cents bands above that until reaching $5.00 above m. There is also a dummy representing the minimum wage itself which actually corresponds to the minimum wage plus or minus 10 cents to allow for some measurement error in reporting. The estimated coefficient on the minimum wage variable corresponds to how much the baseline gets shifted with the presence of the minimum wage and provides an estimate of the "spike" in the hazard (and, in consequence, the density) at the minimum wage. The coefficients on the other variables in D_{mp} have a similar interpretations and, using those estimates, we can map out the effect of the minimum wage on the shape of the wage distribution. In this specification, minimum wages have the same proportional effect on the hazard at any point in the distribution, but we also present specifications in which the $D_{mp(y)}$ vector is interacted with the value of the minimum wage itself to allow for the possibility that minimum wage effects could be, for example, smaller when the minimum wage is smaller. We estimate the models separately for males and females, allowing minimum wage effects to vary by gender.

Given the combination of the assumption that minimum wages do not affect the wage distribution beyond a specific threshold ($5 above the minimum wage in our main specification) and the complete set of province-by-year dummies in the first covariate segment, our most general specification is similar to a conventional "triple-differences" design. We demonstrate the identification strategy in Figure 2 where we plot imaginary estimates in a specification with baseline wage intervals that are 25 cents wide, a flat baseline hazard, the minimum wage allowed to have effects on the hazard rate up to 2 dollars above the minimum wage, and all of the wage intervals shown in the figure assumed to be within the first covariate segment. Thus, there is a common year by province effect helping to determine the height of hazard rates in all the wage intervals to a common degree. We plot the hazard rate in each wage interval for the cases of an $8 minimum wage (dark bars) and a $9 minimum wage (lighter bars) under an assumed pattern in which the minimum wage induces a spike (considerable extra mass) in the wage interval that contains the minimum wage and a spillover pattern with declining impacts as we move further above the minimum wage.

In this simple example, the province-by-year baseline hazards can be identified in the "control" zone above the upper limit of the spillover effects.[9] The spillover effects can then be estimated by comparing the difference in hazards slightly above the minimum wage to those in the "control" region to the right. Consider, for example, the $10 to

[9] A complete triple difference specification in our context would include a complete set of wage bin by province effects and wage bin by year effects. Because we use a large number of wage bins (164), including all of these effects is computationally infeasible. Our flexible province by year by covariate segment effects approach is an approximation to this.

$11 wage range where there are some spillover effects in the province-year combination with the $9 minimum wage (spillover between $1 and $2 above the minimum wage), but no spillover effects in the province-year combination with the $8 minimum wage. The spillover effects in the $1 to $2 above the minimum wage range can then be estimated by comparing the difference in the hazard in the $10 to $11 range when the minimum wage is $9 to the hazard in the same range when the minimum wage is $8. Having established that spillover effect, we can use that to establish the height of the baseline hazard between $9 and $10 when the minimum wage is $8 and use the comparison of that estimated baseline to the height of the hazard in the same wage range when the minimum wage is $9 to establish the spillover effect of the minimum wage in the 0$ to 1$ above the minimum wage range. The same arguments continue down the wage range, resulting in an identification of the effect of minimum wages down to the lowest wage through a cascading argument. In essence, our approach enables us to fully control for province-year effects in the baseline level of the hazard while exploiting shape differences to estimate the spike and spillover effects.[10] Importantly, what we are identifying is the impact of minimum wages measured relative to the baseline hazard. We return to this point in the next section.

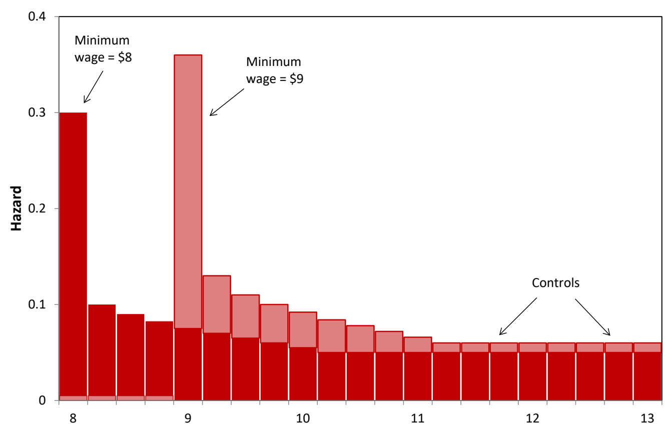

Figure 2: Triple-Differences Identification Strategy

Several points follow from this design. First, because of the recursive nature of the identification, the parts of the minimum wage effects that are nearest to the upper cut-off tend to be more precisely estimated. Second, we can vary the range over which the spillovers are allowed to occur. In Appendix A, we compare results using a $5 and $7

[10]Note that our use of the hazard is important in this identification approach. If truncation spillover effects are included, as they would be if we worked with density functions, there would be no region of the wage range that is unaffected by the minimum wage since the truncation spillover effect inflates the entire density.

cut-off. We show that the estimated effects between \$5 and \$7 above m are very close to zero and statistically insignificant when we use the \$7 cut-off. Moreover, the estimated spike and spillover effects are extremely similar using the two cut-offs. We conclude from this that our identifying assumption that the minimum wage has no effect above \$5 above m is likely to hold.[11] We use the \$5 rather than \$7 above m in our main results because, from the first point, this allows for more precise estimates of minimum wage effects in the region where they seem to actually arise.

3.1 Dynamic Responses to Minimum Wage Changes for Stayers, Joiners, and Leavers

We turn, now, to describing how we implement our estimator for our different subgroups (Leavers, Joiners, and Stayers) and how we introduce lagged effect dynamics. We capture those dynamics using the dummy variables that make up the columns of the D^c_{mp} matrix. In particular, in the case of Leavers, we are interested in whether the wage CDF for Leavers who separate from their job at the time of a minimum wage increase (called $F_{L,t}(y|m_0 \to m_1)$, above) is different from the CDF of Leavers at other times ($F_{L,t}(y|m_0)$). To capture this, we create three groups of wage-range variables. The first group consists of dummy variables indicating if the wage bin, p, is below the old minimum wage (m_0)): 50 cents or more below m_0, 30 to 50 cents below, and 10 to 30 cents below. The second group of dummies allows for a flexible impact in the range going from the old to the new minimum wage. Three dummy variables are used for indicating if a given wage bin is at m_0, at m_1, or in the range between m_0 and m_1. The final group of dummies captures wage ranges that are above the new value of the minimum wage - the one that will be in place in month t+2. These are the same five dummies going up to \$5 above the new minimum wage that are used to capture spillover effects in D^c_{mp}.

In our main specification, we enter each of these wage range dummies in interaction with each of three event time dummies: one, K_{at}, which equals 1 if the separation happens at the time of a minimum wage change; a second, $K_{1-2,t}$, which equals 1 if there was a minimum wage change 1 to 2 months before the separation; and a third, $K_{3-12,t}$, which equals 1 if there was a minimum wage change 3 to 12 months before the separation. In

[11] We cannot, of course, prove that there is not some wage range higher up where the minimum wage has an effect (even if we set the cut-off at \$50 above m, that doesn't prove that there aren't effects still higher), but we view these results as highly suggestive. It is worth highlighting that our estimator includes very flexible time and location effects that vary by wage ranges and, thus, we believe our estimated minimum wage effects are unlikely to include spurious associations with other events.

terms of the example in figure 1,: K_{at} equals 1 for workers on a job in t (March) who are not on that job in t+2 (May) (i.e., period t Leavers) when there is a minimum wage change between March and May; $K_{1-2,t}$ equals 1 for job Leavers in t+2 (May) when there was a minimum wage change between March and May; and $K_{3-12,t}$ equals 1 for job Leavers in April through the following February when there was a minimum wage change between March and May. The interaction of the wage range dummies with K_{at} allows for a very flexible impact of an upcoming minimum wage change on the Leavers' wage distribution. The associated elements of the β^c vector capture a change in the CDF of wages on the previous job for those observed to separate at the time of a minimum wage change relative to the base CDF (i.e., the pre-separation CDF for job Leavers in periods with no upcoming or recent minimum wage changes). If firms lay off some of the workers earning wages below the new minimum wage in response to a minimum wage increase, we should see a sharp increase in the hazard in this wage range. For example, the coefficient on the dummy indicating wages in between the old and new minimum wages should be positive and significant.

The interactions of the wage range variables with $K_{1-2,t}$ and $K_{3-12,t}$ capture the possibility that firms take time to respond to a minimum wage increase, gradually laying off more workers. Notice that we interact these time lag dummies with wage range variables defined at time t, which allows for the possibility that firms take time to bring their wages in line with the new m value. Given these interactions, the coefficients on the D^b_{mp} capture the "steady-state" effects of the minimum wage in periods where the minimum wage has not changed for at least a year, and all the elements of D^c_{mp} are equal to zero. We have implemented specifications in which we allow for further lags in minimum wage effects but found no significant effects at lags beyond 1 year.

We utilize a similar specification for Joiners - designed to capture short term impacts of minimum wage changes on who firms hire. In this case, the K_{at} dummy equals 1 in periods when there has been a minimum wage increase in the previous 2 months (in our example, new hires in May when there was a minimum wage change between March and April). The D^b_{mp} matrix continues to be defined in terms of the previous minimum wage for consistency. $K_{1-2,t}$ and $K_{3-12,t}$ are defined in the same way as for the Leavers. This allows for scenarios such as workers being laid off by low productivity firms at the time of a wage increase being initially hired at the new minimum wage by higher productivity firms, and then the labour market returning to a steady state equilibrium. Such a scenario would show an abnormal concentration of new hires around the minimum wage just after a minimum wage increase that would gradually dissipate.

The specification for Stayers is complicated by the nature of our data. With true panel data recording wages in each month, we would follow Job Stayers and see how their CDF changed at the time of a minimum wage change. As we describe below, the LFS does not re-ask the wage on a job after the first month in the survey for Job Stayers and, so, it cannot be used in this way. Instead, we estimate the impact of a minimum wage change on the CDF for Stayers using statistical arguments. If we think about a minimum wage increase happening between months t and t+2, we can define what we call "Forward" and "Backward" Stayers. Forward Stayers are workers who enter the LFS in month t and who have job tenure of at least 2 months when they are observed in month t+2. Backward Stayers are workers who enter the LFS in month t+2 and who have job tenure of at least 2 months at that time. Importantly, the Backward Stayers at time t+2 are statistically identical (drawn from the same population of Stayers) to the Forward Stayers at time t. That means that we can estimate the Stayers' minimum wage effect using the Backward Stayers alone as long as there were no anticipatory effects altering the Stayers' CDF just before a minimum wage change. We can check for anticipatory effect by estimating the same specification we used for Leavers with the Forward Stayers sample. In Appendix A, we present results from this exercise and show that there is little evidence of such anticipatory effects.[12] Given that, we focus on the Backward Stayers estimates in the main body of the paper, interpreting the β^b estimates as the steady state effects of minimum wages and the β^c estimates as the immediate impacts of a minimum wage increase for a set of workers who do not change jobs as a result of the minimum wage change. The interactions with $K_{1-2,t}$ and $K_{3-12,t}$ then allow for the possibility that minimum wage impacts on Stayer's wages evolve more gradually with specifications allowing for even longer adjustments again showing no longer term effects.

4 Comparisons Across Estimators

Our estimator adds to two other main estimators that have been used to examine the impact of minimum wages on the wage distribution: a quantile based estimator (e.g., Lee (1999) and Autor et al. (2016)); and the estimator introduced in Cengiz et al. (2019). We view each of these estimators as primarily suited to a different question: the impact of minimum wages on overall wage inequality (the quantile estimator); the impact on

[12]The only exception to this statement is that the estimate of β^c at the upcoming minimum wage is statistically significantly positive. In other words, it appears that firms start moving up some workers to this new minimum wage shortly before it comes into effect.

employment of different types of workers defined by their wage (the Cengiz et al. (2019) estimator); and the impact on the shape of the wage distribution (our hazard estimator). In this section, we present a Monte Carlo exercise to compare the estimators and to highlight what each identifies.

In our Monte Carlo exercise, we generate 1 million observations from a log normal distribution with a mean of 3 and a standard deviation of 1 (which delivers a distribution that approximately matches the location and shape of our actual wage data in the middle of our sample period if we interpret the units as dollars). We refer to this as the underlying wage distribution. We start by assuming a minimum wage of m = \$12 and impose minimum wage effects in a series of scenarios defined by what we do with the wages below \$12: Scenario 1, drop observations at wages below m with probability p_0 (keeping them with probability $(1 - p_0)$) - generating a simple truncation with no spike or spillover; Scenario 2, either drop below m observations with probability p_0 or assign them a new wage value as a draw from a negative exponential distribution anchored at m with probability p_1 (keeping below m wage values at the same value with probability $(1 - p_0 - p_1)$); Scenario 3, the same as Scenario 2 but with $p_0 = 0$ (no disemployment effects). The negative exponential parameters in Scenarios 2 and 3 are chosen to generate a spike at the minimum wage and a declining spillover with its main mass reaching to about \$4 above m. Scenario 3 is intended to reflect empirical findings that the lay-off effects of minimum wages are minimal. We repeat these scenarios with a variety of values of m and implement each of the three estimators with each generated dataset. Here, we present results with a few scenario/minimum wage combinations chosen to highlight the advantages and disadvantages of each estimator.

The first estimator group involves examining the impact of minimum wage changes on moments of the conditional (on working) wage distribution. One version of this uses quantile regressions to examine the impact at various percentiles of the distribution and, so, we will call it the quantile regression approach. The strength of this approach is in answering the first question- what do minimum wages do to wage inequality - since, of course, it provides direct estimates of exactly that. On the other hand, it does not provide any evidence on the second - employment effects - question. It is also not well suited to the third (shape impact) question because of the re-scaling issue described earlier. In Appendix C, we present results from implementing the quantile estimator using Scenario 1 data and show that the re-scaling problem makes it appear as if there are spillover effects when none exist in the underlying data.

4.1 Cengiz et al. (2019) Estimator

The second set of estimators estimate changes in employment all along the wage range at the time of a minimum wage change. This is an approach that began with Meyer and Wise (1983). We focus on the most salient version which is found in Cengiz et al. (2019).

Cengiz et al. (2019) divide the wage range into bins, in the same way we do, and form a dependent variable as the number of employees in a bin, j, in state, s, and time, t, (E_{jst}) divided by the population of s in t (N_{st}). The division by N_{st} is done to provide comparability across states of different sizes. Because it is not total employment, what results is not a density. This allows them to examine the employment impacts of a minimum wage change along the wage range without the distraction of the re-scaling we described earlier. Thus, the estimator is ideally constructed to answer questions about shifts in employment related to a minimum wage change.

We examine the Cengiz et al. (2019) estimator using our second Monte Carlo scenario with p_0=0.5 and p_1=0.3. We present the density, proportion employed, and re-scaled density for a \$12 minimum wage under this scenario in the top left panel in figure 3. The figure shows the spike and spillover induced by the minimum wage in this scenario. We implement the Cengiz et al. (2019) estimator with two periods: the first with m=\$12 and the second with m=\$14. With effectively only one 'state' and just two periods, we implement a simple version of the estimator with the following regression:

$$\frac{E_{jt}}{N_t} = \beta_0 + \sum_{k=-4}^{k=217} \alpha_k I_{jt}^k + \mu_j + \gamma_t + u_{jt} \tag{11}$$

where: we define 10 cent wide 'bins' spanning the range of wages indexed by j; E_{jt} is the number of workers in bin j in period t (with 2 periods in our simple case); N_t is the population (which we treat as fixed); I_{jt}^k is a set of dummy variables centred on the minimum wage covering minimum wage effects at, above and below the minimum wage; and μ_j and γ_t are a complete set of bin and time fixed effects, respectively. Following Cengiz et al. (2019), the I_{jt}^k dummies each cover a \$1 range, with the ranges extending from \$4 below the minimum wage to \$17 above the minimum wage. The α_k's are the coefficients corresponding to each of these dummies.

We plot the estimated α_k's along with the actual changes in employment in the wage ranges from the Monte Carlo data in the bottom left panel of figure 3. The Cengiz et al. (2019) estimator effectively captures the patterns in the data. Notice that in both the data and the estimates there are no changes in employment below \$2 below the

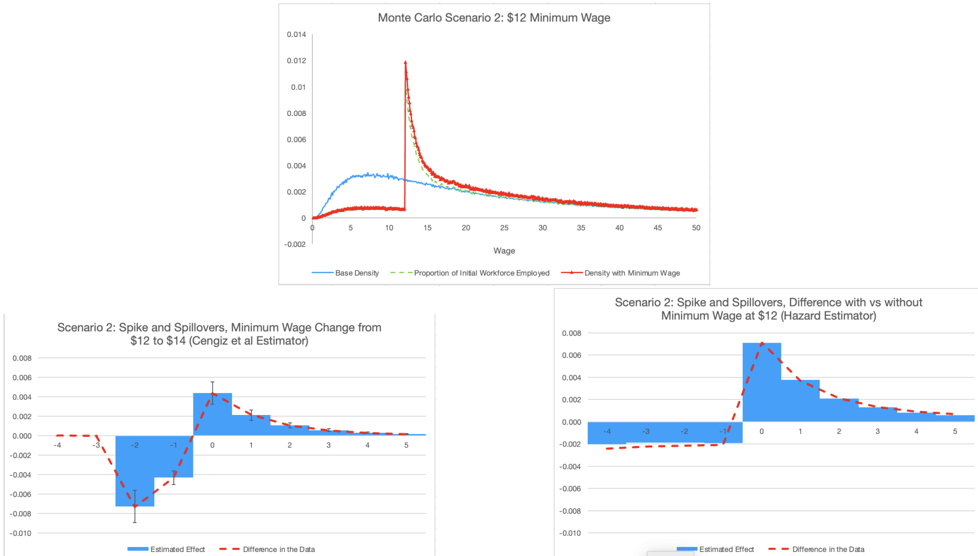

Figure 3: Scenario 2, Monte Carlo Densities and a Comparison between the Cengiz et al. and Hazard Estimators

new minimum wage, i.e., below the old minimum wage. This is because nothing has changed with the change in m in the wage range - the remaining number of workers is still $(1 - p_0 - p_1)$. In contrast, there is a large negative effect at $2 below m because that is where the minimum wage spike was under the old minimum and, so, the change in m implies a large decline in observations at that point. The spike at 0$ and the positive changes in employment for the $3 above it capture the spike and spillover effects from the minimum wage.

4.2 Hazard Estimator

We next implement our hazard estimator using the same data as we used for the Cengiz et al. (2019) exercise. We implement the hazard estimator using the set of bins described in Section 3 but present the results here in the $1 wide bins used by Cengiz et al. (2019) for comparability. We utilize an assumption that the minimum wage has no effect on wages above $7 above m (which is true in our data scenario) and do not include the time lag dummies since there are no longer time spans in our generated data. We plot the implied estimated effects on employment in each bin divided by the total population in the lower right panel in figure 3. The plotted effects are similar to those from the Cengiz et al. (2019) estimator in the bottom left panel in some ways (notably, in the existence of the spike and spillover patterns) but quite different in others. In particular, the hazard estimates show a constant negative effect on employment right down to \$0 while the Cengiz et al. (2019) estimates reveal no effects below \$2 below m, a large negative spike at \$2 below m and a smaller negative effect at \$1 below m. The height of the spike and spillover effects are also much smaller in the Cengiz et al. (2019) estimates and the declines just below m much larger than in the hazard estimates.

The differences between the estimates from the two estimators arise from what each is identifying and, along with that, differences in the identifying variation being used. The Cengiz et al. (2019) estimator is a difference-in-differences estimator that identifies the change in employment in each bin that is due to the change in m. The large negative effect at \$2 below m from that estimator exists because that is the location of m prior to the change and the estimate captures the difference between the spike at the old m and roughly zero employment at that location after the change. In contrast, the hazard estimator identifies the difference between the hazard (and with it the density) in a situation with a minimum wage of a certain level versus a situation with no minimum wage, and that difference is smaller and constant below m. On the other side, the hazard estimated difference is large at the spike and spillover points while in the Cengiz et al. (2019) estimator the difference between the spike at the new m and the employment level at that point before the minimum wage change (which includes spillover effects from the prior minimum wage) is smaller.[13]

In essence, the Cengiz et al. (2019) estimator effectively answers the question of the impact of a change in the minimum wage in terms of changes in employment at each point in the wage range, which is useful for understanding who benefits and who loses from a minimum wage change. But it is less useful for understanding the answer to the third research question in this area: how does the minimum wage alter the shape of the wage distribution. If we took the Cengiz et al. (2019) estimates as an answer to that question then we would conclude that minimum wages do not alter the wage density for wages well below m and have somewhat small spillover effects. In our data generation process, however, neither is the case. Of course, the total shape effect is a more complicated entity to identify - it can't be seen by just directly comparing employment levels when minimum wages change - and to get it we require an extra assumption. In particular, we require the assumption that the minimum wage has no effect on the wage distribution

[13]Going back to the case illustrated in Figure 2, the Cengiz et al. (2019) estimator would recover the difference in spillovers at the old (\$8) and new (\$9) minimum wage in the \$9 to \$10 range, which could be arbitrarily small depending on the precise shape of the spillover effects. In contrast, the hazard approach recovers the underlying shape of the spillover effects using the recursive argument discussed in Section 3.

beyond some threshold above the minimum wage in combination with our cascading identification strategy. We view this as credible but, as with all identifying assumptions, cannot be directly tested. We provide suggestive evidence in favour of it in Appendix A by showing that our estimates do not change when we increase the value of the threshold beyond which minimum wages are imposed to have no effect for values beyond $5 above the minimum wage.[14]

Our conclusion from the Monte Carlo exercise is that the quantile estimator is effective for examining overall movements in wage inequality and the Cengiz et al. (2019) estimator is effective for identifying the employment change effects of minimum wages. Neither, though, is well suited to the question we seek to address: how minimum wages alter the shape of the wage distribution. On the other hand, the hazard estimator does capture those shape effects, and we focus on it through the rest of the paper.

5 Data and Minimum Wage Variation

5.1 Wage Data

Our wage data comes from the Canadian Labour Force Survey (LFS) master files.[15] The LFS is a large Canadian household survey interviewing approximately 50,000 households per month that is similar in nature to the US CPS or UK LFS. We restrict our LFS sample to individuals aged 15 to 60. The self-employed are not included in the analysis as wage information is not available for these workers in the LFS. We further exclude full-time students because working is not their main activity. The LFS began to include a question on wages each month starting in November 1996. For our main estimates, we work with the data at the monthly level, with our data period running from January 1997 through December 2016. We construct education dummy variables for four groups: high school drop outs; high school graduates plus those with some but not completed post-secondary; people with a post-secondary diploma or certificate; and people with a university degree. We also use a set of age dummy variables corresponding to: 15 to 19;

[14]Notice that we don't expect the hazard estimates to be unchanging with *any* value of the threshold. If we set the threshold at a place where there are ongoing spillover effects (e.g., at $1 above the minimum wage in our Monte Carlo data), we get different estimates - as we should. In that case, we are treating an area where there is still a minimum wage effect as if it corresponds to the baseline hazard, resulting in small estimated effects below that point.

[15]The master files were accessed on site at the Manitoba and Carleton, Ottawa, Outaouais local Research Data Centre (COOL RDC).The research and analysis are based on data from Statistics Canada and the opinions expressed do not represent the views of Statistics Canada.

20 to 24; 25 to 34; 35 to 54; and 55 to 60. Our wage variable is the hourly wage deflated to 2002 dollars using the consumer price index (CPI). Hourly wages are reported directly for workers paid by the hour but constructed using reported hours of work for those paid by the week, month or year. We have examined the robustness of our results to using three different wage samples: all earners; hourly paid earners; and hourly paid earners excluding those who receive tips or commission. We use the all earners sample for our main estimates but found similar results for the two other samples in specifications not reported here. We also include a dummy variable equal to 1 when an integer value is included in a wage bin and 0 otherwise to address potential measurement error concerns arising out of the fact that a high proportion of workers (30%) report an integer valued wage.[16]

As we describe in Section 2, minimum wage effects for groups of workers defined by movements into and out of jobs are central to our discussion. To construct those groups, we take advantage of the rotating panel design of the LFS. Individuals remain in the sample for six consecutive months, and every month one-sixth of the panel is replaced. As such, one can link consecutive months of the LFS thereby creating six-month mini panels.[17] Combining this with the consistent LFS job tenure question which indicates for how long workers have been with their current employer, we construct three sub-groups:[18] 1) Job Leavers - workers who are observed on a job in the reference week, but who are not in that job two months later, whether because they have left employment or moved to a different job; 2) Job Joiners - workers who are on a job in the reference week, but have been on that job for two months or less; 3) Job Stayers - workers who have been on the same job for at least three months.

Note that we look two months instead of one month ahead to identify Job Leavers due to the timing of minimum wage changes that happen on the 15th of the month in a few cases (the minimum wage is raised on the first day of the month in most cases). Since the LFS reference week is typically defined to include the 15th of the month, it is unclear how a minimum wage increase on the 15th day of a month would affect reference

[16] Dube et al. (2020) show that bunching at integer values is pervasive in wage data, though not all of the bunching is due to measurement error. Some bunching also occurs in high-quality administrative data.

[17] A detailed description of how the data was linked can be found in Appendix D.

[18] The LFS asks when individuals started working for their current employer. Statistics Canada uses this information to construct a job tenure variable indicating the number of months workers have been with their employer. What distinguishes the LFS from other Canadian data sets, and US data sets, is that the job starting data question (with no change in wording) has been asked every month since 1976. See Brochu (2013) for a detailed discussion of the limitations of other North American data sets.

week wages. To avoid this complication, we use a two-month window in defining Leavers, Stayers, and Joiners to insure a minimum wage change clearly happens between the two reference weeks.

Another challenge with the LFS is that respondents are only asked about their wage on a job at the first interview or if they take up a new job. Thus, if they do not change jobs, their recorded wage becomes "stale" in their subsequent months in the survey. For that reason, we restrict our attention to the incoming sample members each month. This means that we cannot construct panel estimates of, for example, the change in the wage for a Stayer at the time of the minimum wage change. Instead, as discussed in Section 3.1 we estimate the net effect of the minimum wage on the Stayer wage distribution by taking advantage of the mini-panel aspect of the LFS and the job tenure question to compare a statistically similar group of Job Stayers before and after a change in the minimum wage.

5.2 Minimum Wages

We use provincial minimum wage data for the period covered by our wage data: 1997-2016. The minimum wage is set at the provincial level in Canada, which introduces considerable variation in the minimum wage.[19] At times, some provinces have implemented lower rates for special classes of workers (e.g. students in Ontario). Most empirical studies indicate that firms do not take advantage of these special categories (e.g., Card and Krueger (1995)). Thus, we will focus on the general adult minimum wage for each province.[20] As mentioned above, in a few cases provinces increase their minimum wage on the 15th day, as opposed to the 1st day, of the month. For the sake of consistency, we define the monthly minimum wage as the one prevailing on the first day of the month. We convert both minimum wages and our LFS wage measure to real terms using the monthly CPI (using 2002 as the base year).

Figures 4 and 5 show the evolution of real minimum wages by province for our sample period. Figure 4 shows the patterns for the four largest provinces. It is clear from this figure that there is considerable variation within provinces over time and, since the provincial trends often move in different directions, this variation will be available

[19]There are a set of workers who fall under federal jurisdiction (e.g. air transport). Since 1996, the federal government has mandated that these workers be covered by the minimum wage in the province where the employee is usually employed.

[20]We have estimated specifications in which we allow for different minimum wage effects when special lower minima were in place. The main patterns and conclusions from those specifications are the same as from the simpler specification without controlling for the lower minima that we present here.

even after controlling for general time effects. In the period between 2001 and 2009, for example, Ontario and Quebec showed flat trends in real terms while the BC real minimum wage was strongly declining and the Alberta minimum wage was strongly increasing. The patterns in Figure 5 for the six smaller provinces are much more similar across provinces, though even here there are some clear trend differences (e.g., Newfoundland between 2005 and 2011). Altogether, there are 157 nominal minimum wage changes in our time period. Many of these are small - the median nominal change is 3.8% - but 47 of the changes are 5% or larger and the largest is an 18% change in Alberta in 2006.

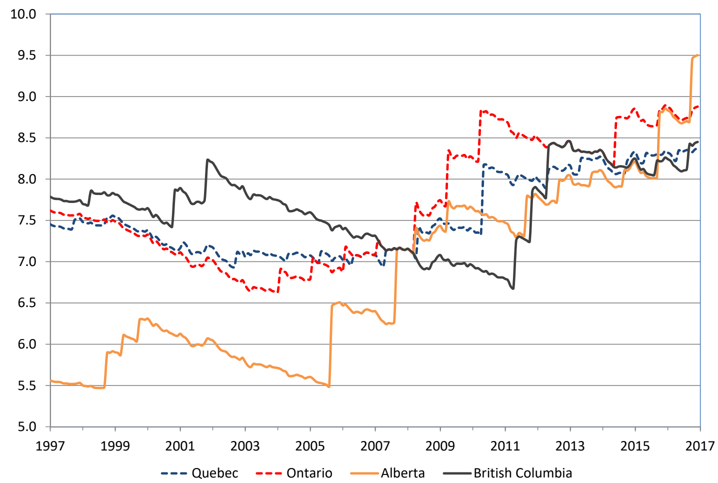

Figure 4: Real Value of the Minimum Wage ($2002), Larger Provinces

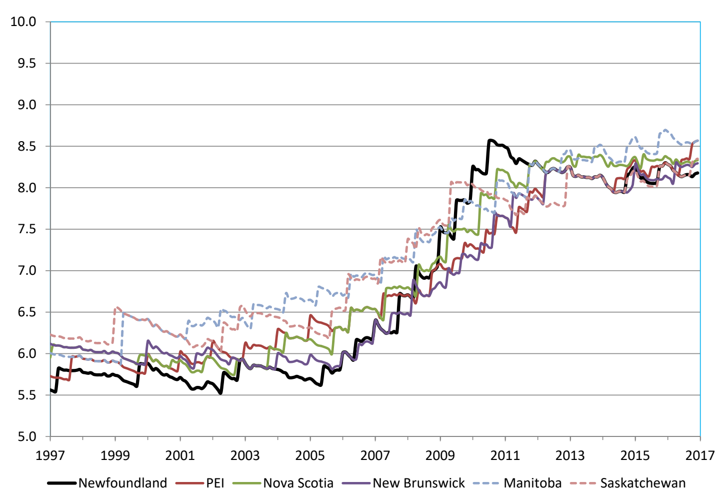

Figure 5: Real Value of the Minimum Wage ($2002), Other Provinces

Before turning to our main estimates, we present results from an event study specification in which we examine movements in the 5^{th}, 10^{th}, and 15^{th} percentiles of the wage distribution around the time of a minimum wage increase in Appendix E. We treat province x quarter observations in which the minimum wage increased by 5% or more as treatments and all other observations as controls. There is no evidence of pre-trends and clear increases in the 5^{th} and 10^{th} percentiles (the range in which the minimum wages occur), limited 15^{th} percentile effects and none at the 20^{th} percentile. We interpret these results as indicating that the variation in the minimum wages that we use is plausibly exogenous and that the minimum wage has substantial effects on the wage distribution at and just above the minimum wage. From our Monte Carlo results with quantile estimators, though, it is not possible to discern the actual impact of minimum wages on the shape of the wage distribution from these results.

6 Estimation Results

6.1 Implementation Details

In implementing our estimator with the LFS data, we top-code wages at \$20 per hour since the computing effort required to estimate the hazard beyond that point is not useful when estimating minimum wage effects.[21] We employ 164 baseline wage bins and 5 covariate segments. The first three baseline bins correspond to: $\leq$ \$3.00; \$3.01 to \$3.50; and \$3.51 to \$4.00. The remaining bins up to \$20.00 are 10 cents wide. The covariate segments consist of a first segment that goes up to \$10.00, and four other segments that equally divide the \$10.00 to \$20.00 range into \$2.50-wide bands. Each contains about 20% of the wages within our wage range (i.e., the wages between \$0 and \$20). The minimum wage values all fall in the first segment, but our estimated spillover effect ranges extend into higher segments (up to the fourth segment when using spillovers up to \$7 above the minimum wage in Appendix A).

Combining the baseline hazard coefficients, the regular covariates, various types of time and province effects, and the minimum wage variables, the total number of estimated coefficients is as high as 769, depending on the specification. In the actual implementation of the estimator, we first count the number of observations in cells defined by the complete interactions of the baseline wage bins, age, education, province, year, and month of the

[21]Approximately 35 percent of workers earn below \$20 an hour in our estimation sample (recall that the wage is expressed in real, 2002 dollars). Note that all observations, including those above \$20 an hour, are used to compute the hazard rates. As such, the estimation procedure accounts for the fact that a fraction of observations "survive" above the \$20 threshold.

year (recalling that all our estimation is done separately for males and females). Thus, we group our data rather than working directly with individual observations, which is a first step in obtaining standard errors on our estimated coefficients that are at the right level of clustering. In addition, working with the data in this way allows us to take advantage of the insight in Ryu (1992) that a proportional hazard model with grouped data can be rewritten as a Generalized Linear Model in which the dependent variable is the transformation of the hazard rate in a cell given by $\ln(-\ln((1-h(y)))$. This allows us to reduce the computational burden and work with the GLM command in Stata.[22]

6.2 Results for all Workers

In order to ease exposition, we begin by reporting the results that include short term but not medium term dynamics, (i.e., the K_{at} interactions are included but not the $K_{1-2,t}$ and $K_{3-12,t}$ interactions). We focus on the medium term dynamics in the next subsection. Table 1 reports the estimated "steady-state" minimum wage coefficients (i.e., the coefficients corresponding to periods when there is no recent or upcoming minimum wage change) for all workers and for our three subgroups, broken down by gender. The set of dummies capturing short term effects, D^c_{mp}, are included in all the models, and we report the coefficients on these variables in the next table.

We begin by considering the estimates for all workers, pooled together. Based on the discussion in Section 2, the "all workers" specification estimates will reflect a combination of selection and shape effects. Looking at the results for all female workers in the first column of Table 1, we see strong reductions in the hazard rate below the minimum wage, with the effects being larger the farther below the minimum wage a wage is. There is a large, statistically significant spike in the distribution at the minimum wage and spillover effects up to \$2.50 above the minimum wage. The spillover effects decline in size as we move farther above the minimum wage but are all statistically significant at at least the 10% level up to \$2.50 above the minimum wage, with further offsetting positive and negative effects in the \$3.50 to \$4.50 range. In Appendix F, we present results using the Cengiz et al. (2019) estimator with our data. Those results are very similar to their Figure 2, showing a large negative spike at \$1 below m, a large positive spike at m, and some spillovers up to about \$2 above m. Fitting with what we observed in the Monte

[22]If there were wage observations in all of our cells (i.e., if the hazard rate were non-zero in each cell) then this estimator could be implemented with simple OLS. Because this is not the case here, we are forced to use the maximum likelihood option in GLM, which slows down the estimation time. Thus, there is a potential trade-off between having a more smoothed baseline estimate (i.e., having fewer cells) and speed of estimation. We chose to work with a more flexible baseline specification.

Table 1: Minimum Wage Effects on the Hazard
with No Current Change in the Minimum Wage

Wage Range	**All Workers**		**Stayers**		**Leavers**		**Joiners**	
	Females	Males	Females	Males	Females	Males	Females	Males
Over \$0.50	−1.76**	−1.39**	−1.72**	−1.37**	−1.46**	−1.24**	−1.91**	−1.37**
below	(0.28)	(0.29)	(0.3)	(0.33)	(0.32)	(0.42)	(0.32)	(0.29)
\$0.30 to 0.50	−0.73	-0.32	−0.76	-0.33	-0.53	-0.22	-0.41	-0.33
below	(0.42)	(0.45)	(0.40)	(0.45)	(0.34)	(0.51)	(0.33)	(0.42)
\$0.10 to 0.30	−0.54**	−0.30**	−0.55**	−0.29**	−0.18	-0.06	−0.43	−0.29**
below	(0.11)	(0.11)	(0.13)	(0.09)	(0.17)	(0.27)	(0.29)	(0.09)
At min wage	1.95**	2.03**	1.89**	1.98**	2.19**	2.25**	2.39**	1.98**
	(0.23)	(0.27)	(0.25)	(0.21)	(0.31)	(0.25)	(0.27)	(0.21)
\$0.10 to 0.30	0.87**	0.84**	0.84**	0.83**	1.00**	0.94**	1.11**	0.83**
above	(0.14)	(0.11)	(0.15)	(0.15)	(0.18)	(0.28)	(0.23)	(0.15)
\$0.30 to 0.50	0.74**	0.74**	0.72**	074**	0.86**	0.82**	0.88**	0.74**
above	(0.10)	(0.17)	(0.11)	(0.15)	(0.19)	(0.31)	(0.17)	(0.16)
\$0.50 to 1	0.51**	0.45**	0.50**	0.35**	0.54**	0.54**	0.64**	0.46**
above	(0.07)	(0.12)	(0.08)	(0.11)	(0.12)	(0.21)	(0.17)	(0.11)
\$1 to 1.50	0.33**	0.25	0.32**	0.25	0.35**	0.29	0.46*	0.25
above	(0.07)	(0.13)	(0.08)	(0.13)	(0.13)	(0.22)	(0.17)	(0.13)
\$1.50 to 2.00	0.34**	0.27	0.34**	0.27	0.36*	0.32	0.49*	0.27**
above	(0.07)	(0.14)	(0.07)	(0.14)	(0.14)	(0.26)	(0.13)	(0.14)
\$2.00 to 2.50	0.20**	0.08	0.20	0.11	0.10	0.11	0.24	0.11
above	(0.05)	(0.08)	(0.05)	(0.08)	(0.06)	(0.18)	(0.13)	(0.08)
\$2.50 to 3.00	0.09	0.04	0.09	0.07	0.00	0.11	0.12	0.07
above	(0.06)	(0.06)	(0.06)	(0.06)	(0.07)	(0.17)	(0.14)	(0.06)
\$3.00 to 3.50	0.12	0.13	0.12	0.14	0.07	0.17	0.18	0.14
above	(0.08)	(0.08)	(0.08)	(0.07)	(0.07)	(0.18)	(0.11)	(0.07)
\$3.50 to 4.00	0.06**	0.04	0.05**	0.06	0.02	0.02	0.16*	0.06
above	(0.02)	(0.06)	(0.03)	(0.07)	(0.05)	(0.14)	(0.07)	(0.07)
\$4.00 to 4.50	−0.05**	−0.05	−0.05**	−0.04	−0.06	−0.10	0.00	−0.04
above	(0.02)	(0.03)	(0.02)	(0.03)	(0.04)	(0.08)	(0.10)	(0.03)
\$4.50 to 5.00	−0.01	0.00	-0.02	0.00	−0.04	0.01	0.06	0.00
above	(0.02)	(0.02)	(0.02)	(0.07)	(0.06)	(0.04)	(0.06)	(0.03)
Num Obs	701137	685396	660231	640589	59894	63596	40906	44807

*,** significant at 1%, 5% level

All specifications are based on 164 baseline wage bins and 5 covariate segments. All include controls for age and education, a dummy indicating whether a particular bin corresponds to an integer in the nominal wage data, a complete interaction of province by year effects in the first covariate segment and separate province and time effects in other segments, the bin number interacted with province dummies, and quarter dummies.

Carlo exercise, the spillovers are smaller and extend less far up the distribution than what we estimate using our hazard estimator. Thus, the differences have to do with the estimator rather than the different countries (Canada vs. the United States).

In Figure 6, we plot the exponentiated coefficients along with the estimated CDF for 20-24 year old high school graduates living in the province of Ontario with and without the minimum wage effects enacted, where the latter serves as a benchmark for interpreting the size of our estimated minimum wage effects.[23] Those effects are sizeable. The counterfactual CDF without the minimum wage is at 0.103 just below the minimum wage while the CDF including minimum wage effects is at 0.032 at that point. The CDF with the minimum wage more than doubles at the minimum wage point, rising to 0.074. Spillovers take the CDF higher and the two lines converge at $2 above the minimum wage. The exponentiated minimum wage effect coefficients show that the minimum wage generates over a 400% increase in the wage hazard rate at the minimum wage and the spillover effects in the region above the minimum wage correspond to an 80% increase in the hazard rate just above the minimum wage that gradually declines as we move up in the wage distribution.

The results for men (column 2 of Table 1 and Figure 7) are remarkably similar to those for women. Recall that our estimated minimum wage effects represent proportional shifts in the underlying hazard function. Thus, the similarity of the estimated male and female minimum wage effects imply that estimates such as those in DiNardo et al. (1996) showing larger effects of the minimum wage on the female wage distribution arise because the female wage distribution has more mass in the region near the minimum wage rather than because minimum wages have a larger effect on the female labour market, per se. This is an intriguing result that suggests that employers do not assign female low wage workers to the minimum wage more frequently than they do low wage male workers, that is, that any discrimination does not take this specific form.

The remaining columns of Table 1 contain the estimated coefficients for the "steady state" effects of minimum wages on the CDFs for Stayers, Leavers and Joiners. Plots of the CDF without minimum wage effects (not shown here) reveal, not surprisingly, that both the Leavers and Joiners wage distributions have more weight at lower wages than the distribution for Stayers. The results in Table 1 show, in addition, that both the Leavers and Joiners distributions also have larger spikes at the minimum wage and more mass just above the minimum wage than for the Stayers distribution. This seems

[23] To simplify the exposition, we present estimates from a model with spillovers running up to $2 above the minimum wage.

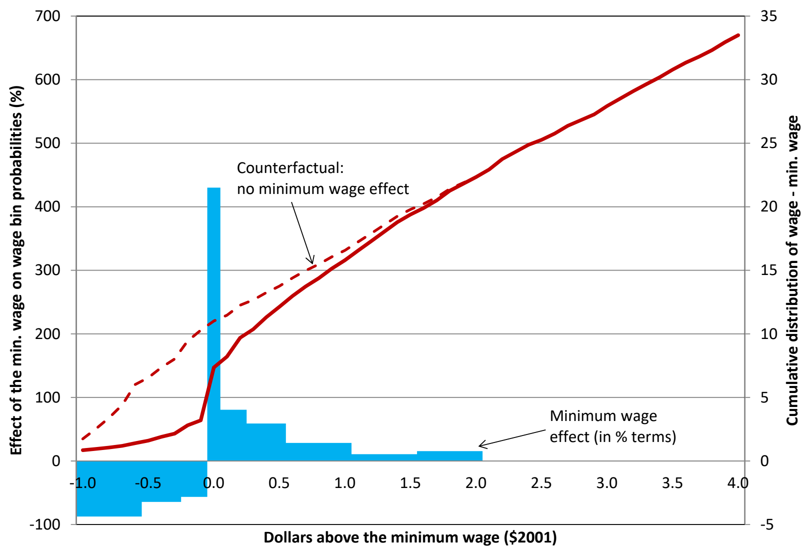

Figure 6: Effect of the Minimum Wage on the Female Wage Distribution

reasonable because it indicates that when the distribution is compressed with a minimum wage and workers whose wages would otherwise be below the minimum wage are now at or just above the minimum wage, layoffs will be more likely in this compressed area. But it is equally the case that new hires are likely to happen in this bottom region of the distribution and, in fact, the comparable columns in Table 1 contain quite similar estimates for the effects at and above the minimum wage for Leavers and Joiners.

6.3 Allowing Minimum Wage Effects to Vary with the Minimum Wage Level

To this point, we have estimated minimum wage effects as proportional shifts in the baseline hazard. This means, for example, that the minimum wage will increase the baseline hazard by the same 400% whether the level of the minimum wage is high or low and that the spike will be larger in absolute terms in regions in which the wage distribution has more mass. The imposed assumption that the minimum wage effect is proportional in this way seems potentially overly restrictive. To examine whether this is the case, we implement our main specification including interactions of the value of the minimum wage with each of our minimum wage range dummy variables. This allows the minimum wage spike, spillovers, and the reduction in mass below the spike to vary with the value of the minimum wage.

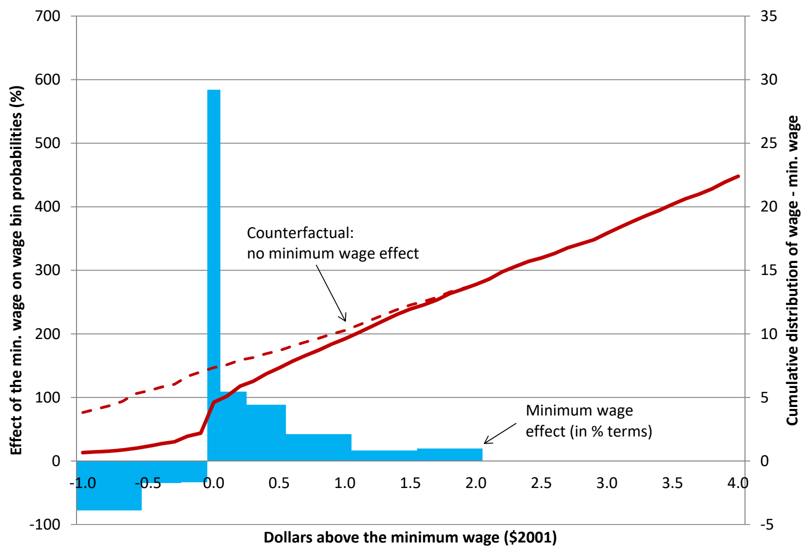

Figure 7: Effect of the Minimum Wage on the Male Wage Distribution

In Table 2, we present the estimated coefficients of the interactions of m with each of the minimum wage range dummies in a model with males and females pooled.[24] The estimated effects are mainly statistically insignificant, with the exception of the spillovers between \$0.50 and \$2.0 above the minimum wage, where the negative estimated coefficients imply that there are smaller spillovers at higher minimum wages, and a significant effect between \$3.50 and \$4.00 above the minimum wage which has the opposite sign. It is difficult, of course, to understand the size of the interactions from these estimates and, so, we present fitted densities based on these estimates in figure 8. In particular, we predict the density setting m = \$6.75 and m=\$7.75. We then recentre the \$7.75 plot on \$6.75 so that the two spikes are in the same place and we can see the differences in the shapes of the wage distribution for the different m values. Importantly, the two densities are substantially similar: in both cases, we see a spike at the minimum wage and a spillover region extending to about \$2.50 above m. At a more granular level, the minimum wage spike is slightly higher with the higher minimum wage (0.064 versus 0.059) and the spillover effect is smaller by about 0.004 across the range from \$1 to \$2 above the minimum wage. The offsetting higher spillover between \$3.50 and \$4.00 above m is hard to see. Interestingly, the total mass in the density between m and (m + \$5) is almost identical with the two minimum wage values (0.59 with m=\$6.75 and 0.57 with

[24]The remaining estimates from the model - the uninteracted wage range effects and the coefficient on the minimum wage variable itself - are available upon request.

m=$7.75). Our conclusion is that higher minimum wages result in slightly higher spikes at m and an offsetting reduction in spillovers just above m, though not to a particularly large degree. We will return to this result when we discuss implications for different models. In the meantime, we will continue our discussions with the model without interactions with m since they are easier to interpret and the interaction effects do not alter our main conclusions

Table 2: Coefficients on the Interactions of the Minimum Wage Value with Minimum Wage Range Dummy Variables

Over \$0.50	-0.019	\$0.10 to 0.30	-0.023	\$2.50 to 3	0.029
below	(0.048)	above	(0.136)	above	(0.026)
\$0.30 to 0.50	-0.158	\$0.30 to 0.05	-0.205	\$3 to 3.50	0.050
below	(0.173)	above	(0.138)	above	(0.051)
\$0.10 to 0.30	-0.253	\$0.50 to 1	-0.226^{*}	\$3.50 to 4	0.068^{**}
below	(0.196)	above	(0.126)	above	(0.025)
At min wage	0.080	\$1 to 1.50	-0.181^{**}	\$4 to 4.50	0.045
	(0.150)	above	(0.090)	above	(0.029)
		\$1.50 to 2	-0.104^{**}	\$4.50 to 5	-0.003
		above	(0.052)	above	(0.025)
		\$2 to 2.50	-0.065		
		above	(0.060)		

*,** significant at 10%, 5% level

The table entries are estimated coefficients on variables constructed as the interaction of the value of the minimum wage with dummy variables corresponding to the stated ranges relative to m. Standard errors are in parentheses. See Table 1 for further estimation details.

6.4 Dynamic Effects for Stayers, Joiners, and Leavers

We turn, next, to the effects for the Stayers, Leavers, and Joiners samples. The responses to minimum wage changes, as summarized by the β^c coefficients on the K_{at} variables are reported in Table 3 (again, for a specification with no further lagged effects).

6.4.1 Stayers: Evidence on Shape Changes

Changes in the wage CDF for Stayers, because they are by definition a fixed set of people before and after a minimum wage change, do not reflect selection effects. Instead, if their wage CDF changes at the time of a minimum wage increase (relative to changes in jurisdictions without a minimum wage change), the change must reflect a difference

Table 3: Minimum Wage Effects on the Hazard
with Recent or Upcoming Minimum Wage Change

Wage Range	**All Workers**		**Stayers**		**Leavers**		**Joiners**	
	Females	Males	Females	Males	Females	Males	Females	Males
Over $0.50 below	0.10	−0.19*	0.09	-0.15	0.15	-0.25	0.11	−0.15
old min wage	(0.07)	(0.09)	(0.06)	(0.08)	(0.11)	(0.15)	(0.23)	(0.08)
$0.30 to 0.50 below	−0.70**	−0.67**	−0.67**	−0.57**	0.15	-0.10	−0.88**	−0.57**
old min wage	(0.11)	(0.06)	(0.08)	(0.07)	(0.13)	(0.10)	(0.27)	(0.07)
$0.10 to 0.30 below	-0.24	-0.14	-0.25	-0.18	0.01	0.29	-0.13	-0.18
old min wage	(0.29)	(0.17)	(0.30)	(0.23)	(0.25)	(0.31)	(0.32)	(0.23)
At old min wage	−1.41**	−1.37**	−1.33**	−1.30**	-0.10	-0.07	−1.94**	−1.30**
	(0.15)	(0.09)	(0.14)	(0.08)	(0.11)	(0.11)	(0.18)	(0.08)
Between old and	−0.74**	−0.66**	−0.74**	−0.62**	-0.01	0.05	−0.74**	−0.62**
new min wage	(0.10)	(0.05)	(0.10)	(0.05)	(0.11)	(0.19)	(0.16)	(0.05)
At new min wage	0.87**	1.12**	0.82**	1.05**	0.30*	0.50**	1.18**	1.06**
	(0.07)	(0.12)	(0.08)	(0.12)	(0.15)	(0.17)	(0.11)	(0.12)
$0.10 to 0.30 above	0.46**	0.52**	0.42**	0.52**	0.11	0.34*	0.63**	0.52**
new min wage	(0.07)	(0.06)	(0.07)	(0.06)	(0.10)	(0.14)	(0.12)	(0.06)
$0.30 to 0.50 above	0.14*	0.15*	0.15**	0.14	0.08	0.09	0.06	0.15
new min wage	(0.05)	(0.07)	(0.04)	(0.08)	(0.09)	(0.05)	(0.09)	(0.09)
$0.50 to 1 above	0.09*	0.06	0.08*	0.05	0.18**	0.14*	0.12	0.06
new min wage	(0.03)	(0.04)	(0.04)	(0.04)	(0.05)	(0.05)	(0.07)	(0.04)
$1 to 1.50 above	-0.04	-0.10	-0.02	-0.09	-0.05	-0.01	−0.27**	-0.09
new min wage	(0.03)	(0.04)	(0.03)	(0.05)	(0.08)	(0.09)	(0.08)	(0.05)
$1.50 to 2.00 above	-0.07	0.00	-0.05	-0.03	-0.05	0.02	−0.31**	-0.03
new min wage	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.05)	(0.09)	(0.03)
$2.00 to 2.50 above	0.05	0.01	0.04	0.02	0.06	-0.04	0.03	0.02
new min wage	(0.02)	(0.02)	(0.03)	(0.03)	(0.06)	(0.05)	(0.09)	(0.02)
$2.50 to 3.00 above	-0.05	0.01	-0.04	0.00	0.13	0.02	−0.19**	0.00
new min wage	(0.03)	(0.06)	(0.04)	(0.05)	(0.04)	(0.07)	(0.07)	(0.05)
$3.00 to 3.50 above	0.02	-0.08	0.02	-0.08	-0.10	0.06	-0.07	-0.08
new min wage	(0.05)	(0.06)	(0.05)	(0.07)	(0.11)	(0.08)	(0.07)	(0.07)
$3.50 to 4.00 above	-0.01	-0.01	-0.02	-0.01	0.03	0.05	0.09	-0.01
new min wage	(0.03)	(0.02)	(0.02)	(0.03)	(0.06)	(0.09)	(0.14)	(0.03)
$4.00 to 4.50 above	0.04*	0.01	0.04	0.03	0.03	0.09	-0.01	0.03
new min wage	(0.02)	(0.03)	(0.03)	(0.04)	(0.03)	(0.06)	(0.08)	(0.04)
$4.50 to 5.00 above	-0.02	-0.06	-0.02	-0.06	0.00	-0.16	−0.19*	-0.06
new min wage	(0.03)	(0.05)	(0.02)	(0.06)	(0.07)	(0.09)	(0.09)	(0.06)
Num Obs	701137	685396	660231	640589	59894	63596	40906	44807

*,** significant at 1%, 5% level
See the footnote to Table 1 for details on the specifications.

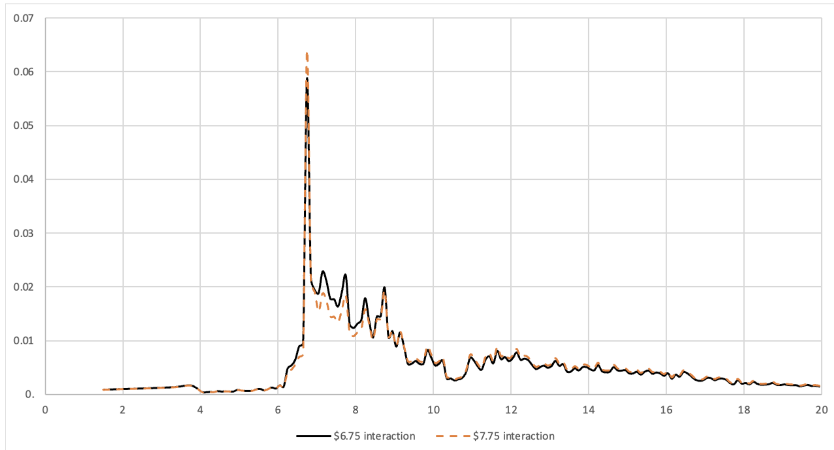

Figure 8: Fitted Densities from Interaction Specification
with m = \$6.75 and m = \$7.75

in the way continuing workers are paid. Most directly, this could reflect moving workers whose initial wage was below the new minimum wage up to the minimum to meet legal requirements. But changes in the shape of the CDF for Stayers could also reflect relative shifts in bargaining power and/or fairness concerns.

In the case of Stayers, we graphically illustrate the impact of a 50 cent minimum wage increase in Figure 9, for women. The patterns for men are very similar, and we present results for them in Appendix A. The effect of the minimum wage on the wage bin probabilities before a minimum wage increase is based on the steady-state coefficients β^b estimated for the "backward" Stayers. We then add the β^c coefficients applied to a 50 cent change to obtain the second set of effects reported in the figure. For example, in the case of women we add the -1.33 change in the probability of being at the old minimum wage (Table 3) to the 1.89 steady-state effect of being at the minimum wage (Table 1). We also plot the error bars to indicate whether the changes are statistically significant. For example, the standard error on the change at the old minimum wage is 0.14, yielding an error bar of $+/-$ 0.28. The steady-state estimate of 1.89 falls well outside of the error bars, indicating that the 50 cents minimum wage increase significantly reduces the probability of being at the old minimum wage.

Interestingly, the results for both women and men indicate there is still a small abnormal concentration of observations at the old minimum wage after a minimum wage increase, suggesting some inertia in how firms adjust their wages. This may be particularly possible for workers who are paid at the weekly or lower frequency rather than hourly since their compliance with minimum wages may be more difficult to assess.

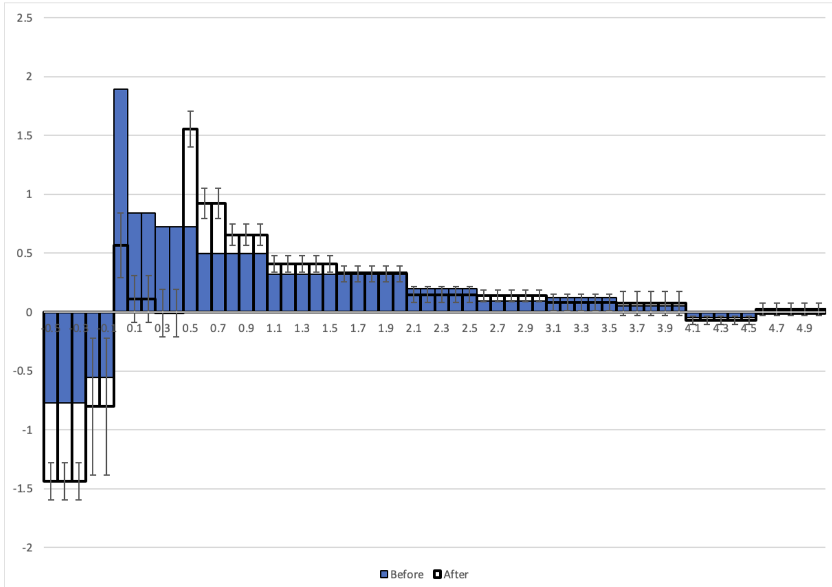

Figure 9: Effect of the minimum wage on wage bin hazards just after a 50 cent increase in the minimum wage (Stayers): Women

Comparing the hazard effects before versus after a minimum wage increase suggests some similarity in shape. It seems possible, however, that this is only because we are, so far, only examining immediate impacts of a minimum wage change and firms may take time to adjust their wage policies (Sorkin (2015)). We investigate this by estimating specifications in which we include the 1 to 2 month and 3 to 12 month lag interactions in the $K_{1-2,t}$ and $K_{3-12,t}$ matrices. In figure 10, we plot fitted densities from the Stayers dynamic specification, examining the impact of a minimum wage change from \$6.75 to \$7.25 over time. The solid line with no markers corresponds to the fitted density when the minimum wage (m) has been at \$6.75 for at least 12 months. That density shows a now familiar pattern with a spike at m and extra mass just above it. By two months after the change in m, the former spike has been reduced to about one quarter of its former size and the mass between \$6.75 and \$7.25 has also been reduced. Meanwhile a new spike has emerged at \$7.25 along with spillover effects above the new m. By 3 to 12 months after the increase, the mass at the old m and between the old and new m values has decreased further but the height of the spike at \$7.25 has increased only marginally and the spillover pattern above the new m is virtually unchanged. Finally, it is difficult to discern any difference between the density at 3 to 12 months after the m increase and the density at \$7.25 when m has not changed for over 12 months, especially in the range at and above m. Thus, our conclusion is that the wage policy responses evident for the

Stayers emerge very quickly, with much of the change having occurred even by 2 months after a minimum wage increase.

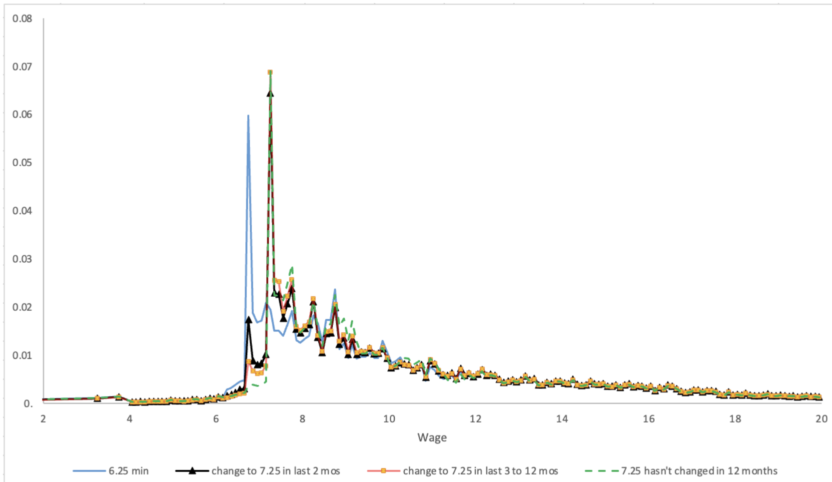

Figure 10: Minimum Wage Dynamics: \$6.75 to \$7.25 Increase (Stayers)

6.4.2 Leavers: Evidence on Selection

Assuming that firms don't alter wage policies in advance of minimum wage increases, the impact of the minimum wage on the job Leaver CDF identifies selection effects.[25] For example, if firms respond to a minimum wage increase by laying off their lowest ability/lowest wage workers then we should see extra mass at the low end of the wage distribution of Leavers in minimum wage increase periods relative to other periods. In terms of our estimator, this would show up as an increase in the spike at the old minimum wage and in the mass between the old and new minimum wages. The Leavers columns in Table 3, however, show small and statistically insignificant coefficient estimates for both of these effects. In fact, the estimated effects of the minimum wage increase on the Leavers' pre-separation distribution are nearly all statistically insignificant for both men and women. [26] The only exceptions being a positive and statistically significant effect at the new minimum wage.

Many of the theoretical claims about the effect of a minimum wage increase in terms of worker separations focus on either weak firms shutting down or ongoing firms laying

[25] As noted earlier, we find limited evidence of anticipatory effects in wage setting before a minimum wage increase for Job Stayers. We assume that conclusion carries over to the distribution for Job Leavers.

[26] A possible explanation for these findings is that firms do not layoff additional workers in response to a minimum wage increase, but reduce hiring to keep employment constant since workers are now less likely to quit. These predictions are consistent with the dynamic model of Sorkin (2015) and the empirical findings (lower turnover) of Brochu and Green (2013).

off their less able workers. In figure 11, we present the effect of the minimum wage on the hazard rate before and after a \$0.50 increase in the minimum wage (similar to Figures 9 for female workers who were laid off).[27] The similarity between the No Minimum Wage change plot and the plot corresponding to an upcoming increase in the minimum wage is remarkable, apart from some extra mass at the new minimum wage. Further, in our extended, dynamic specification, there is also no evidence of changes in who is laid off in the year following a minimum wage increase. We conclude that there is essentially no selection effect of a minimum wage increase in terms of the pre-separation wages for workers losing their jobs due to a layoff or firm closure.

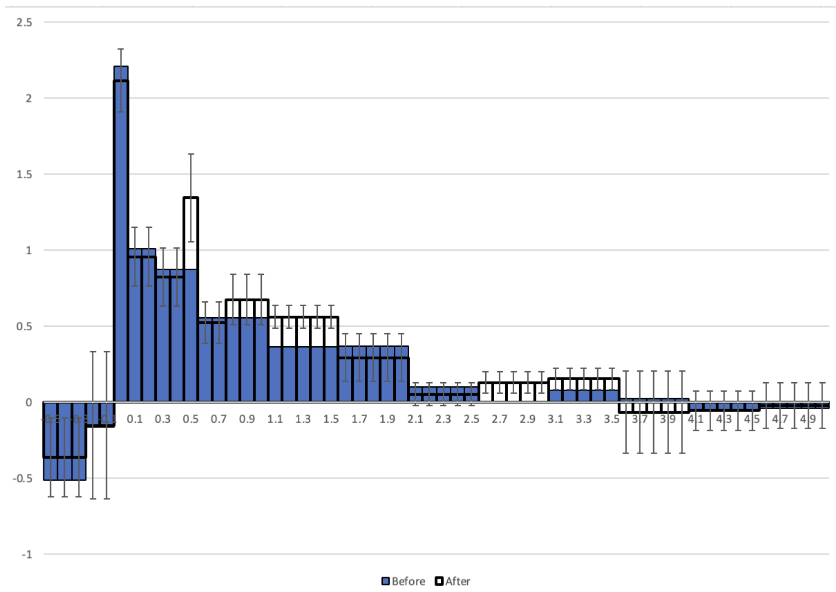

Figure 11: Effect of the minimum wage on wage bin hazards just after a 50 cents increase in the minimum wage (Layoffs): Women

In Figure 12, we present the same plot for workers who quit their previous job.[28] Because of small numbers of male job quitters, the plot for them has very large standard errors and is somewhat erratic. To provide more readable results (and because our results for men and women are very similar), we present a plot for men and women combined. Even then, the error bars in the figure are large, showing few statistically significant

[27] We define lay-offs as job separations for which the worker gave as a reason for losing their previous job as one of: end of seasonal job; end of temporary or casual job; company moved or out of business; business conditions; or dismissal or other reasons. Thus, layoffs include separations due to business closings, the end of temporary jobs, and being let go by an ongoing firm. Results for men are, again, found in the appendix.

[28] Quitters are job separators who give the reason for job loss as one of: Own illness or disability; caring for own children; pregnancy; other personal or family responsibilities; going to school; dissatisfied; retired.

effects. There is, though, some evidence of positive effects at and just above the new minimum wage - a point to which we will return in our discussion of implications for economic models.

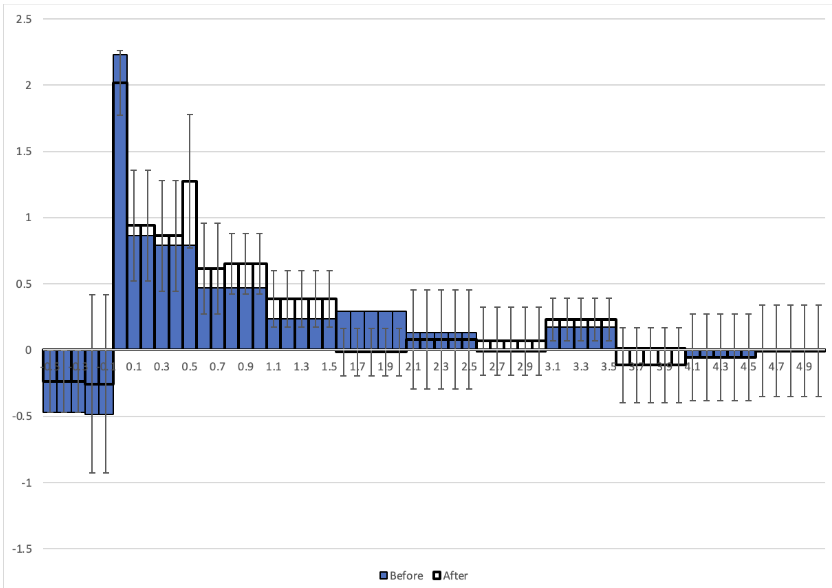

Figure 12: Effect of the minimum wage on wage bin hazards just after a 50 cents increase in the minimum wage (Quits): Men and women combined

6.4.3 Joiners

For completeness, we also present the results for new job Joiners. Again, we present a plot of the impacts of a minimum wage on the hazard rate before and after a \$0.50 minimum wage increase (in Figure 13). We just present the plot for women, putting the figure for men (which is very similar) in the Appendix. The figure shows how a minimum wage increase affects wages for those who were hired in the previous 3 months and how that changes just after a minimum wage increase. As such, our estimated effects for Joiners reflect a combination of changes in who is hired (selection) and in how wage policies change (what we call wage policy shape changes).

The Joiners plot shows a pattern that is quite similar to that for Stayers. In the wake of a minimum wage increase, there are substantial decreases in mass at the old minimum wage and between the old and new minimum wage. At the same time, a new spike grows in at the new minimum wage and the spillover effects just above it re-emerge. As with the overall sample, a specification that includes interactions with the value of the minimum wage shows relatively small differences in the proportional impact on the hazard of minimum wages at different levels.[29]

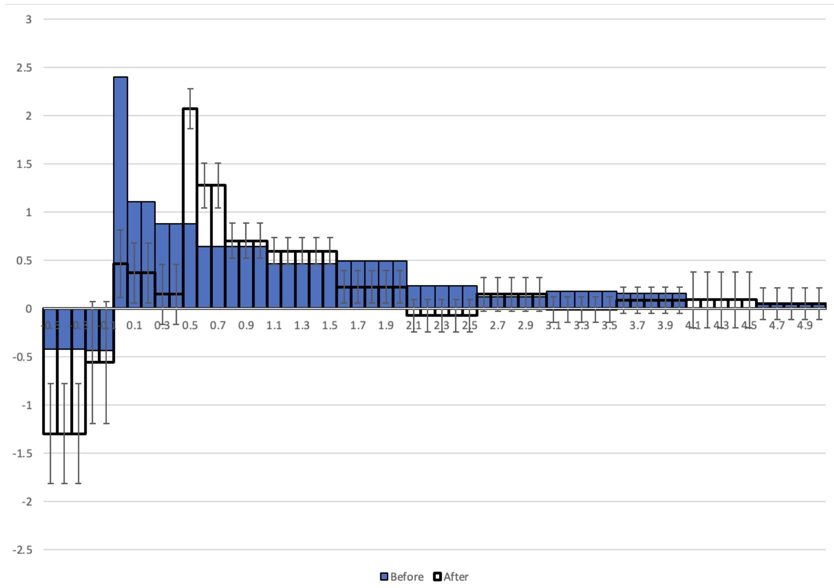

Figure 13: Effect of the minimum wage on wage bin hazards just after a 50 cents increase in the minimum wage (Joiners): Women

In Figure 14, we plot the minimum wage effects on the hazard just after a $0.50 increase in the minimum wage for female Stayers and Joiners on the same figure. The effects are quite similar in the decrease in mass below the new minimum wage. But at and just above the new minimum wage, there is more mass for Joiners than Stayers. That is, new hires are particularly likely to be hired near and just above the minimum wage. Since both Joiners and Stayers show limited changes in the shape of the minimum wage effect after versus before a minimum wage change, this is not a statement about the impact of a recent minimum wage increase but about the impact of minimum wages on Joiners even in steady state. Firms use the minimum wage more for new hires and then move workers to higher wages when they stay with the firm.

6.5 Results by Job Tenure

Differences between Joiners and Stayers raise the interesting possibility that the impact of a minimum wage increase could vary by job tenure. We examine this by re-estimating for Stayers with under 1 year of job tenure, 1 to 5 years of job tenure, and over 5 years of job tenure. In order to highlight the tenure related differences, we plot the estimated

[29]Our results for Joiners differ from those in Cengiz et al. (2019), who find no spillovers for this group. We obtain the same result when applying their estimator to our data (see Appendix C), implying that the difference stems from the different identification sources discussed in the Monte Carlo exercise.

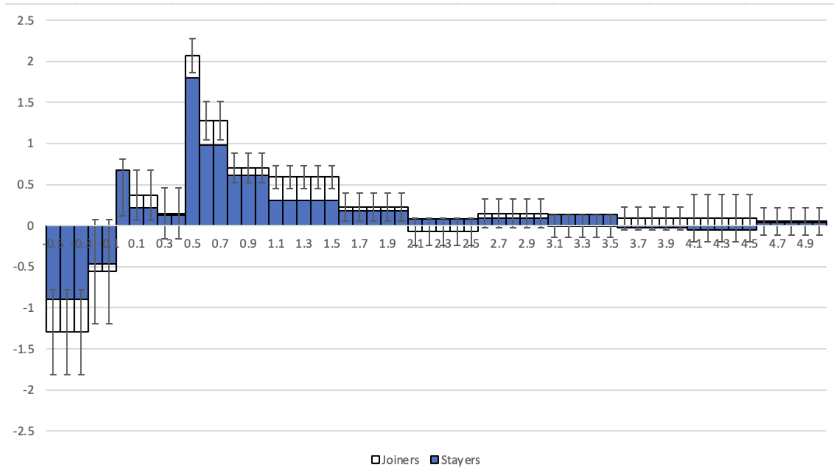

Figure 14: Effect of the minimum wage on wage bin hazards just after a 50 cents increase in the minimum wage (Joiners and Stayers): Women

effects for female workers with under 1 year of job tenure and workers with over 5 years of tenure in Figure 15. The results for are striking. The plot for workers with under 1 year of tenure, naturally, looks like the one for Joiners, with a substantial increase in the hazard at the minimum wage and spillovers reaching several dollars above the minimum wage. But for workers with over 5 years of tenure, the effect at the minimum wage is half that for new workers and the spillover effects are much smaller. For males (reported in the appendix), both the spike and spillovers are also reduced, though to a much smaller extent.

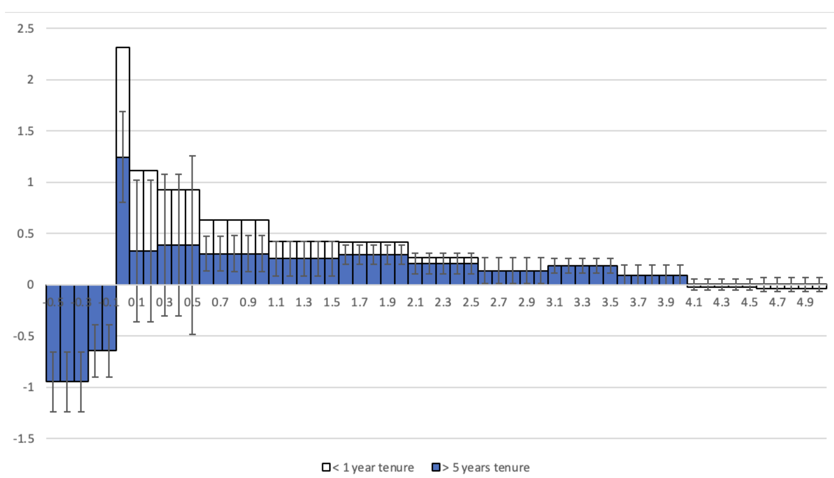

Figure 15: Effect of the Minimum Wage on Stayers by Job Tenure: Women

7 Using the Estimated Effects of Minimum Wages to Examine Models of the Labour Market

In this section, we use the key patterns identified in the previous section, treating them as core moments that models of the labour market ought to be able to match. To the extent models don't match all of them, the moments can be used as a guide to thinking about how to adjust or combine models. We start with a summary of the key empirical patterns we have identified and then move sequentially through some of the main models of the labour market.

7.1 Summary of Key Patterns

We view the empirical results in the previous section as pointing to a set of key patterns:

1. A spike at the minimum wage
2. Spillovers up to $2.50 above the minimum wage.
3. No impact of an increase in the minimum wage on selection out of jobs.
   3a Quits and Layoffs: The lack of impact is particularly evident for workers who are laid off. For workers who quit their job at the time of a new minimum wage increase, there is some evidence of extra mass at and just above the new minimum wage.
4. Changes in the wage distribution occur quickly - in less than 2 months for most responses.
5. Spike and spillover patterns for Stayers do not vary substantially with the level of the minimum wage.
6. Long tenure workers show smaller minimum wage effects. Spillovers are quite small for longer job tenure women.

7.2 Economic Models

7.2.1 Neoclassical Models

We start by considering the simplest neoclassical model: one in which the wage equals the value of the marginal product of a worker. The observed distribution of wages in this model comes from a fixed distribution of productivities across workers. The introduction of a minimum wage, m, should then price all workers with productivity below m out of the market, resulting in a truncated distribution of wages. There is no reason for a spike in the wage distribution at m nor for spillovers above m. In the Leavers distribution, we should see an extra mass of layoffs between the old and new minimum wages at the time of a minimum wage increase. Our estimated patterns effectively reject this model. This is a well-known result that might almost be elevated to the status of a folk theorem: the spike in the wage distribution at the minimum wage implies a rejection of the simple neoclassical model (Card and Krueger (1995)).

Teulings (2000) develops a more complete neoclassical model in which workers are characterized by skills, with a large number of types of skills, and shows that substitution across worker types will lead to an extra mass of workers being paid just above m. But there is still no reason for a spike at m and workers with wages below the new m before a minimum wage increase should still be disproportionately laid off (though to a lesser extent than in the simple model).[30] Again, the patterns that we estimate indicate a rejection of the model.

Haanwinckel (2024) generalizes Teulings' model by introducing monopsony power and rich heterogeneity in firm production functions. As in other monopsony models we discuss next, this richer model yields a spike at the minimum wage. But as in Teulings' model, the key force behind spillovers remains dis-employment effects among low productivity workers that increase the demand for slightly more productive workers. This prediction is at odds with our findings for Leavers.

7.2.2 Monopsony and Bargaining Models

The evident rejection of simple neoclassical models based on minimum wage impacts on employment and the wage distribution has been partly responsible for a surge in interest in models with imperfect competition in the labour market. In several key types of these models, the wage can be written as:

$$w(p, w^*) = (1 - \alpha)w^* + \alpha p, \tag{12}$$

[30] Brown (1988) makes the related point that if firms paid workers with varying abilities the minimum wage (in the period just after an increase), the spike would eventually unravel as firms compete for the more able workers by paying them just above the minimum wage in order to keep them from being poached. In contrast, we find that the minimum wage spike is permanent and is remarkably similar in the short and long runs for Stayers.

where p is the value of the worker marginal product in the firm; w^* is the value of the worker's next best alternative; and α is the share of surplus in the match going to the worker. Specifically, this is the form of the wage equation in both search and bargaining models (e.g., Flinn (2006)) and monopsony models (Manning (2003), Card et al. (2018), Lamadon et al. (2022)). In the bargaining models, α is a parameter reflecting worker bargaining power while in classical monopsony models, it is the inverse elasticity of the supply of labour to the firm. Whether w^* is worker specific or varies at the level of the market or a type of worker varies depending on the model.[31]

In contrast to neoclassical models, these models, when there is heterogeneity in p, can generate a spike at the minimum wage. To see this, consider the productivity threshold p^m, where $w(p^m, w^*) = m$, and m is the minimum wage. To make the exposition straightforward, assume that w^* is the same for all workers. p^m is then the productivity at which the wage determination process would generate a wage of m. For all firms with productivity above m (so that it is still profitable for them to operate) and below p^m, the minimum wage forces them to increase the wage they pay but they are still getting some positive surplus. Thus, all these matches will continue and will pay exactly the minimum wage. Some of the models can also generate a spillover effect on above-minimum wages. In classical monopsony models, this occurs when the minimum wage affects the elasticity of labour supply facing firms (i.e., α) and does so more for firms paying just above the minimum wage than for firms paying higher wages (Manning (2003)).[32] In bargaining models this may happen if a minimum wage improves the outside option of workers (implying an increase in w^*)(Flinn (2006)).

While these models can explain the broad shape pattern we observe, they face two key problems in matching our results. First, if the minimum wage increases from m to m', we would expect that there are some firms with $m < p \leq m'$, and those firms will shut down because they no longer generate a positive surplus to a match with a worker (Aaronson et al. (2018),Dustmann et al. (2022)). In these models, these should be the

[31] Lamadon et al. (2022) derive a specification that also includes a compensating differential term and tax structure parameters, but the general form of the markdown on productivity related to worker preferences over heterogeneous firm attributes continues to hold.

[32] In the Card et al. (2018) version of New Classical Monopsony models, there are no spillovers since neither w^* nor α vary across workers and firms. This is due to the convenient but simplifying assumption that workers choose among all possible jobs based on preferences that follow an extreme value distribution. In a more realistic version of the model where only employers in a certain segment of the labour market (e.g. around the minimum wage) compete for a given group of workers, employers that used to pay just above the minimum wage will need to compete more aggressively and raise wages when their close competitors increase their wages to comply with the minimum wage. This version of the model yields spillover effects, as in Manning (2003).

firms initially paying m, since it is the lowest productivity firms that pay exactly the minimum wage. Thus, we should see an increase in the size of the spike in the wage density for Leavers at the time of a minimum wage increase, but we do not observe this.

The second stems from the fact that all workers whose wages are initially between m and m' should be paid exactly m' after the minimum wage increase in these models. The hazard based model is particularly useful for investigating this implication. To see this, consider a situation in which the m increase does not cause any layoffs (i.e., $p > m'$ for all firms) and there is an underlying wage distribution in the absence of minimum wages given by $F(w)$. Then the spike at the initial minimum wage will equal $F(m)$ in these models, and the spike at the new minimum wage will equal $F(m')$. To relate this to our hazard estimator, consider a baseline hazard, $h_0(w)$ of unrestricted shape and assume that a minimum wage increases the hazard at m, so that the observed hazard at m is given by:

$$h_1(w) = h_0(w)e^{\beta_1} \tag{13}$$

where, $\beta_1 >> 0$ if $w = m$. Under our scenario with no layoffs, the hazard at m equals:

$$h_1(m) = \frac{F(m)}{(1 - F(m))} \tag{14}$$

Because the haazard in the absence of the minimum wage would equal $\frac{f(m)}{(1-F(m))}$,

$$h_1(m) = h_0(m)\frac{F(m)}{f(m)} \tag{15}$$

i.e., $e^{\beta_1} = \frac{F(m)}{f(m)}$. Importantly,

$$\frac{\partial \frac{F(m)}{f(m)}}{\partial m} = 1 - \frac{F(m)f'(m)}{f(m)^2} \tag{16}$$

Thus, the proportionality factor at m will be a function of m unless the percentage change in F equals the percentage change in f for all m. Since the percentage change in F must always be positive while the percentage change in f must sometimes turn negative, this is not possible. Moreover, in the normal and log-normal distribution, this derivative will be increasing in m. Our results from the specification with the minimum wage effects interacting with m show, instead, that we cannot reject that the equivalent of β_1 does not vary with m. The point estimate is positive but the implied effect is too small relative to the theoretical prediction from these models. Those predict that the spike should grow by $F(m') - F(m)$ (given that our Leavers results imply no added layoffs). In our example with m = \$6.75 and \$7.75, this amounts to an increase in the spike at m from 0.06 to 0.23 - much larger than the increase depicted in figure 8. Thus, our estimates do not fit with this key implication of frictional models.

Our estimator is useful for capturing this implication because the proportional form of the minimum wage effects focuses attention on the minimum wage effects themselves and away from the distraction of the height of the underlying density in the region where the minimum wage lies. More generally, the link function that corresponds to a proportional hazard model in a GLM representation ($\ln(1 - \ln((1 - h(w))))$) includes the hazard and, thus, the CDF evaluated at the wage point, w. That is, it naturally focuses attention on the location relative to the rest of the wage distribution, as captured in the CDF. In contrast, the Cengiz et al. (2019) estimator uses the employment level in a bin as its dependent variable and, so, does not use variation related to relative position. Since the implications concerning movements in the spike and the relative position arguments that we discuss in the next set of models relate to the position of an observation in the distribution, the hazard estimator is better suited for examining implications from these models.

Other models with imperfect labour markets also struggle to account for our findings for Stayers. Consider what Manning (2021) calls Modern Monopsony models, which are built on the Burdett and Mortensen (1998) job ladder model. In these models, there is a disperse wage distribution in equilibrium even when there is no heterogeneity in firm productivity through a mechanism based on firms poaching employed workers from other firms. Only the relative position of a firm in the offered wage distribution matters and spillover patterns above a minimum wage will get re-established in the same form when a minimum wage increases. But, as is well known, there are no spikes in the wage distribution in these models, including at the minimum wage. Engbom and Moser (2022) adjust the classic Burdett and Mortensen (1998) model by incorporating a set of workers who do not search on the job. Firms can identify who those workers are and, as a result, can pay them their reservation wage, inducing a spike at the minimum wage for workers with low reservation wages. The result is a model that would both predict a reassertion of a spillover pattern when a minimum wage increases and includes a spike. However, the spike should grow in size with minimum wage increases as a new set of workers with reservation wages between m and m' are moved right to the minimum wage. Again, we do not observe this in our Stayers estimates.

Based on these arguments, the standard imperfect labour market models do not fit easily with our estimated minimum wage impacts, though they come closer than the simple neoclassical models. In particular, what these models struggle with is the finding for Stayers that the spike and spillover pattern reasserts itself in the same form after a minimum wage increase. That is, our results appear to imply a wage structure anchored on the minimum wage. Perhaps more fundamentally, these models are not well suited for understanding the evolution of wages within firms. In particular, bargaining and job ladder models focus on wage setting at the beginning of an employment relationship.

7.2.3 Models with Fairness Considerations and Internal Labour Markets

The idea that fairness considerations could explain minimum wage spillovers was first proposed by Grossman (1983). The core concept is that workers provide more effort if they believe they are being treated fairly and their concept of fairness is based, in part, on how their wage relates to those of other workers in their firm. In the Akerlof and Yellen (1990) version of the model, workers' standard of wage fairness is written as a weighted average of the going wage in the market and the wages of others in their firm. A minimum wage then serves to move that standard for the lowest paid workers and then, as the wage of the lowest paid moves up, workers just above them require a wage increase to re-establish what they see as a fair wage premium. Since it is the comparison with the lowest paid workers that determines effort, firms will need to re-establish pay rankings relative to the minimum wage when it increases, which is what we observe. A spike would arise in this model if there is a set of workers whose ability is sufficiently low that, even including comparisons to others, their fair wage benchmark would be below the minimum wage. These workers would provide full effort when paid the minimum wage and, so, will all be paid at that level. This model could also explain the excess quits for workers who were paid at the level of the "new" minimum wage after a minimum wage increase. Those workers would have been paid more than the old minimum wage before the change and if their wage was not raised when the minimum wage increased then they could see that as unfair, inducing them to quit. This fits with Dube et al. (2019), who find that differential pay increases that were mechanically linked to the minimum wage at a large retail firm in the US caused workers with lower wage increases than identical colleagues to quit at a higher rate.[33] The importance of fairness considerations in the context of the minimum wage is also supported by experimental studies such as Falk et al. (2006) which find that workers' effort is related to the difference between their

[33]It also fits with evidence on effort and quit intentions in higher paid settings (e.g., Card et al. (2012) and Cullen et al. (2021)).

wage and the minimum wage. Because firms recognize this, in equilibrium, wage premia paid in order to induce effort will be paid relative to the minimum wage. This would generate a spillover pattern that moves up with the minimum wage, as we find.

Internal labour markets provide another potential mechanism for explaining our results. In this broad class of models, the internal wage structure is used as a way of incentivizing workers to provide more effort and to stay with the firm. In the classic Lazear and Rosen (1981) paper, for example, firms set up a contest for a promotion among workers in order to extract optimal effort. A key result in these models is that effort depends only on the spread between the wage workers get if they win the competition, W_2 and the wage they get if they lose, W_1, not on the wage levels (Malcomson (1984)). Where the levels of W_1 and W_2 are set then depends on assumptions about the market structure. If the firm is a monopsonist, it will set the wages to take all of the surplus. A minimum wage set above the losing wage then allows workers to capture some of the surplus. As in the classic monopsony models, the jobs will continue because the firm is still earning positive profits (even if they are smaller than without a minimum wage). That is, we would, again, observe a spike at the minimum wage. As the minimum wage increases, the spread between the lower paid wage and the wage paid after promotion to the tournament winner will be re-established in order to continue to extract maximum effort. Thus, this simple form of the model can explain the result that the spike and spillovers are anchored on the minimum wage and are reproduced when the minimum wage increases. It also provides a potential explanation for why the spike and spillover patterns are lower for high tenure workers since the tournament will already have been run for them and they are either at their higher wage or will have left the firm if they lost. On the other hand, these models do not provide an explanation for the spillover effects of the minimum wage that we observe for Joiners since the tournament will not yet have been run for new hires and, so, there is no reason for the firm to pay them different wages. In contrast, a Burdett-Mortensen type job ladder model can readily account for spillover effects in entry wages for Joiners as firms compete to maintain their position in the wage distribution.

Our main conclusion is that the patterns we estimate require some level of labour market imperfection and internal (to the firm) considerations stemming either from fairness or efficiency considerations. Our results are consistent with monopsonistic firms competing with each other when trying to hire workers, but internal considerations (fairness or internal labour markets) driving the evolution of wages within firms in response to minimum wage changes. Unlike a conventional job ladder model where workers are constantly searching for jobs, in a hybrid model where it is costly to change jobs, firms have incentives to set up an internal wage structure that incentivizes and helps retain workers. Such a hybrid model could also be viewed as an extension of Burdett and Coles (2003) where firms compete over contracts that offer both an entry wage and a tenure profile.

Finally, it is worth noting that all the models we examine have the feature that an increase in the minimum wage will cause a shut down in the least productive jobs/firms. That should show up as increased mass at the bottom end of the Leavers wage distribution - something we don't observe. This could reflect the relatively high frequency of our data variation or imply that we need always to be thinking about output prices and the ability of firms to pass wage costs along to consumers when examining wage impacts of minimum wages.

8 Conclusion

In recent years, our understanding of the effects of minimum wages has been substantially increased through the application of more sophisticated estimators to better data. In the US context, Cengiz et al. (2019) and Gopalan et al. (2021) have provided convincing estimates showing that minimum wage induce a sizeable spike in the wage density and spillover effects on wages up to about $2 above the minimum wage. As they show, these results have important implications for both the employment effects of minimum wages and their impact on inequality. Reaching back at least to the 1980s, the variation induced by minimum wages has also been used to differentiate among different possible models of the labour market (Card and Krueger (1995), Manning (2003)). Structural modelling of the labour market in papers such as Engbom and Moser (2022) and Haanwinckel (2024) have taken the implications of minimum wages for the shape of the overall wage distribution seriously.

Our contribution in this paper stems from our use of panel data to decompose the effects of a minimum wage change on the wage distribution into a selection effect (in the form of changes in who separates from jobs as measured by where they were in the wage distribution before separation) and a wage distribution shape effect (measured by changes in the shape of the wage distribution for people who remain on their job at the time of a minimum wage increase). We argue that what happens with these different components can help further differentiate among competing models of the labour market. Estimating these components amounts to estimating the effects of minimum wage changes on the whole shape of the wage distribution for different groups of workers. We accomplish this with a hazard function based estimator that allows for very flexible minimum wage effects and that we argue estimates how minimum wages change the shape of the underlying wage distribution, which differs from the employment change effects estimated in Cengiz et al. (2019) and Gopalan et al. (2021). Nonetheless, we similarly find evidence of a sizeable spike at the minimum wage and spillover effects up to about $2 above the minimum wage in the Canadian context.

Our first key finding is that there is no evidence of changes in selection - i.e., in where in the wage distribution job separators come from - when minimum wages increase. In particular, we do not find any evidence of increased separations for workers who were paid the pre-increase minimum wage or who had wages that place them between the old and new minimum wages. This is clearly a problematic finding for standard neoclassical models (as is the existence of the spike and spillovers) but it is also hard to reconcile with most imperfect labour market models.

Our second main finding is that there are wage distribution shape changes with an increase in the minimum wage. In particular, for people who remain on their job, we observe that the spike and spillover patterns that existed around the old minimum wage are re-established around the new, higher minimum wage in very much the same form and magnitude. This does not fit with the main imperfect labour market models which have the implication that people with wages at the old minimum wage and between the old and the new minimum wage should move directly to the new minimum wage (if they are not laid off). Instead, we argue that they fit best with models of wage setting within the firm such as job ladder and fairness based models. Even these models, though, have difficulty fitting some of our results - especially in terms of the impact of minimum wages on the wage distribution for new hires - and, in the end, we argue that the best model is likely one that combines models based on competition among firms for workers (Burdett and Mortensen (1998)) with models of wage setting within firms for incumbent workers.

We see our results as putting a new set of facts about minimum wages on the table to help in thinking about the best way to model the labour market. Our estimated effects for Job Stayers forces us to think about the internal workings of the employment relationship - something that perfect and imperfectly competitive models of the labour market do not really address. Our findings point to a need to broaden the imperfect competition models that have done well in explaining some key wage and employment patterns to incorporate considerations explored in earlier generations of models on efficiency wages and job ladders.

References

Aaronson, D., E. French, I. Sorkin, and T. To (2018). Industry dynamics and the minimum wage: A putty-clay approach. *International Economic Review 59*(1), 51–84.

Akerlof, G. A. and J. Yellen (1990). The fair wage-effort hypothesis and unemployment. *Quarterly Journal of Economics 105*(2), 255–283.

Alvarez, J., F. Benguria, N. Engbom, and C. Moser (2018). Firms and the decline of earnings inequality in Brazil. *American Economic Journal: Macroeconomics 10*(1), 149–189.

Autor, D., A. Manning, and C. Smith (2016). The contribution of the minimum wage to US wage inequality over three decades: A reassessment. *American Economic Journal: Applied Economics 8*(1), 58–99.

Brochu, P. (2013). The sources of the new Canadian job stability patterns. *Canadian Journal of Economics 46*(2), 412–440.

Brochu, P. (2021). A researchers guide to the labour force survey: Its evolution and the choice of public versus master files. *Canadian Public Policy 47*(3), 335–357.

Brochu, P. and D. A. Green (2013). The impact of minimum wages on labour market transitions. *The Economic Journal 123*(3), 1203–1235.

Brown, C. (1988). Minimum wage laws: Are they overrated? *Journal of Economic Perspective 2*(3), 133–145.

Burdett, K. and M. Coles (2003). Equilibrium wage-tenure packages. *Econometrica 71*(5), 1377–1404.

Burdett, K. and D. T. Mortensen (1998). Wage differentials, employer size, and unemployment. *International Economic Review 39*(2), 257–273.

Card, D., J. Heining, A. R. Cardoso, and P. Kline (2018). Firms and labor market inequality: Evidence and some theory. *Journal of Labor Economics 36*(S1), S13–S70.

Card, D. and A. B. Krueger (1995). *Myth and Measurement: The New Economics of the Minimum Wage*. Princeton University Press.

Card, D., A. Mas, E. Moretti, and E. Saez (2012). Inequality at work: The effect of peer salaries on job satisfaction. *American Economic Review 102*(6), 2981–3003.

Cengiz, D., A. Dube, A. Lidner, and B. Zipperer (2019). The effect of minimum wages on low wage jobs. *Quarterly Journal of Economics 134*(3), 1405–1454.

Cullen, Z., S. Li, and R. Perez-Truglia (2021). What's my employee worth? the effects of salary benchmarking. Technical Report 30570, National Bureau of Economic Research.

Dickens, R. and A. Manning (2004a). Has the national minimum wage reduced UK wage inequality? *Journal of the Royal Statistical Society, Series A 167*(1), 613–626.

Dickens, R. and A. Manning (2004b). Spikes and spill-overs: The impact of the national minimum wage on the wage distribution in a low wage sector. *The Economic Journal 114*(494), 95–101.

Dickens, R., A. Manning, and T. Butcher (2012). Minimum wages and wage inequality: Some theory and an application to the UK. Technical Report 1177, Centre for Economic Performance.

DiNardo, J., N. M. Fortin, and T. Lemieux (1996). Labor market institutions and the distribution of wages, 1976-1992: A semiparametric approach. *Econometrica 64*(5), 1001–1044.

Donald, S. G., D. A. Green, and H. J. Paarsch (2000). Differences in wage distributions between Canada and the United States: An application of a flexible estimator of distribution functions in the presence of covariates. *Review of Economic Studies 67*(4), 609–633.

Dube, A., L. Guiliano, and J. Leonard (2019). Fairness and frictions: The impact of unequal pay raises on quit behavior. *American Economic Review 109*(2), 609–633.

Dube, A., A. Manning, and S. Naidu (2020). Monopsony and employer mis-optimization explain why wages bunch at round numbers. Technical Report 24991, National Bureau of Economic Research.

Dustmann, C., A. Lindner, U. Schoenberg, M. Umkehrer, and P. von Berge (2022). Reallocation effects of the minimum wage. *Quarterly Journal Labor Economics 137*(1), 267–328.

Engbom, N. and C. Moser (2022). Earnings inequality and the minimum wage: Evidence from Brazil. *American Economic Review 112*(12), 3803–3847.

Falk, A., E. Fehr, and C. Zehnder (2006). Fairness perceptions and reservation wages - the behavioral effects of minimum wage laws. *Quarterly Journal of Economics 121*(4), 1347–1381.

Flinn, C. J. (2006). Minimum wage effects on labor market outcomes under search, matching, and endogenous contact rates. *Econometrica 74*(4), 1013–1062.

Fortin, N. M. and T. Lemieux (2015). Changes in wage inequality in Canada: An interprovincial perspective. *Canadian Journal of Economics 48*(2), 682–713.

Fortin, N. M., T. Lemieux, and N. Lloyd (2021). Labor market institutions and the distribution of wages: The role of spillover effects. *Journal of Labor Economics 39*(S2), S369–S412.

Gopalan, R., B. Hamilton, A. Kalda, and D. Sovich (2021). State minimum wages, employment, and wage spillovers: Evidence from administrative payroll data. *Journal of Labor Economics 39*(3), 673–707.

Green, D. A., R. Morissette, and B. M. Sand (2019). Economy wide spillovers from booms: Long distance commuting and the spread of wage effects. *Journal of Labor Economics 37*, S642–S687.

Grossman, J. B. (1983). The impact of the minimum wage on other wages. *Journal of Human Resources 18*(3), 359–378.

Haanwinckel, D. (2024). Supply, demand, institutions, and firms: A theory of labor market sorting and the wage distribution. Technical report, University of California Los Angeles.

Jales, H. (2017). Estimating the effects of the minimum wage in a developing country: A density discontinuity approach. *Journal of Applied Econometrics 33*(1), 29–51.

Lamadon, T., M. Mogstad, and B. Setzler (2022). Imperfect competition, compensating differentials, and rent sharing in the US labor market. *American Economic Review 112*(1), 169–212.

Lazear, E. P. and S. Rosen (1981). Rank-order tournaments as optimum labor contracts. *Journal of Political Economy 89*(5), 841–864.

Lee, D. S. (1999). Wage inequality in the United States during the 1980s: Rising dispersion or falling minimum wage. *Quarterly Journal of Economics 114*(3), 977–1023.

Malcomson, J. (1984). Work incentives, hierarchy, and internal labor markets. *Journal of Political Economy 92*(3), 486–507.

Manning, A. (2003). *Monopsony in Motion: Imperfect Competition in Labor Markets.* Princeton University Press.

Manning, A. (2021). Monopsony in labor markets: A review. *ILR Review 74*(1), 3–26.

Meyer, B. D. (1990). Unemployment insurance and unemployment spells. *Econometrica 58*(4), 757–782.

Meyer, R. H. and D. A. Wise (1983). Discontinuous distributions and missing persons: The minimum wage and unemployed youth. *Econometrica 51*(6), 1677–1698.

Neumark, D., M. Schweitzer, and W. Wascher (2004). Minimum wag effects throughout the wage distribution. *Journal of Human Resources 39*(2), 425–450.

Ryu, K. (1992). Group duration analysis of the proportional hazard estimator: Minimum chi-squared estimators and specification tests. Technical report, University of California Los Angeles.

Sorkin, I. (2015). Are there long-run effects of the minimum wage? *Review of Economic Dynamics 18*(2), 306–333.

Stewart, M. B. (2012). Quantile estimates of counterfactual distribution shifts and the effect of minimum wage increases on the wage distribution. *Journal of the Royal Statistical Society, A 175*(1), 263–287.

Teulings, C. N. (2000). Aggregation bias in elasticities of substitution and the minimum wage paradox. *International Economic Review 41*(2), 359–398.

A On-line Appendix: Other Estimation Results

A.1 Estimated Effects for Forward Stayers

Recall that given our data do not allow us to follow the wages of job Stayers, we generate samples of Forward Stayers (workers who just joined the sample and who will be on the same job two months later) and Backward Stayers (workers who just joined the sample and have job tenure of at least 2 months). We use the Backward Stayers in our estimates to see if minimum wage changes alter the wages of Stayers. Forward Stayers are useful for seeing whether firms alter their wage policies in anticipation of a minimum wage change. In particular, we estimate our model for them in the same way as for Leavers: allowing for effects from a minimum wage change that will occur in the next two months.

Figures A1 and A2 present the estimated effects on the hazard when there is no minimum wage change and when there is an upcoming minimum wage of 50 cents for women and men, respectively. The plots show minimal changes in the minimum wage effects when there is an upcoming minimum wage change, implying limited anticipatory changes in wage setting by firms.

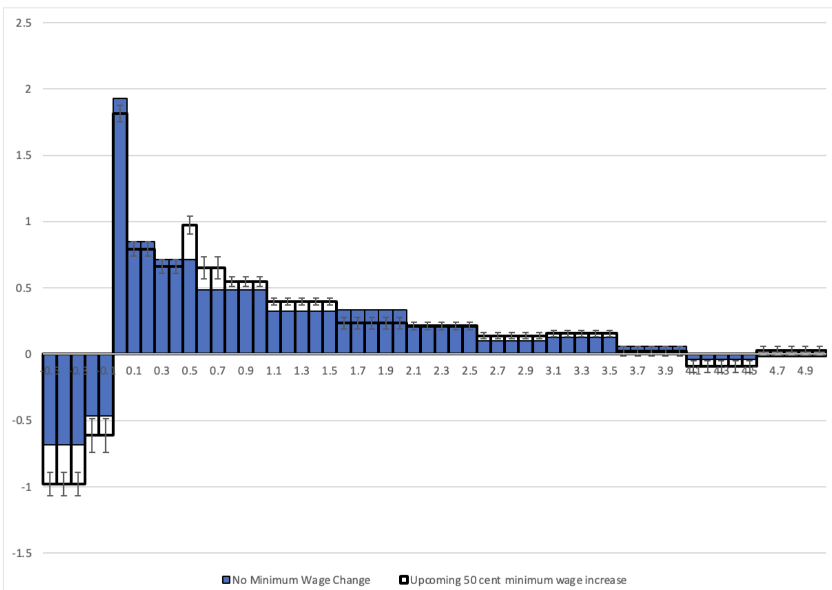

Figure A1: Effect of the minimum wage on Stayers with and without an upcoming 50 cent increase in the minimum wage: Women

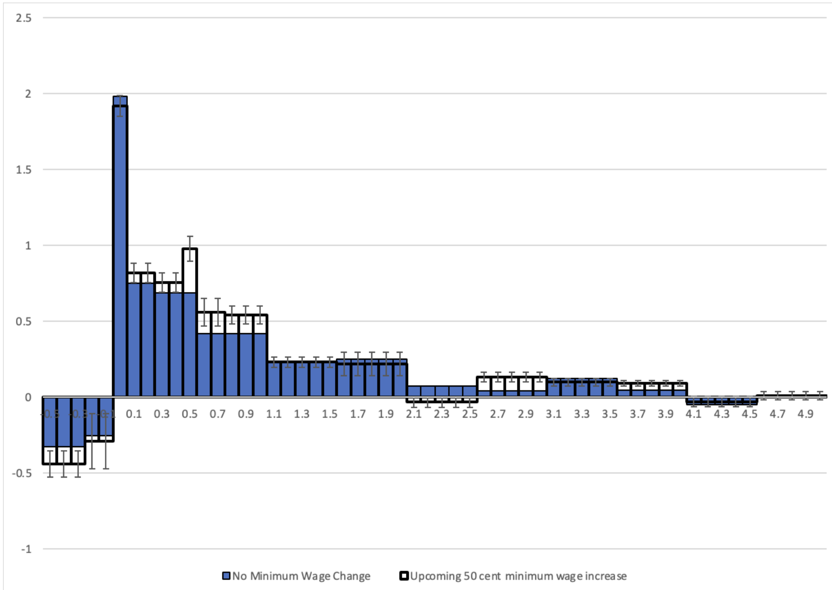

Figure A2: Effect of the minimum wage on Stayers with and without an upcoming 50 cent increase in the minimum wage: Men

A.2 Estimates with \$5 and \$7 Cut-offs

In this section, we present the estimated minimum wage hazard effects using the \$5 above m cut-off (the one used for our main results) and a \$7 above m cut-off. Figure A3 contains the estimated effects when there is no upcoming minimum wage change for females. The estimates using either the \$5 or \$7 cut-offs are remarkably similar. Both show small and statistically insignificant effects above about \$2 above m and the estimated spike and spillover effects up to that point are very close to one another (and each lies well within the 95% confidence interval for the other). The estimates below m show more differences but the basic picture of large negative effects that are essentially flat across the below m range is the same. The same picture emerges in Figure A4, for men, with the differences below m being smaller.

Most importantly for our estimation approach, the estimated effects above \$5 above m that we obtain when we use a cut-off of \$7 above m are all essentially zero and statistically insignificant. We conclude from this that our assumption that there are no effects of the minimum wage on wages above \$5 (the assumption we require for the main estimates in the text) is very likely to hold. We cannot, of course, prove that there is not some wage range higher up where the minimum wage has an effect (even if we set the cut-off at \$50 above m, that doesn't prove that there aren't effects still higher), but we view these results as highly suggestive. It is worth highlighting that our estimator includes very flexible time and location effects that vary by wage ranges and, thus, we believe our estimated minimum wage effects are unlikely to include spurious associations with other events.

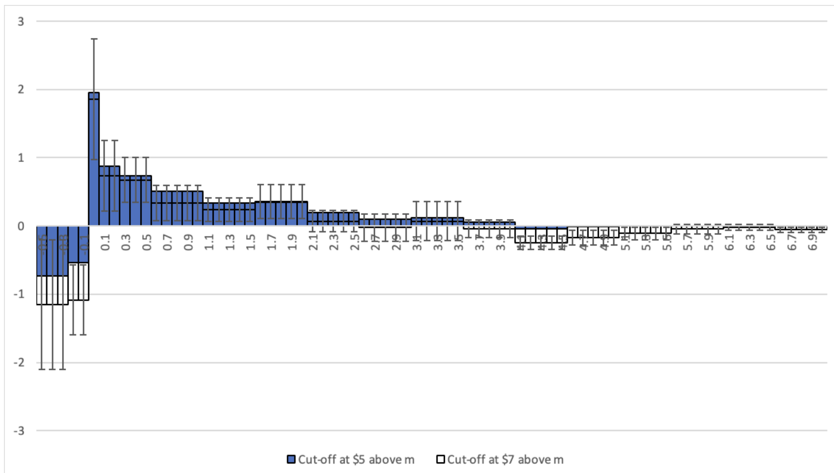

Figure A3: Effect of the minimum wage on all workers with no upcoming increase in the minimum wage, with different cut-offs: Women

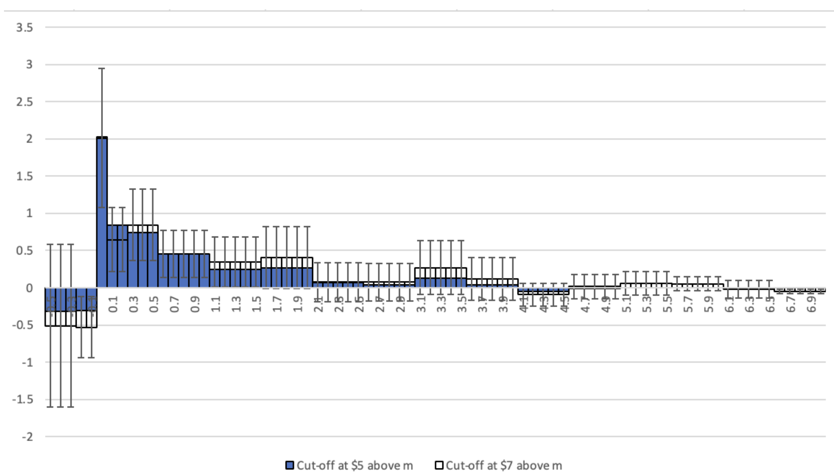

Figure A4: Effect of the minimum wage on all workers with no upcoming increase in the minimum wage, with different cut-offs: Men

A.3 Results for Males

In this section, we present the plotted minimum wage effects for males to match the figures for females in the text.

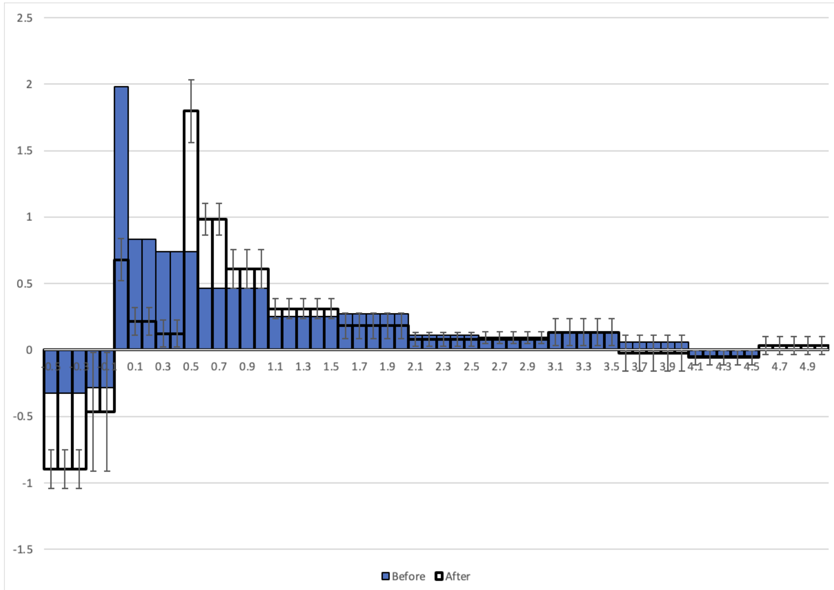

Figure A5: Effect of the minimum wage on wage bin hazards just after a 50 cents increase in the minimum wage (Stayers): Men

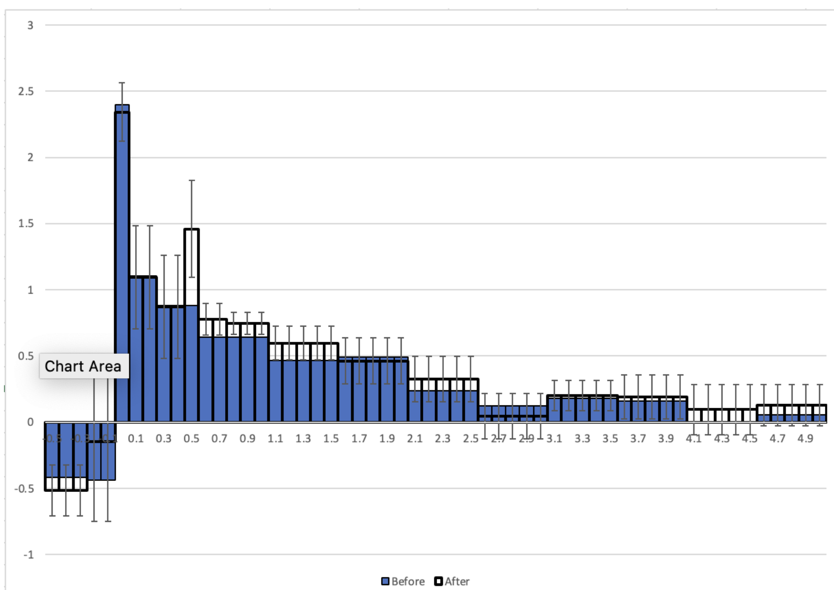

Figure A6: Effect of the minimum wage on wage bin hazards just after a 50 cents increase in the minimum wage (Layoffs): Men

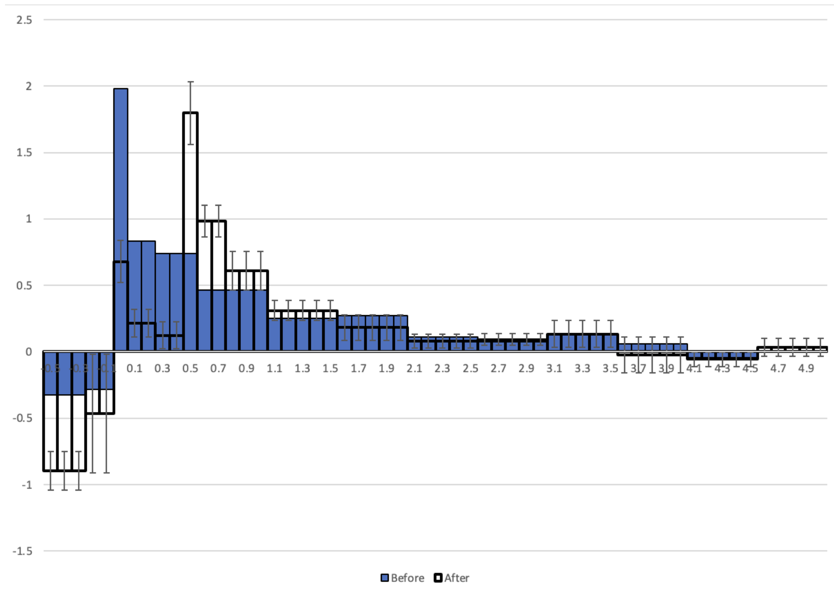

Figure A7: Effect of the minimum wage on wage bin hazards just after a 50 cents increase in the minimum wage (Joiners): Men

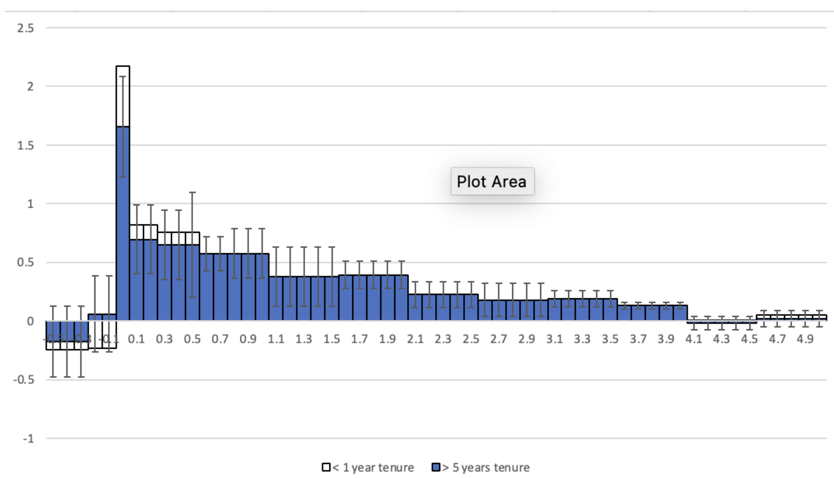

Figure A8: Effect of the Minimum Wage on Stayers by Job Tenure: Men

B On-line Appendix: The Hazard Estimator and the Mechanical Channel

In this Appendix, we show that a hazard estimator allows us to eliminate the effects of the third, mechanical, channel when estimating minimum wage effects - at least, in parts of the distribution that are not directly affected by minimum wage changes through the other two channels. The re-scaling effects of the third channel can serve to obscure the shape and selection effects of minimum wage changes.

Consider a situation with two periods, 0 and 1, and two provinces: NM (which has no minimum wage change between periods 0 and 1) and M (which has a minimum wage increase between 0 and 1). Assume that the two provinces are identical in period 0 in terms of the wage CDF ($F_0^j(w)$, j=NM, M) and its component elements ($N_{L0}^j, N_S^j, F_{S0}^j(w), F_{L0}^j(w), E_0^j$). Assume, further, that the minimum wage has no impact on the component CDF's above a wage level, w^*. Expressed in terms of some wage interval $[w_1, w_2]$, where $w_1 > w^*$, this assumption can be written as:

$$N_{S1}^M[F_{S1}^M(w_2) - F_{S1}^M(w_1)] = N_{S1}^{NM}[F_{S1}^{NM}(w_2) - F_{S1}^{NM}(w_1)] \tag{17}$$

$$N_{J1}^M[F_{J1}^M(w_2) - F_{J1}^M(w_1)] = N_{J1}^{NM}[F_{J1}^{NM}(w_2) - F_{J1}^{NM}(w_1)] \tag{18}$$

$$N_{L0}^M[F_{L0}^M(w_2) - F_{L0}^M(w_1)] = N_{L0}^{NM}[F_{L0}^{NM}(w_2) - F_{L0}^{NM}(w_1)] \tag{19}$$

The first of these expressions says that there is no shape change in the $[w_1, w_2]$ interval induced by the minimum wage increase. The second and third expressions say there is also no selection effect in this range. In all three cases, the assumption is expressed in terms of the number of workers observed in the given wage range. Writing the assumption in terms of the number of workers addresses the issue that the component distributions in the M province (e.g., $F_{S1}^M(w_2)$) have to be re-scaled relative to the matching NM distribution if the number of workers in the category (N_S in this case) is changed by the minimum wage increase.

First, consider the proportion of workers in the $[w_1, w_2]$ in each province in period 1:

$$F_1^M(w_2) - F_1^M(w_1) = \frac{1}{E_1^M}[N_S^M(F_{S1}^M(w_2) - F_{S1}^M(w_1)) \tag{20}$$

$$+N_{J1}^M(F_{J1}^M(w_2) - F_{J1}^M(w_1))]$$

and,

$$F_1^{NM}(w_2) - F_1^{NM}(w_1) = \frac{1}{E_1^{NM}}[N_S^{NM}(F_{S1}^{NM}(w_2) - F_{S1}^{NM}(w_1)) \tag{21}$$

$$+N_{J1}^{NM}(F_{J1}^{NM}(w_2) - F_{J1}^{NM}(w_1))]$$

Under the no-minimum-wage effects on shape and selection assumption above, the right hand sides of equations 20 and 21 are equal except for the scaling factors, $\frac{1}{E_1^M}$ and $\frac{1}{E_1^{NM}}$. These scaling factors mean that the observed probability distribution function in this interval will change with the change in the minimum wage even though there was no direct economic effect of the minimum wage change in this part of the wage range. This is equivalent to what we described in the simple truncation case in the text.

Now, consider the hazard rate for this wage range for both provinces. The denominators in the two cases are given by:

$$1 - F_1^M(w_1) = \frac{1}{E_1^M}[N_S^M(1 - F_{S1}^M(w_1)) \tag{22}$$

$$+N_{J1}^M(1 - F_{J1}^M(w_1))]$$

and,

$$1 - F_1^{NM}(w_1) = \frac{1}{E_1^{NM}}[N_S^{NM}(1 - F_{S1}^{NM}(w_1)) \tag{23}$$

$$+N_{J1}^{NM}(1 - F_{J1}^{NM}(w_1))]$$

The hazard rate for M is given by the ratio of (20) to (22) and the hazard rate for NM is given by the ratio of (21) to (23). Note that the scaling factors $\frac{1}{E_1^M}$ and $\frac{1}{E_1^{NM}}$ cancel out, leaving terms that are equivalent according to our assumption that the minimum wage has no scale or selection effects in this range. Thus, while the pdf would show scaling effects that might be mistaken for spillovers from the minimum wage into the $[w_1, w_2]$, the hazard rate would show that there are no selection or shape channel effects.

C On-line Appendix: Monte Carlo Results with the Quantile Estimator

In this appendix, we demonstrate the extent and nature of the re-scaling issue for the quantile estimator. Recall that the re-scaling issues arises if minimum wage increases induce lay-offs since the wage density need to be re-scaled by 1 minus the proportion laid-off to insure that the density continues to integrate to 1, and that implies movements in the quantiles.

To demonstrate the nature of the problems for quantile estimators that arise from the re-scaling issue, in figure A9, we show the Monte Carlo exercise for Scenario 1 with p_0=0.5 and m=\$12 in the left panel. The solid line represents the underlying density while the dashed line corresponds to the proportion of workers at each wage who remain employed after imposing the minimum wage. Note that the latter line coincides with the underlying density above m. The third line (with the symbols) is the density after introducing the

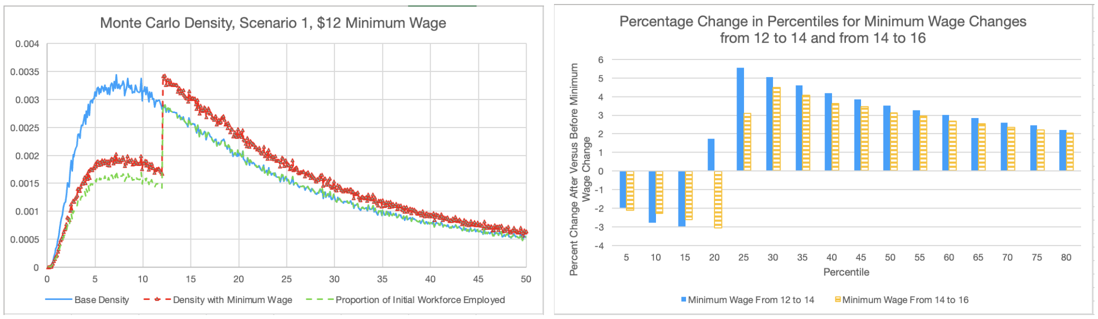

Figure A9: Scenario 1, Monte Carlo Densities and Quantile Estimates

minimum wage and corresponds to the proportion employed line re-scaled to integrate to 1. The right panel shows quantile estimates of the impact of a minimum wage change from \$12 to \$14 (solid bars) and from \$14 to \$16 (striped bars) in 5 percentile wide bins. These are simply the percentage change in the location of each percentile bin after versus before the minimum wage increase. Notice that the proportion of workers who remain working at each wage below \$12 does not change (it is always p_0) but the total proportion laid off rises when we shift m from \$12 to \$14 and that results in a different re-scaling for the densities at each value of m. The re-scaling means that the percentiles for wages below \$12 shift to the left when m increases to \$14. The new lay-offs in the range between \$12 and \$14 imply that the percentiles in that range shift to the right and this, combined with the rescaling, implies that all the higher percentiles also shift right at a declining rate. The pattern is somewhat different in shape for the minimum wage change from \$14 to \$16. In each case, the pattern is heavily determined by the re-scaling that is evident in the left panel. From the quantile estimates, those re-scaling effects might lead us to conclude (incorrectly) that there are reductions in mass below \$12 and spillovers throughout the above-m range even though neither are true in Scenario 1. The same holds in Scenario 2 but in Scenario 3 (with no layoffs and, therefore, no re-scaling needed), the quantile estimator provides an accurate picture of the change in the density when the minimum wage changes.

In addition to the re-scaling problems, with multiple jurisdictions with different minimum wages, what we will observe at any given percentile could be an average of the direct minimum wage (if the percentile coincides with the minimum wage in some states) and spillover effects (if it is above the minimum wage in others). Working directly with percentiles makes it very difficult to assess whether minimum wages generate spillovers and what changes for wage setting below the minimum wage. For both reasons, quantile estimators are not well suited to estimating the impacts of minimum wages on the shape of the wage distribution.

D On-line Appendix: Data Linkage Details

In this appendix, we describe the process by which we link individuals across Labour Force Survey (LFS) months. The LFS has a rotating panel design in which households remain in the sample for six consecutive months, with 1/6 of the sample replaced each month. Attempts to create a panel from waves of the LFS are complicated by the fact that it follows dwellings rather than individuals. Individuals who change dwellings exit the survey. Fortunately, we can uniquely identify individuals across monthly files with a combination of LFS variables. Changes over time in geographical identifiers (e.g., EI regions) mean that different identifying variables must be used for different periods but since our sample period begins in 1996, we only need to use one set of consistent variables (Brochu (2021)). In our description of the procedure for identifying individuals and linking them across months that follows we name variables (in capital letters) using the same names provided in the LFS codebook.

We linked observations across months using the month, the one-digit province code (PROV1), the pseudo UIC regions (PSEUDOUI), the regional strata (FRAME), the superstratum (STRAFRAM), the sample design type (TYPE), the first-stage sampling unit (CLUST), the rotation number (ROTATION), the number assigned to dwellings within a cluster (LISTLINE), the multiple dwelling code for structures that have more than one dwelling (MULT) and the LINE variables. After we created our data, the LFS has added a new variable, HHLDID, which makes the linkage much more straightforward.[34]

We dropped individuals that had incompatible tenure spells across the two periods of the panel. That is, for an individual that worked in period 1, she must have one more month of tenure in period 2 (i.e. continued with the same employer), one month of tenure in period 2 (i.e. started a new job), or no job tenure in period 2 (i.e. is out of work). We also dropped individuals who were full time students at any of their six interviews. This means we do not include students on summer jobs between their school years. Finally, as described in the text, we do not include self-employment spells and, so, we dropped

[34]See Brochu (2021) for a complete discussion of linkage issues in different periods and other issues with LFS data.

transitions into self-employment.

E On-line Appendix: Details of the Event Study Implementation

E.1 Event Study

In order to present the key features of the wage data in as transparent a way as possible, we present, here, an event study of the impact of larger minimum wage changes on percentiles of the wage distribution. More specifically, we run a set of regressions of the form:

$$y_{qpt} = \eta_p + \psi_t + \sum_{\tau=-8}^{8} D_{\tau pt}\theta_\tau + \epsilon_{qpt} \tag{24}$$

where y_{qpt} is the value of the q^{th} percentile of the real log wage distribution in province p in quarter t, η_p is a complete set of province effects, ψ_t is a complete set of time effects, $D_{\tau pt}$ is a dummy variable that equals 1 if quarter t in province p corresponds to period τ before or after a minimum wage change, and θ_τ are the set of estimated coefficients of interest to us. Monthly data are collapsed at the quarter level to improve precision, with the quarterly minimum wage defined as the maximum value of the minimum wage prevailing in the quarter. As such, we are likely understating the impact of a minimum wage change in the quarter when the minimum wage first increases ($\tau = 0$), since the new minimum wage doesn't typically prevail in all months of the quarter.

We define the relevant events as changes in the nominal minimum wage that are 5% or greater in size, treating periods with smaller or no nominal minimum wage changes as the control periods.[35] We allow for the minimum wage changes to have effects starting as early as 8 quarters before the actual change to test for possible pre-trends linked to anticipatory changes or endogenous changes in the minimum wage that may be more likely to take place when the local economy is doing well and real wages are growing fast. We also look at wage effects as late as 8 quarters after that minimum wage change and leave out the dummy for period $\tau = -1$ so that the θ_τ's correspond to effects relative to this pre-event benchmark. We provide more details of our implementation of the event study in the following subsection.

Figure A10 shows the plots of the estimated θ_τ coefficients for the 5th percentile for

[35] Adding a set of dummies to control for small changes in the minimum wage has little impact on the results and slightly reduces the precision of the main estimates.

women. There is no evidence of a pre-trend, as the estimated values of θ_τ's prior to the minimum wage increase are closely scattered around 0. None of these estimated coefficients are significantly different from zero at conventional significance levels. After the event, the 5th percentile rises, quickly reaching a level of approximately 4% after a few quarters. As conjectured above, the estimated effect is smaller in the quarter when the minimum wage increase takes place ($\tau = 0$) since the new minimum wage may not be prevailing for the entire duration of the quarter. Note also that there is a slightly upward trend in the estimated effect in the second year after the minimum wage increase. It is unclear whether this reflects growing dynamic impacts of the minimum wage, or some unmodelled trend growth in wages.

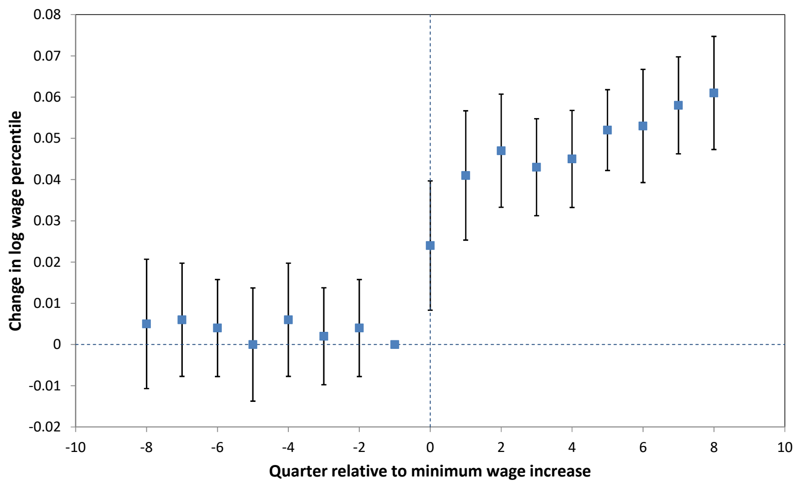

Figure A10: Event Study for women, log 5th percentile

For the 10th percentile (Figure A11), there is again little in the way of a pre-trend. With the exception of the quarter when the minimum wage takes place, the estimated effects are all positive and statistically significant after the increase in the minimum wage. Although the estimates are smaller than for the 5th percentile, there is again a clear trend in the size of the effects over time. By the 15th percentile (Figure A12), the post-event effects are smaller still, and only become statistically significant in the fourth quarter after the minimum wage increase. Estimation for the 20th percentile (not shown here) reveals no significant post-event effects. For men (again, not shown here, for brevity) the estimated effects are smaller (often less than half those for women) and are not statistically significant past the 10th percentile.

These results show that there is evidence of effects of minimum wage changes on the wage distribution in our data, especially for women. Focusing on percentile movements

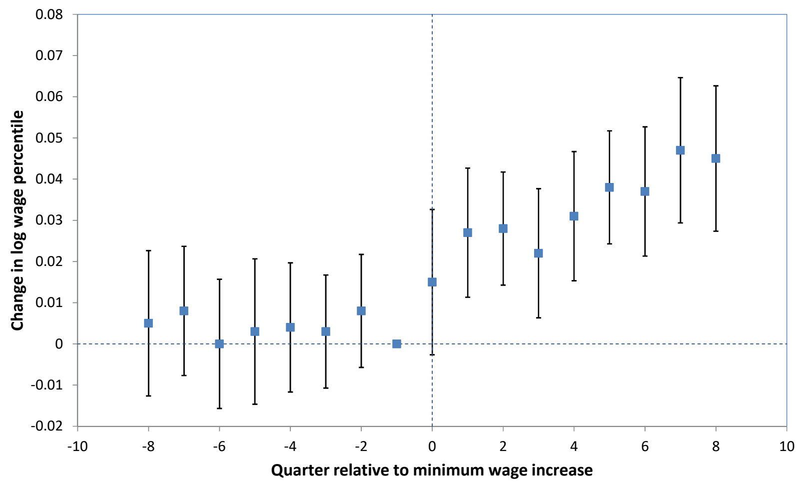

Figure A11: Event study for women,log 10th percentile

means that we cannot be precise about the extent of spillovers versus direct effects of the minimum wage change for the 5th and, for women, even the 10th percentiles since the minimum wage falls in this part of the distribution in much of our sample period.[36] The fact that there are some movements at the 15th percentile for women and the 10th percentile for men is suggestive of spillover effects but we need to turn to our main estimator in order to get direct estimates of shape effects of minimum wage changes.

E.2 Implementation Details

In this section, we present extra implementation details for the event study exercise.

One complication faced in running the event study is that there is overlap in the event-study windows when two large changes in the minimum wage take place less then 4 years apart. One possible approach for dealing with this issue is to restrict the analysis to minimum wage changes that take place more than four years away from each other. This would, unfortunately, result in a sharp drop in the number of large minimum wage changes to be used in the event-study analysis. A more efficient approach (and the one we use here) is to add a set of "treatment window dummies" to control for the possible overlap in the event-study windows.

In the case where the minimum wage only changes once, say in January 2007, we in-

[36]By direct effect we mean that when a wage percentile –say the 5th percentile– is equal to the minimum wage, there is a one-for-one relationship between the minimum wage and the wage percentile even in absence of spillover effects.

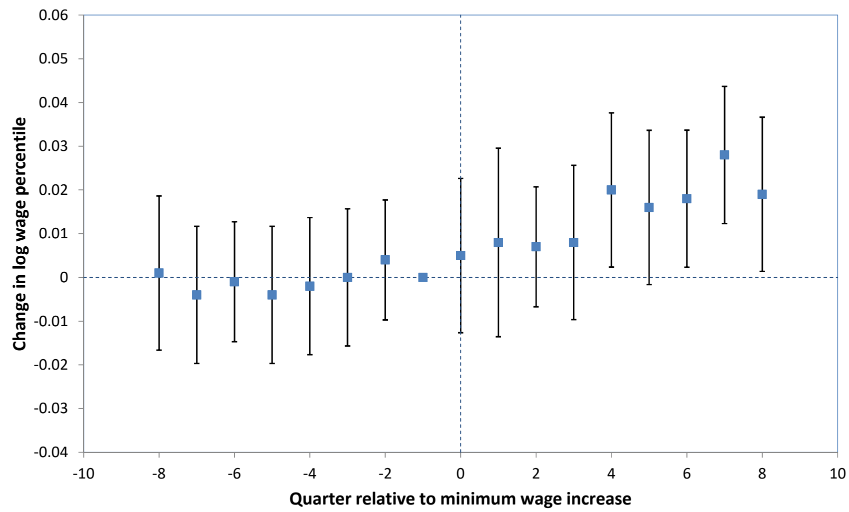

Figure A12: Event study for women, log 15th percentile

clude dummies for three province-specific treatment windows: a "pre-treatment" window for the 1997-2004 period; a "treatment" window going from January 2005 to December 2008; a "post-treatment" window going from 2009 to 2016. In this simple example, the "trigger" for creating a new window is hitting 8 quarters before or after the minimum wage change.

We use the same "triggering" approach in more general cases with several minimum wage changes. For example, consider the case where there is an additional increase in the minimum wage in January 2010. We now create five window treatment dummies. The pre-treatment window is still 1997-2004, and the post-treatment window is now 2012 to 2016. We also have three other windows: a pure first-change window going from January 2005 (8 quarters before the first change) to December 2007 (4 quarters after the first change); a "mixed" window in 2008 that is within 8 quarters of both the first and second minimum wage changes; a pure second-change window going from January 2009 (4 quarters before the second change) to December 2011 (8 quarters after the second change).

Doing so insures that the event-study coefficients are identified using the variation within the time windows (the first- and second-change windows in the above example) where the impact of a minimum wage change is not confounded by previous or upcoming changes in the minimum wage.

F On-Line Appendix: Cengiz et al. (2019) Estimator Applied to Canadian Data

In this section, we present the results of applying the Cengiz et al. estimator to our Canadian data. We use their baseline specification

$$\frac{E_{pjt}}{N_{pt}} = \sum_{\tau=-3}^{4} \sum_{k=-4}^{17} \alpha_{\tau k} I_{pjt}^{\tau k} + \mu_{pj} + \Omega_{pjt} + u_{pjt}, \tag{25}$$

where E_{pjt} is the employment count in 0.25 wage bin j in province p in quarter t and N_{pt} is the adult population in the same quarter and province. The treatment dummies, $I_{pjt}^{\tau k}$ are a series of dummies that equal one if the minimum wage was increased τ years from the current quarter for those wage bins that lie between k and $k+1$ dollars relative to the new minimum wage. $\tau = 0$ indicates the quarter in which the minimum wage was increased, as well as the three quarters following the increase. Any minimum wage spike will be included in the coefficient α_τ corresponding to $k = 0$. In addition, μ_{pj} are province-bin fixed effects while Ω_{pjt} are a set of controls for small minimum wage increases[37] As with Cengiz et al., we cluster our standard errors at the level of the province.

Following Cengiz et al., we report our results graphically by calculating and plotting the five-year post-treatment average employment change, $\frac{1}{5}\sum_{\tau=0}^{4} \alpha_\tau$. Following Cengiz et al., these are normalized by dividing the change in employment in each bin by the average overall pre-treatment employment rate.[38] The results are shown in Figure A13, which is our replication of Figure 2 of Cengiz et al. using the same Canadian data and study period that we use in our hazard function model. Although the wage bins and the minimum wage effects are based on wider bins than we use in our analysis, and the dynamics are of a longer duration than we use, the results are broadly similar. The Cengiz et el. estimator also shows a decrease in the employment share below the new minimum wage and spill-overs up to \$2 above the minimum wage. The spike from \$0 to \$1 above the minimum wage includes our spike at the minimum wage as well as the spillover effects in the three ranges above (i.e. \$0.10 to \$.30, \$.30 to \$.50 and \$.50 to \$1.00). This differs somewhat from Cengiz et al.'s own results using American data, as they found positive spill-overs up to three dollars above the minimum wage.

We also estimated our model separately for Stayers and Joiners. This corresponds

[37]See footnote 4 of Cengiz et al., 2019, for additional details.

[38]The average pre-treatment employment rate over our study period is .642.

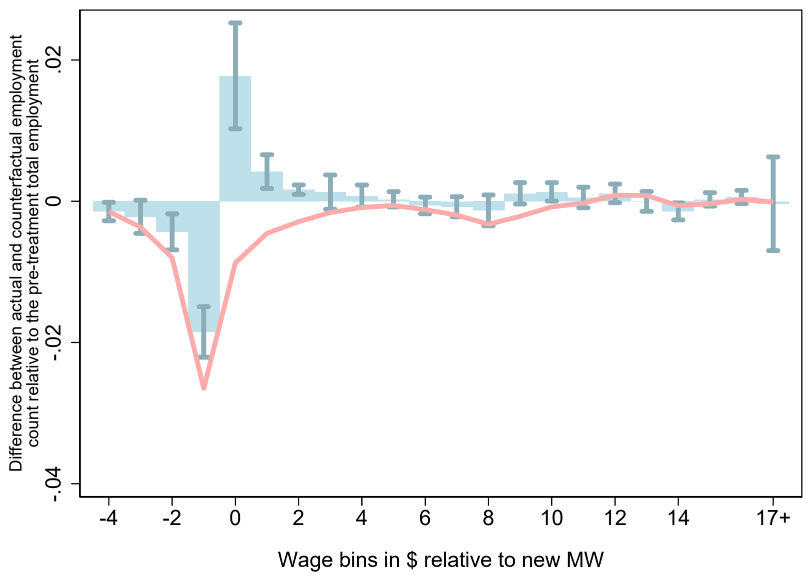

Figure A13: Impact of the Minimum Wage on the Wage Distribution
The bars show the estimated average employment changes in each bin during the five-year post-treatment period relative to the average total employment rate across provinces one year before the treatment. The error bars show the 95% confidence intervals using standard errors that are clustered at the provincial level. The red line shows the cumulative employment changes up to the corresponding wage bin.

to the specifications used to produce Figure 4 in Cengiz et., where effect of a minimum wage change on the wage distribution in the year following an increase (i.e. $\tau = 0$) are presented for incumbents and stayers.[39]. We continue to use our definition of stayers (greater than two months of job tenure) and joiners (up to two months of job tenure). Again, all figures are normalized by the average employment rate in the year immediately proceeding a minimum wage increase.[40] The results are presented in Figures A15 and A14 respectively.

The results in the bin below the new minimum wage and at the new minimum wage are larger in magnitude both Figures are larger than the five-year averages displayed in Figure A13, which suggests that the minimum wage effects attenuate over time.[41] For

[39]As the Current Population Survey does not include a tenure variable, these classifications are based on a comparison of employment status in the fourth and eighth month in which the CPS follows an individual. Unlike their results in Figure 2, the results in Table 4 of Cengiz et al. is based only on information from the eighth month, provided that their algorithm provides a match between the records from the fourth and eighth month.

[40]For Stayers, the average pre-treatment employment rate is .607, while for Joiners, it is .036.

[41]This may also be an artifact of minimum wages changing more frequently in Canada than the United

Joiners, positive spillovers are only seen up to a dollar above the minimum wage. Again, this parallels the results in Cengiz et al., where for entrants, the minimum wage effect is confined to the bin in which the minimum wage occurs. This differs somewhat from our results obtained using the hazard estimator, where we find minimum wage effects up to $2 above the minimum for Joiners.

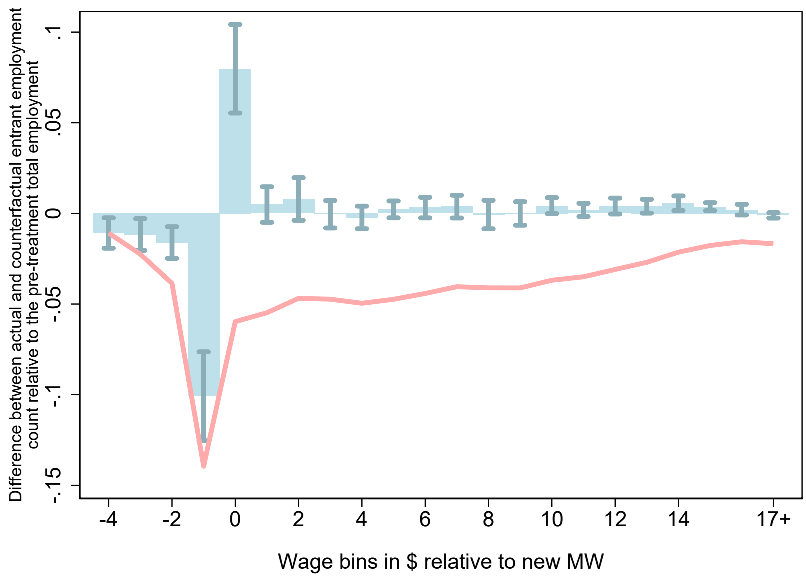

Figure A14: One-year Impact of the Minimum Wage on the Wage Distribution of Joiners
The bars show the estimated average employment changes in each bin in the year immediate following treatment relative to the average total employment rate across provinces in the year leading up to the treatment. The error bars show the 95% confidence intervals using standard errors that are clustered at the provincial level. The red line shows the cumulative employment changes up to the corresponding wage bin.

For Stayers, there is a larger decrease in employment below the new minimum wage, which is partly offset by a large increase in the bin containing the minimum wage and a smaller but statistically significant increase in employment in the bin $1-$2 above the minimum wage. This is again comparable to our results, with an increase in the proportional hazard up to $2 (see Table 1). It is also similar to Figure 4(b) in Cengiz et al., with the caveat that the magnitude of the effects we find in the bin including and the bin just below the minimum wage are approximately three times as large .

States.

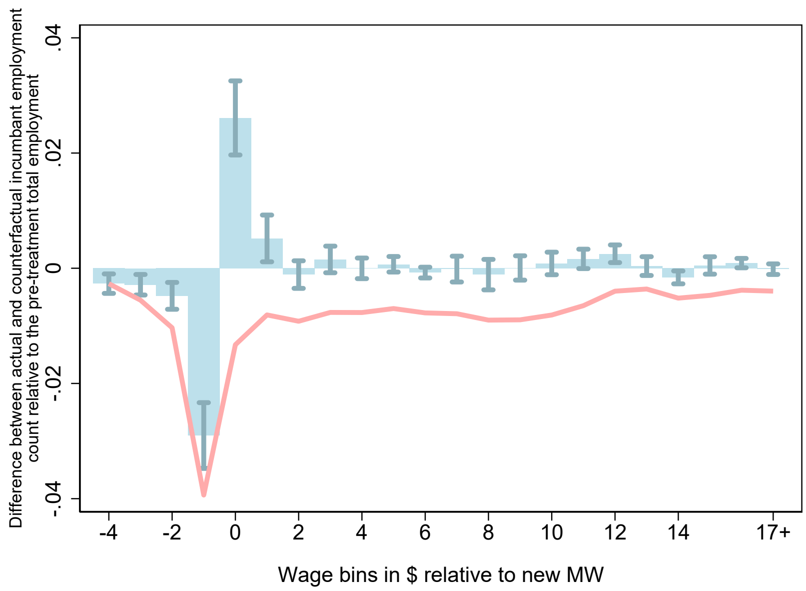

Figure A15: One-year Impact of the Minimum Wage on the Wage Distribution of Stayers
The bars show the estimated average employment changes in each bin in the year immediate following treatment relative to the average total employment rate across provinces in the year leading up to the treatment. The error bars show the 95% confidence intervals using standard errors that are clustered at the provincial level. The red line shows the cumulative employment changes up to the corresponding wage bin.