The Minimum Wage, Turnover, and the Shape of the Wage Distribution

Pierre Brochu, David A. Green, Thomas Lemieux, and James Townsend

March 16, 2015

Abstract

This paper proposes an empirical approach for jointly modelling the impact of the minimum wage on the wage distribution, and on movements in and out of the workforce. 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 in the presence of employment effects linked to the minimum wage. We use the estimates 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 find significant differences across the three groups.

1 Introduction

There is now a growing number of papers examining the impact of the minimum wage on the wage distribution. In a key paper, Teulings(2000) introduces that puzzle that, working with US data, there is a tendency to find small or no employment effects from a minimum wage increase for most workers but substantial impacts on inequality. One possible answer to this puzzle is that minimum wage effects on employment are not as insubstantial as they first appear. Brochu and Green(2013) argue that the lack of impact of minimum wages on employment rates for adults actually reflects a tendency for both hiring rates and separation rates from jobs to decline in an almost exactly offsetting manner in high versus low minimum wage regimes. Lemieux(2011) argues, further, that a joint consideration of distributional and employment effects of minimum wages is needed to truly understand minimum wage effects. Based on this, we decompose the distributional effects of minimum wages into effects for workers moving out of employment, workers for those moving into employment, and for workers continuing in employment. We find significant differences across the three groups.

Methodologically, we use an estimator proposed in Donald, Green and Paarsch(2000) and used in a more limited examination of wage effects of minimum wages for teenagers in Green and Paarsch(1996). The key feature of this estimator is that it focuses on the effects of the minimum wage on the hazard rate for wages rather than on the density or distribution function. This is useful because it immediately eliminates effects from the re-scaling of the distribution that follow either because of associated layoffs (Meyer and Wise(1983)) or monopsony effects (Butcher et al(2012)). In fact, we find that the minimum wage generates a decrease in the hazard rate just above the minimum wage. But when we translate these effects to impacts on the density, the truncation readjustment leads to an observed lack of impact on the density in the same wage range. We view this as supporting earlier work showing limited spillover effects on the spread of the wage density but still implying that minimum wage increases do affect above minimum wages (as revealed in the hazard rate effects). Thus, our results do not alter prior conclusions about the impact of minimum wages on the extent of inequality in the wage distribution but do provide interesting new evidence on how wages are set in an economy.

2 A Decomposition of Short Run Minimum Wage Effects on the Wage Distribution

We will start with a decomposition of the change in the wage distribution between the period just before to the period just after a minimum wage increase. For simplicity, we will initially assume that the environment is otherwise stationary, i.e., that there would not have been any change in the wage distribution in the absence of the minimum wage change. More specifically, we will consider an environment in which there are ongoing job separations and hirings, with the numbers of separations and hires being equal. For the wage distribution to be stationary, the distribution of wages for those separating and those being hired will be the same and both will be equal to the (unchanging) wage distribution for those remaining on the job between periods.

To write out the decomposition, we require notational definitions. We will consider hourly wages, w, with a cumulative distribution function, F(.). We are interested in how F(.) changes between period t and t+1, bracketing a change in the minimum wage from m_0 to m_1. The distribution can change shape for three reasons: 1) because firms lay off workers in response to the minimum wage increase; 2) because firms change who they hire and/or how they pay new hires in response to the minimum wage increase; and 3) because firms change how they pay workers who remain with the firm. To implement a decomposition into these three components we will define a set of dummy variables: S(t,t+1) = 1 for workers who are employed in both periods t and t+1 ("Stayers") and 0 for anyone else; L(t,t+1) = 1 for workers who are employed in period t and non-employed in period t+1 ("Leavers") and zero for anyone else; J(t,t+1) = 1 for individuals who are non-employed in period t and employed in period t+1 "Joiners" and 0 for anyone else; and E(t) = 1 for people employed in period t and 0 for non-employed individuals.

Using that notation, we can write the value of the CDF at some wage, w, in period t as:

$$F_t(w|m_0) = Pr(S(t,t+1)|E(t)=1,m_0) \cdot F_t(w|S(t,t+1)=1,m_0) \tag{1}$$
$$+Pr(L(t,t+1)=1|E(t)=1,m_0) \cdot F_t(w|L(t,t+1)=1,m_0)$$

and the value of the CDF in period t+1 as:

$$F_{t+1}(w|m_1) = Pr(S(t,t+1)|E(t+1)=1,m_1) \cdot F_{t+1}(w|S(t,t+1)=1,m_1) \tag{2}$$
$$+Pr(J(t,t+1)=1|E(t+1)=1,m_1) \cdot F_{t+1}(w|J(t,t+1)=1,m_1)$$

Combining these expressions, the change in the value of the CDF is:

$$F_{t+1}(w|m_1) - F_t(w|m_0) = \tag{3}$$

$$+(Pr(J(t,t+1|E(t+1)=1,m_1) - Pr(L(t,t+1|E(t)=1,m_0))$$

$$\cdot[F_t(w|L(t,t+1,m_0) - F_t(w|S(t,t+1)=1,m_0)]$$

$$+[F_{t+1}(w|J(t,t+1)=1,m_1) - F_t(w|L(t,t+1)=1,m_0)]$$

$$\cdot Pr(J(t,t+1|E(t+1)=1,m_1)$$

$$+(F_{t+1}(w|S(t,t+1)=1,m_1) - F_t(w|S(t,t+1)=1,m_0))$$

$$\cdot Pr(S(t,t+1)|E(t+1)=1,m_1)$$

Thus, changes in the CDF can be broken into three components: 1) the difference in proportion of workers who are Joiners in t+1 and the proportion of workers who are Leavers in t weighted by the extent to which the CDF for times Leavers in t differs from the Stayer distribution (captured in the first two lines of the decomposition); 2) the difference between the wage distributions of the Joiners and Leavers times the proportion of joiners in t+1 (lines 3 and 4); and 3) the change in value of the CDF for Stayers times the proportion of workers who are Stayers in t+1 (lines 5 and 6); 3).

We can also consider longer term changes in the wage distribution. In particular, consider comparing the wage distribution at t (which we assume is a pre-change steady state) with the wage distribution at t+s. Between t-1 and t, the wage distribution will not change, implying that the proportion Joiners and Leavers will be equal, the wage distribution for Stayes will be unchanged, and the wage distributions for Joiners and Leavers will be unchanged and equal to one another.[1] The same will be true between t+s-1 and t+s. But while the wage distribution is stationary in each case, the overall wage distribution and, associated with it, any of its component parts may differ between the two steady states. Thus, the change in the steady state wage distributions can be

[1]The proportion of Joiners and Leavers are equal in the steady state because we have not so far allowed for any dynamics within employment. If, instead, we were to think of initial hires having some probability of being promoted to jobs paid from a superior wage distribution then stationarity would require flows into and out of the initial hire state to be equal and for flows out of both initial and later employment states into non-employment to be equal to the flows into initial employment from non-employment. Allowing for these dynamics would not change the main points of our discussion. The states we present here fit with our empirical implementation.

decomposed as,

$$F_{t+s}(w|m_1) - F_t(w|m_0) = (Pr(S(t+s-1,t+s)|E(t+s) = 1, m_1) \tag{4}$$

$$-Pr(S(t-1,t)|E(t) = 1, m_0)) \cdot F_t(w|S(t-1,t) = 1, m_0)$$

$$+(F_{t+s}(w|S(t+s-1,t+s) = 1, m_1) - F_t(w|S(t-1,t) = 1, m_0)) \cdot Pr(S(t+s-1,t+s)|E(t+s) = 1, m_1)$$

$$+Pr(J(t+s-1,t+s|E(t+s) = 1, m_1) \cdot [F_{t+s}(w|J(t+s-1,t+s) = 1, m_1) - F_t(w|J(t-1,t) = 1, m_0)]$$

$$+(Pr(J(t+s-1,t+s|E(t+s) = 1, m_1) - Pr(J(t-1,t|E(t) = 1, m_0)) \cdot F_t(w|J(t-1,t,m_0)$$

where, we have written this imposing the restrictions that the proportion of Joiners equals the proportion of Leavers and that their wage distributions are equal in each steady state. Thus, the overall wage distribution can change because the flows in and out of work are different between the two equilbria, because the wage distributions for the Joiners (and Leavers) are different, or because the wage distributions for the Stayers are different.

3 Models of the Impact of Minimum Wages on the Wage Distribution

Equations (3) and (4) are mechanical decompositions of observed changes in the wage distribution corresponding to a minimum wage change in an otherwise stationary environment. It is useful to consider the decompositions in light of models of the impact of the minimum wage on the wage distribution. The simplest of these is a neoclassical competitive labour market model in which workers are paid the values of their marginal products, the wage dispersion corresponds to the dispersion of worker skills, and skills are not substitutable in production. In that case, the equilibrium wage distribution at time t will be a truncation of the underlying skill distribution at the skill level corresponding to m_0. The increase in the minimum wage implies moving to a new equilibrium where the same distribution is truncated at the higher point, m_1. If adjustment costs are such that firms make the transition to the new stationary state immediately then all the workers with period t wages between m_0 and m_1 will be laid off.

The minimum wage increase has three effects seen in terms of the short run decomposition. The first is the difference in the wage distributions for the Leavers in period t and Joiners in period t+1. In particular, the wage distribution for Leavers in t will be heavily weighted toward workers who earned wages between m_0 and m_1 in period t while the Joiners distribution in t+1 will simply be the distribution of skills truncated at m_1. The second is the difference in the proportion of Leavers in t and Joiners in t+1. This also reflects the fact that there will be many more Leavers in t as firms drop workers who formerly earned below the minimum wage. Because Stayers are the complements of Joiners and Leavers in the distributions, this necessarily implies the proportion of Stayers will also change between the period (i.e., that the first line in (3) will not equal zero).

The third effect is the one emphasized in Meyer and Wise(1983), that the increase in the truncation point implies an associated inflating of the wage density above the new cut-off point even if nothing else changes in order to insure the new, more truncated distribution integrates to 1. In parts of the previous literature on minimum wage effects, any impact on the wage distribution above the new minimum wage after a minimum wage increase is called a "spillover effect" (e.g., Butcher et al(2012)). In line with the discussion in Stewart(2012), we will refer to an effect which changes the scaling of the wage distribution above the new minimum wage but does not change its shape as a truncation spillover effect. We will refer to changes in the shape of the wage distribution above the new minimum as a shape spillover effect. It is not easy to see the truncation spillover effect in the decomposition in equation (3). In fact, it is not really a separate effect in its own right but arises as a result of the difference in the proportion of Leavers versus Joiners. In the context of this simple model, there is no shape spillover effect. This means that if we considered the isomorphic description of the wage distribution in terms of hazard rates instead of densities, we would find no effect on the hazard for wages above m_1. For example, the hazard rate at m_1+1 (i.e., the probability a worker has a wage of m_1+1 conditional on having a wage of at least m_1+1) is the same before and after the minimum wage change. We will carry out our estimation in terms of hazard rates which has this advantage of allowing us to differentiate between truncation spillover effects (which do not alter hazard rates above the new minimum wage) and shape spillover effects (which alter at least some of those hazard rates).

The long run steady state decomposition contrasts with the short run decomposition because in the simple model we have adopted to this point with separation and joining rates being independent of the wage conditional on the wage being above the minimum wage, the separation and joining rates are the same in both steady states and are equal to each other. That means that the component related to the initial layoffs (implying that layoffs before the minimum wage exceed the proportion joining employment afterwards) is not present. Instead, the differences come down simply to the Joiners, Leavers and

Stayers wage distributions all being truncated at a higher point. As with the short run case, there will be no change in the hazard rate above m_1 for any of the sub-groups or for the overall distribution.

The simple model has the significant disadvantage that it is not even a complete description of minimum wage effects within a standard competitive model. Teulings(2000) develops a model with a continuum of worker skills and of task complexity. Under a standard assumption on skill-complexity complementarity, the most skilled workers match with the most complex tasks and the substitutability between workers of differing skill levels declines in the difference between those skill levels. When the minimum wage increases, firms shift demand away from the workers whose wages were raised the most toward the most substitutable workers. This will have the effect of raising the wages of workers with skills that would earned them wages just above the new minimum wage. As that happens, there is an incentive for firms to reach into the set of workers who earned just less than m_1 before the minimum wage increase since they are highly substitutable for the above minimum wage workers whose wages have increased. The end result is fewer layoffs than predicted in the simple model, a truncation of the post-change wage distribution at m_1 and added mass in the wage distribution just above m_1.

In terms of the short-term decomposition, there will again be an excess of Leavers relative to Joiners (though not by as much as in the simple model). The wage distribution for Leavers will again be heavily weighted toward those with wages between m_0 and m_1 in period t (though with the mass more weighted toward the bottom end of that range). The major difference is the Stayers' wage distribution. Since the Stayers are the same set of people in period t and t+1, the increase in demand for them will imply that their wage distribution will lose mass just above m_1 as the equilibrium price for their skills increases. In terms of hazard rates, the hazard rate will decline just above m_1, recovering gradually back to its initial level as we move up the wage distribution toward workers with skills that are less substitutable for minimum wage workers. That is, the Stayers experience shape spillovers and that carries over to the overall distribution.

It is interesting to think of the wage distribution changes experienced by each subgroup in terms of selection effects (i.e., changes in the set of skills that are employed) versus changes in the pay policy of firms (i.e., changes in the prices for the skills). For Leavers, the effect on the shape of their wage distribution is entirely due to selection effects - the tendency of firms to lay off less skilled workers - since we examine their wage distribution before the minimum wage change. For Stayers, in contrast, the shape change is entirely due to pay policy changes since they are the same set of workers before and after the minimum wage change. The change in the wage distribution for Joiners represents a complex combination of the two since a different set of worker skills are hired after the minimum wage increase but the prices paid for those skills have also changed. Because of that, we cannot make any clear predictions about changes in the shape of the Joiners distribution.

For the long run decomposition, there is again a difference relative to the short run approach because the set of layoffs that constitute part of the immediate adjustment to the new minimum wage does not play a role. There will be a difference in the wage distribution for Stayers in the two steady states but now we can no longer claim that the difference is due only to changes in skill prices. The Stayers skill distribution in the new steady state will come to reflect the different set of Joiners and Leavers under the new minimum wage. As with the short run case, we will see effects in the hazard rates for the distributions of workers in all sub-groups and for the overall wage distribution above m_1. That is, there will be both truncation and shape spillover effects. Concentrating on the hazard rate allows us to get a clearer picture of the extent of the latter.

Both of the neoclassical, competitive models discussed so far imply a simple truncation of the wage distribution. However, in many instances (including, as we will see, Canada in the time period we study) there is as substantial spike in the wage distribution at the minimum wage. A variety of authors (e.g., Teulings(2000), Manning(2003), Flinn(2006)) have argued that the existence of the spike points to a non-competitive model of the labour market. Butcher et al(2012) present a non-competitive model that has some similarities to the simple competitive model. In the basic form of their model, workers are homogeneous but firms differ in productivity and each firm is modelled in a reduced form way as facing a less than perfectly elastic labour supply curve. This implies that firms act in the same way as monopsonists, following a policy in which workers are paid less than their marginal product. The imposition of a minimum wage will cause some firms (those with productivity below the minimum wage) to go out of business. At the same time, we can define a productivity level, A^* at which firms would pay exactly the minimum wage if the minimum wage did not exist. All firms with productivity between the minimum wage and A^* will stay in business but pay the minimum wage, generating a spike in the wage distribution. Importantly, firms with productivity above A^* will pay the same wage as before. This means that, just as in the simple competitive model, the model only contains truncation spillover effects, with the wage distribution above the minimum wage maintaining its shape but being re-scaled by a factor that depends on the extent of the "bite" of the minimum wage relative to the underlying wage distribution.

Building on the insights from this model, Butcher et al(2012) construct a more realistic model in which workers have different skill types. For a given skill type, the minimum wage effects are the same as in their basic model but with the minimum wage having a different bite for different skill groups, a minimum wage change can have a more complex effect on the shape of the overall distribution above the minimum. For our purposes, if we were able to focus on one skill group, the impact on the number of Joiners and Leavers in the short run is not clear. On one hand, a certain number of firms (those with productivity below m_1 will leave the market. On the other hand, the set of firms with productivity near A^* will increase their number of employees, just as in a monopsony model. In terms of the wage hazard rates, the firms put out of business by the minimum wage increase will all have been paying m_0 before the change and so the Leavers should have an extra large spike in the hazard rate at m_0. Above the minimum wage, the hazard should be constant for Leavers, Joiners and Stayers since there are only truncation spillover effects. The distribution of wages between m_0 and m_1 in period t will map directly into firm productivity levels. The firms paying close to m_1 will increase employment due to standard monopsony effects, so we should see a reduction in the hazard for Leavers in that range. Overall, the effect of an increase in the minimum wage on the hazard rate for Leavers should be positive at and possibly just above m_0 then decreasing and turning negative as we approach m_1 from below, and, finally, becoming zero at and above m_1. For Stayers, there should be no effect on the hazard because their distribution is truncated as m_1 in both periods. For Joiners, the effect on the size of the spike in the hazard at m_1 is uncertain because of offsetting firm exit and monopsony effects but, again, there should be no effect above the minimum wage. If we allow for different skill types but no substitution among the types then these results won't carry over exactly to the full distribution because the mixture of hazards for the different types is not equal to the hazard of the mixtures but for small enough changes in the minimum wage (which would mostly have bite on a small set of low skilled types) the implications might not be very different. If we go further and allow the type of substitution discussed in the competitive model in Teulings(2000) then there might be shape spillover effects of the type discussed earlier. For Leavers, this would seem to imply the same qualitative effects on the hazard below m_1 as before. For Stayers, in the short run, it would imply declines in the hazard just above m_1 in period t+1 because of increases in skill prices.

Flinn(2006) presents a model in which firms that are again heterogeneous in productivity match with homogeneous workers in a frictional labour market. Wages are determined by individual bargaining. As in the Butcher et al(2012) model, firms with productivity below m_0 will go out of business while those with productivity between m_0 and a threshold level of productivity at which firms and workers would bargain m_1 will all stay in business and pay the minimum wage. Thus, this model also implies a spike at the minimum wage. A key difference relative to the Butcher et al(2012) model is that the change in the minimum wage implies different bargained wages even in matches that paid wages above m_1 in the pre-change equilibrium. This arises because the outside option for workers changes as the minimum wage changes. In particular, if workers leave the current match, they will remain longer in the non-employed state (because of the elimination of firms with productivity below m_1) but will get higher wages once employed because the wage distribution from which they will draw is truncated at a higher point. The net effect is uncertain. Given this, we should again expect an increase in the spike in the hazard for Leavers at m_0 but no effect on the hazard between m_0 and m_1 (since all of these workers would remain employed while having their wages increased to m_1) and no effect on the hazard above m_1 since these workers would also not face any added employment or disemployment effects. For Joiners and Stayers, there could be effects on the shapes of the hazard function above m_1 but the size or even the direction of the effects are uncertain.

Brochu and Green(2013) examine changes in job hiring and separation rates in high versus low minimum wage regimes. They find that turnover rates of both kinds decline when minimum wages are higher. For workers over age 20, these effects are almost exactly offsetting. This fits with the standard finding that minimum wage increases have no effect on employment rates for older workers but implies that underlying this is a shift in the labour market equilibrium toward one with both longer employment and nonemployment spells. They argue that the reduced job separation rates can fit with a model in which outside options for firms decline with the minimum wage since the chance that the next match a firm will make will be a (lower profit) minimum wage match increases. This would imply that wages in above minimum -wage matches (specifically, wages for Stayers in the short run decomposition) shoudl increase. Thus, as in the Teulings(2000) model, we should see declines in the hazard rate for Stayers just above m_1 in the short run.

Finally, several authors have discussed variants on efficiency wage models in which worker notions on what constitutes a fair wage are benchmarked on the minimum wage (e.g., Grossman(1983) and). Falk et al(2006) provide experimental evidence in favour of such a model. The implications of such a model for the shape of the wage distribution, however, have not so far been worked out in detail. Suppose that we follow x in assuming that market wages are generally seen as fair in the sense of balancing the requirements of workers and firms. In that case, workers who formerly earned m_0 and have a new wage of m_1 will be seen as being given something like an unearned gift. Workers formerly earning just above m_1 might, in fairness, expect to get as sizeable an increase but, on the other side, even an increase to m_1 is still a gift relative to the old equilibrium. This would mean that some set of workers earning between m_0 and m_1 in period t would have their wages raised up to but not above m_1. Some of those higher in the m_0 to m_1 interval would likely have their wage raised above m_1 and those with period t wages just above m_1 would likely also get fairness based wage increases. With sub-m_1 wages in period t hitting something of a sticking point at the new minimum wage and wages above m_1 being increased, the net effect could be a decline in the hazard rate just above the minimum wage.

While many of these models are set out with homogeneous workers, a common argument is that there are actually multiple skill types and that the higher in the skill distribution we move, the less the direct or indirect impact of the minimum wage. This is discussed explicitly in Teulings(2000) and in Butcher et al(2012) but it could also be a factor in the fairness based models if higher skilled workers do not compare themselves directly to workers earning near the minimum wage. As Stewart(2012) discusses in detail, this raises the possibility of identifying minimum wage effects by using higher wage workers as a control that identifies general economic forces with comparisons between wage movements for those workers and workers near the minimum wage being used to identify minimum wage effects. But Steward(2012)'s also shows that such a strategy may not work well since wage movements in different parts of the upper part of the wage distribution are often different. We will return to this point in describing our own approach, which uses a variant of this identification strategy.

3.1 Previous Empirical Results

There is a large literature showing that the minimum wage has a substantial impact on the wage distribution. Not surprisingly, the impact is very much concentrated at the bottom end of the distribution. For instance DiNardo, Fortin, and Lemieux (1996) show that the decline in the real value of the minimum wage in the United States during the 1980s had a large impact on the 50-10 gap among women, and a smaller but still substantial impact among men. While DiNardo, Fortin, and Lemieux (1996) only computed simple counterfactual experiments under the assumption of no spillover effects, Lee (1999) proposed a new approach exploiting differences in the relative minimum wage across different regions of the United States to estimate the magnitude of spillover effects. He found large spillover effects and concluded that the decline in the real minimum wage accounted for all of the increase in inequality in the lower end of the wage distribution during the 1980s. More recent research by Autor, Manning, and Smith (2010) concludes that spillover effects, while important, are not quite as large as those found by Lee (1999).

Unlike Lee (1999) who only exploits minimum wage variation linked to the fact that the same national wage is relatively higher in low wage regions, Autor, Manning, and Smith (2010) also use variation in state minimum wages to estimate spillover effects. Likewise, Fortin and Lemieux (2015) estimate a Lee-type model using provincial variation in minimum wages in Canada. One can think of this approach as a difference-in-differences strategy aimed at estimating the impact of the minimum throughout the wage distribution. Like Autor, Manning and Smith (2010), Fortin and Lemieux find some limited evidence of spillover effects, though those dont tend to go above the 10^{th} or 15^{th} percentile of the wage distribution for all workers.

Another interesting set of papers have exploited the introduction of the new U.K. national minimum wage in 1999, and changes in it thereafter, to estimate its impact on the wage distribution. Using a broad cross-section of the U.K. workforce from the Labour Force Survey, Dickens and Manning (2004a) find that the new minimum had essentially no impact on the wages of workers earning more than the minimum wage. Dickens and Manning (2004b) reach a similar conclusion using a new survey data set for workers in the home care sector who should potentially be more affected by the minimum wage, given the low wages paid in that sector. Stewart (2012) also finds that spillover effects are at best modest, and sensitive to specification issues. The only study to find substantial spillover effects of the minimum wage in the U.K. is Butcher, Dickens, and Manning (2012). Interestingly, that study uses a different estimation strategy that exploits regional variation in the relative value of the minimum wage and is, as such, more comparable to the U.S. papers by Lee (1999) and Autor, Manning, and Smith (2010).

One important shortcoming of this literature is that since most studies focus on the wage distribution among workers, some of the distributional effects of the minimum wage may be confounded by possible employment effects of the minimum wage. For instance, consider an extreme version of the competitive model where workers are paid their marginal product, and all workers whose marginal product is smaller than the minimum wage are not employed. In this extreme truncation model, all workers earning less than a new minimum wage lose their jobs when the minimum wage increases. As a result, all percentiles of the wage distribution among employed workers go up, which give the misleading impression that the minimum wage has large spillover effects.

While this example is admittedly an extreme case, it shows that failing to simultaneously consider employment and distributional effects of the minimum wage can bias estimates of the distributional effect of changes in the minimum wage. Existing studies also fail to distinguish whether changes in the wage distribution are driven by what happens to leavers who may have lost their job because of the minimum wage, joiners who may be in a better bargaining position because of a higher minimum wage, or stayers. We next explain how our empirical approach enables us to fill these gaps in the existing literature.

4 Estimation Approach

Our goal is to estimate the various components of the two decompositions (3) and (4). This means, first, estimating the effect of a change in the minimum wage on the overall distribution of wages and on the wage distributions for the three subgroups: Leavers, Joiners, and Stayers. Under an assumption that firms do not alter offered wages before the implementation of a given minimum wage increase, the effect of a minimum wage change on the wage distribution of employment Leavers reflects pure selection effects of the minimum wage increase. If we combine it with estimates of the effect of a minimum wage change on the marginal probability of a person leaving employment, we can pursue one of the goals set out in Lemieux(2011): determining whether all of the workers fomerly earning below the new minimum wage are laid off after a minimum wage increase or, more generally, which set of workers are more likely to be laid off. As we have discussed, understanding these effects will help in distinguishing between competitive and non-competitive models of different types. In contrast, in the short run, changes in the wage distribution for Stayers reveals changes in wage setting by firms since the set of workers being examined doesn't change. From this we can see whether workers who formerly earned below the new minimum wage and remain employed have their wages raised just to the new minimum wage or above it. If we are willing to assume that the changes in wage setting policies identified by the Stayers outcomes apply equally to the employment Joiners then any difference between Stayer and Joiner effects will identify changes in the selection of new hires induced by a minimum wage increase.

Our discussion of models of minimum wage impacts on the wage distribution indicated that there is an advantage to estimating minimum wage effects on the hazard function for wages rather than directly on the distribution. In particular, we are interested in distinguishing between what we have called truncation and shape spillover effects. By working with hazard rates, we effectively eliminate truncation spillover effects and can better identify whether increases in the minimum wage alter wages paid to above-minimum workers. Our approach is a variant of the survivor function approach set out in Lemieux(2011) that is based on the estimator in Donald, Green and Paarsch(2000). As we describe later, this estimator has the advantage of allowing for very flexible effects of covariates on the shape of the wage distribution while also yielding estimates of the distribution function that are consistent in the sense of always lying between 0 and 1.

4.1 The Estimator

The estimator models $F(y|x)$ using a hazard function approach, where $F(y|x)$ is the conditional (on employment) cumulative distribution of wages, y, given a vector of covariates, x. In particular, 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 $F(y|x)$ are also guaranteed to be consistent in the sense that the distribution function always lies between 0 and 1.

A typical approach in the duration literature is to adopt a proportional hazard specification in which $h(y|x) = exp(x\alpha)h_0(y)$, where $h_0(y)$ is the baseline hazard common to all individuals. While this approach allows for a flexible specification of the baseline hazard, it also forces a given covariate to shift all parts of the hazard function up or down by the same proportion. To allow for the effects of x to differ at various points in the wage distribution, we will use a specification which allows for interactions between the covariates x and the baseline hazard.

A straightforward way to introduce such interactions and to allow for a flexible specification of the baseline hazard is to first divide the earnings distribution into P segments:

$$y_0 < y_1 < ... < y_P \tag{7}$$

The survivor function then becomes

$$S(y|x) = exp\left(-\sum_{i=1}^{p-1}\int_{y_{i-1}}^{y_i} exp(x\alpha_i)h_0(u)du - \int_{y_{p-1}}^{y} exp(x\alpha_p)h_0(u)du\right), y \in (y_{p-1}, y_p) \tag{8}$$

As in Meyer(1990), suppose that the effect of x is constant within each segment $(y_{p-1}, y_p), p = 1..P$. Then the hazard function for a given segment is given by

$$h(y|x) = exp(x\alpha_p)h_0(y). \tag{9}$$

where α_p denotes the effects on the hazard of the covariates within the segment. This implies that

$$\int_{y_{p-1}}^{y_p} exp(x\alpha_p)h_0(u)du = exp(x\alpha_p)\gamma_p \tag{10}$$

where γ_p corresponds to the integration of the baseline hazard over the interval (y_{p-1}, y_p). Using this, we can calculate the probability that earnings lie in the interval (y_{p-1}, y_p), which equals $S(y_{p-1}|x) - S(y_p|x)$. With a suitably large number of intervals, P, this framework yields a flexible specification of the conditional earnings distribution.

Flexibility in the effects of the covariates in this approach is generated from the fact that the covariates have a different effect, α_p, in each interval. However, this flexibility comes at a cost: in order to have enough degrees of freedom to effectively estimate the α vectors, we would need to restrict the number of baseline segments, P. The result could be an over-smoothing of the shape of the hazard function. We address this problem by allowing for a large number of baseline segments (P = 164) but restricting $\alpha_p = \alpha_s$ within each of s = 1, ..., S non-overlapping sets of baseline segments. We will call the sets within which the covariate effects are constant "covariate segments" and will allow for

10 such segments in our implementation, details of which are given in the next section. Notice that 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. The estimator also naturally incorporates top-coding, which is just right-censoring in this context.

Within the context of this estimator, minimum wages are analogous to time varying covariates in the standard duration literature. In particular, to estimate the impact of the minimum wage on the distribution right at the minimum wage, we can introduce a covariate that takes a value of one in the baseline segment that contains the minimum wage and zero in all other segments. The coefficient estimated on this variable is 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. We can similarly define variables that take values of one for ranges related to the minimum wage (e.g., from 10 to 30 cents above the minimum wage) and, using those estimates, can map out the effect of the minimum wage on the shape of the wage distribution.

Based on this, we have two types of covariates in our estimator. The first are standard covariates whose effects are allowed to vary between covariate segments but are constant within segments. We will describe the exact set of covariates we include when we discuss our data but we broadly include controls for gender, age and education. In addition, we include a complete set of province and year dummy variables. The second type of covariates are the ones related to the minimum wage, which are imposed to have the same coefficients regardless of the baseline or covariate segment. We will estimate everything separately for males and females so these minimum wage effects are allowed to vary by gender. We could, further, include interactions between the two types of covariates, which would allow for differences in minimum wage effects by, for example, age or education group. To this point, we have not investigated such interactions. This effectively means that we are assuming that minimum wages affect low wage workers in the same way regardless of whether they correspond to very low earning university graduates or relatively average earning high school drop-outs.

Given the inclusion of year and province effects, our identification approach is a variant of a difference in differences approach. In particular, we allow for wage distributions to differ in potentially complex permanent ways across provinces. We also capture common shifts in the provincial wage distributions over time. The effect of the minimum wage is then defined relative to these province and year differences. We restrict the minimum wage to have no effect on the wage distribution beyond $ 4 above the minimum wage in any given province by month cell. As discussed in Stewart(2012), this effectively makes the upper part of the wage distribution part of the control group. The separate identification of the baseline hazard and the minimum wage effects then comes through a recursive argument. Consider, for example, estimating an effect for $3 to $4 above the minimum wage. In one province, the minimum wage could be $8 in one period and 9$ in the next. When the minimum wage is at $8, the wage range between $12 and $13 is assumed not to be affected by the minimum wage and so the baseline for this region can be directly estimated (controlling for common year and province effects). When the minimum wage rises to $9 this same wage range is affected by the $3 to $4 above the minimum wage effect. Thus, the difference between the hazard rates between $12 and $13 under the two minimum wage regimes (controlling for general differences identified off of provinces where there was no minimum wage change) identifies the $3 to $4 above the minimum wage effect. Knowing that effect, we can back out the baseline estimates for the $11 to $12 range by examining the difference between the actual hazard in this range when the minimum wage is 8$ and the estimated $3 to $4 above the minimum wage effect. We can continue working in this way down the distribution. This means, in principle, we can separately identify the hazard rate in the absence of the minimum wage and the effects of a minimum wage on the hazard.

Two further points about identification are of interest. First, as described earlier, Stewart(2012) showed that using changes in high percentiles (e.g., the 50th percentile) could be problematic because they appear to be following different trends from the lower percentiles in his UK data. Our estimator allows for different trends in different parts of the wage range corresponding to the covariate segments. Moreover, our identification is based just on movements in the wage range directly above the top of our minimum wage effect range ($4 above the minimum wage in our case). Second, we could, in principle, allow for the minimum wage to affect the hazard rate throughout the wage range by estimating a differences in differences approach for the hazard at each wage level. Strictly speaking, we would no longer be trying to identify the baseline hazard in the absence of a minimum wage in this case. However, we have tried specifications in which we allow the minimum wage to have effects on wages above $4 above the minimum wage and find no statistically significant or economically substantial effects beyond this range. That means that the two approaches actually yield the same estimates. Imposing the restriction of no effects beyond $4 above the minimum wage allows for an extra source of identification and, as a result, greater efficiency.

It is important for this exercise that we are using hazards rather than directly estimating quantiles of the wage distribution. As discussed earlier, all quantiles of the wage distribution above the minimum wage will move if there are either disemployment effects or the types of monopsony effects described in Butcher et al(2010) in order to insure that the wage distribution integrates to one. Thus, there is no way to use higher quantiles as part of the control group when there are truncation spillovers. However, this does not apply to the hazard. If there are no shape effects on the distribution beyond a certain point then the hazard function beyond that point will not be affected.

5 Data and Implementation Details

5.1 Data

This section contains a brief description of the two main sources of data: provincial minimum wage data and Canadian Labour Force Survey (LFS) data. We also present basic patterns of the key variables of interest.

We use provincial minimum wage data that cover the 1979-2008 period. The minimum wage, as with other labour matters, falls under provincial jurisdiction in Canada.[2] Having each of Canada's ten provinces set their own minimum wage thus provides for a rich source of minimum wage variation. Some provinces have, at various points in time, adopted lower rates for special classes of workers (e.g. students in Ontario). Yet, the evidence shows that firms do not, for the most part, take advantage of these special categories (e.g., Card and Krueger (1995)). As such, this paper focusses on the general adult minimum wage for each province. To match our other data, we focus on monthly frequencies. In particular, we use the minimum wage in force on the 15th of each month as relevant for that month to ensure compatibility with our tenure data.[3]

The key explanatory variable in our regression analysis is the real minimum wage. We construct it by deflating the (nominal) minimum wage by the CPI for the same province and month, using 2002 as our base year. Appendix Figures 1a and 1b show the real minimum wage patterns by province and year. Importantly, the minimum wage shows

[2] Workers under federal jurisdiction (e.g. air transport) were the exception. Prior to 1996, there was a distinct federal minimum wage for those workers. Yet, the federal minimum wage was relevant to only a small subset of workers, and since 1996, the federal rate has adopted the general adult minimum wage of the province where the employer is usually employed.

[3] Tenure information is asked in the week which includes the 15th of the month.

considerable variation over time within each of the provinces.

The second source of data is the Canadian Labour Force Survey (LFS) master files.[4] The LFS is a large Canadian household survey involving interviews with approximately 50,000 households per month. The focus of the LFS is to gather information on labour market activities of Canadians. A critical variable for this study comes from the LFS tenure question which asks, "When did ... start working for his current employer". Based on the answer to this question, the LFS records the number of months of employment. What distinguishes the LFS from other Canadian data sets, and American data sets for that matter, is that this question (with no change in wording) has been asked every month since 1976.[5]

We restrict our LFS sample to individuals aged 15 to 59. We further exclude full-time students, the self-employed, and those in the military. Full-time students are not part of the study because working is not their main activity. The self-employed and those working in the military are removed because the processes that generate their job tenure spells are very different from (non-military) paid employees. The LFS began to include a question on wages each month starting in November 1996 and we use data from the point through to November 2012. Unfortunately, respondents are only asked about their wage on a job at the first interview. 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, we estimate the net effect of the minimum wage effect on the Stayer wage distribution, which could reflect any combination of the workers moving among different wages.

As described earlier, we are interested in decomposing workers into groups based on movements in and out of employment. 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 two-month mini panels.[6] However, we are constrained in some of our definitions by our need to use the incoming sample members

[4]These files were accessed on site at the Carleton, Ottawa, Outaouais local Research Data Centre (COOL RDC). This RDC is run and sponsored by Carleton University, University of Ottawa, Université du Québec en Outaouais, in collaboration with Statistics Canada, Social Sciences and Humanities Research Council, and Canadian Institutes of Health Research.

[5]See Brochu (2006) for a detailed discussion of the limitations of other North American data sets.

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

in a month in order to get "fresh" wages. Based on this, we define employment Leavers as those employed in month 1 but non-employed in month 2 (because we are only concerned with their month 1 wage, we can take advantage of the information from the mini-panel), employment Joiners as those who are employed in month 1 and have job tenure of 1 month, and employment Stayers as those who are employed in month 1 and have job tenure of 2 months or more.

We also use education broken down into 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; 20 to 24; 25 to 34; 35 to 54; and 55 to 59.

In Autor et al(2010)'s examination of the impact of the minimum wage on the US wage distribution, they argue that wage measurement issues can have a substantial impact on estimates of minimum wage effects. We respond to those concerns, in part, by examining 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. For the all earners sample, we calculate an hourly wage as their weekly earnings on their main job in the survey week divided by their hours worked in the survey week. As Autor et al(2010), this has the potential to create a "division bias" effect that could alter our estimates. We work with wages of hourly paid workers who do not receive tips or commission as our core sample, presenting the results from that sample in our main tables and figures. Working with the other samples does not alter the main patterns we discover or our conclusions. We put the results from those alternate sample in an appendix.

5.2 Implementation Details

In implementing our estimator with this data, we top-code wages at \$20 per hour since we don't believe the computing effort required to estimate the hazard beyond that point is useful when estimating minimum wage effects. We employ 164 baseline segments and 10 covariate segments. The first three baseline segments correspond to: $\leq$ \$3.00; \$3.01 to \$3.50; and \$3.51 to \$4.00. The remaining segments up to \$20.00 are 10 cents wide. The thresholds defining the covariate segments are set such that each segment contains 10% of the wages within our wage range (i.e., the wages between \$0 and \$20) for the sample of all earners pooled across years. To capture minimum wage effects, we use a set of dummy variables corresponding to ranges that are anchored on the minimum wage in force in the given province and month. The specific ranges for these dummies are 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 30 cents above; 31 cents to $1 above m; 1.01$ to $2.00 above m; $2.01 to $3.00 above m; $3.01 to $4.00 above m. There is also a dummy corresponding to the minimum wage itself which actually corresponds to the minimum wage plus or minus 10 cents to allow for some measurement error in reporting. Given the set of education, age, year and quarter variables - each with a different estimated effect in each covariate segment - plus the baseline hazard coefficients, the total number of estimated coefficients is 530.

In the actual implementation of the estimator, we first count the number of observations in a large set of cells defined by the complete interactions of the baseline wage segments, age, education, province, year, and quarter of the 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 work with the GLM command in Stata.[7]

Finally, we will use the estimated coefficients to form fitted versions of the hazard, density, and cumulative distribution functions. For both the hazard and density functions, the large number of baseline segments we use means that the simple predicted functions show local spikes to a distracting degree. In response, we form a smoothed hazard in the absence of minimum wage effects by using a 5th order moving average of the predicted hazard points. We then apply the minimum wage effects to that smoothed underlying hazard so that, for example, any spike at the minimum wage is visible.

[7] 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 for us, the dependent variable is not defined for all cells and we are forced to use the maximum likelihood option in GLM, which slows down the estimation time considerably. 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.

6 Results

In Table 1 and 2, we present the estimated coefficients on the set of dummy variables corresponding to the ranges at and around the minimum wage for males and females, respectively. The standard errors reported in this table (and that form the basis of all of our inference in this version of the paper) are clustered at the province level. Recall that we have already grouped our data at the wage bin x province x education group x age group x year x quarter level so that clustering effects that might have arisen from commonalities among individuals within these groups have already been addressed. The clustering at the provincial level allows for the possibility of cross-time correlation within provinces. Of course, given the work in Bertrand et al(2003), there is reason for concern that using this standard clustering estimator is not sufficient to address issues of over-rejection in our context with only 10 provincial units. Our context is somewhat different from that studied in Bertrand et al(2003) in that they examine the effect of a one-time policy shift with each policy unit eventually adopting the policy while we are studying a policy variable that takes many different values and moves both up and down. To check on our clustered standard errors, we have so far obtained bootstrapped standard errors on one of our core specifications. Given our number of parameters, this is a very time consuming exercise and so we haven't yet implemented it for all of our specifications. The result in this one case was that the bootstrapped standard errors were, on average, about 4% larger than the standard clustered standard errors. This was not enough of a difference to over-turn virtually any of the inference conclusions in the remainder of our discussion.

In our proportional hazards model, the actual impact of the minimum wage variables on the hazard is obtained by exponentiating the coefficients reported in Tables 1 and 2 and then multiplying them by the hazard determined by a combination of the baseline hazard and the effects of the other covariates. Given that, we will mainly present our results in the form of figures, but there are some immediate points that can be made examining these coefficients. First, for all four samples for both genders, there is a spike in the hazard at the minimum wage that is statistically significant at at least the 5% level. Interestingly, in all the cases the proportional effect on the hazard is about the same (1.1 or multiplying the underlying hazard by about 3). As we will see, though, the underlying hazards have different shapes for the different samples, resulting in different sized spikes in the densities. The various samples also share the feature that the hazard is substantially reduced below the minimum wage, though with a tendency for this extent of the reduction to increase as we approach the minimum wage from below.

Perhaps the most interesting pattern in the Tables is the tendency for the point estimates for above minimum wage regions up to \$2 to \$3 above the minimum wage to be negative. Only the \$2 to \$3 show up as statistically significantly different from zero (mainly at the 10% level of significance) but a test of the joint null hypothesis that all the coefficients on all the above minimum wage dummies up to and including the one for the \$2 to \$3 above minimum wage range are zero rejects the null at at least the 10% level in all the samples. The implication is that a minimum wage results in a reduction in the mass in the part of the wage distribution just above the minimum wage itself. We will return to discussing that result in greater detail below.

We do not present the long list of other estimated coefficients for brevity. The estimates indicate that the hazard is declining in education in virtually all the covariate segments (i.e., across the distribution), although in the very bottom covariate segment (corresponding to roughly the lowest 10% of our wage observations), the hazard for university graduates is actually slightly higher than that for workers with a post-secondary certificate or diploma below a bachelor's level degree. For the age effects, we observe a declining hazard up to age 54 but a higher hazard for the 55 to 60 year age group than the 35 to 54 year old age group in the lower part of the distribution. At and above the 7th covariate segment, however, the hazard declines continuously with age throughout the age range we consider. The year effects indicate declining mass in the 5 lowest covariate segments throughout much of our sample period but a mix of effects, with some increase in mass over time, in the upper segments. That is, over time, there appear to be fewer workers earnings wages below about \$14 per hour and more earning wages between about \$15 and \$20 per hour. The complexity of these year effects reveals one of the advantages of this estimator: that it allows for us to control for such effects in a flexible way while obtaining our minimum wage effect estimates.

In figures 1-4, we present plots of the wage hazards and densities for a specific person type with and without the minimum wage effects. We chose a person type with a distribution that is likely to be highly affected by the minimum wage (but who isn't simply a teenager, as is commonly done): a 25 to 34 year old high school dropout living in Ontario in November 2009. Figures 1 and 2 correspond to a female with these characteristics while Figures 3 and 4 correspond to a male. The minimum wage in Ontario in November of 2009 was \$9.50, or \$8.35 in the 2002 dollars in which we work.

Figure 1 contains the fitted hazard rate for a female of our specified type. Here one can see graphically what was presented in Table 2. There is virtually no probability of a worker receiving a wage below the minimum wage but a substantial probability of earning the minimum wage conditional on earning at least the minimum wage. Above the minimum, up to about \$11, the hazard rate is lower with a minimum wage of \$8.35. In fact, in the range above the minimum wage and below \$10, the hazard is reduced by approximately 25%. Thus, it appears that there are spillovers above the minimum wage. In earlier specifications, we allowed for effects of minimum wages on the hazard rate up to \$5 above the minimum wage. There, we found that the impact on the hazard rate was small (and not statistically significantly different from zero) but positive in the \$3 to \$4 and \$4 to \$5 above the minimum wage ranges. Thus, the pattern fits with workers who would have earned just above the minimum wage, having their wages increased to a greater extent than workers who would otherwise have earned below the minimum wage have their wages raised to just above the minimum.

In figure 2, we plot the wage density constructed using the estimated hazard rates displayed in Table 2. Strikingly, there is very little difference between the density with and without the minimum wage in the range between \$8.35 and about \$10. Above that point, the density for the case with the minimum wage is higher. Thus, the truncation spillover effect associated with the minimum wage exactly offsets the shape spillover effect just above the minimum wage and then becomes fully evident in the range where the effect of the minimum wage on the hazard is zero. The lack of an observable effect of the minimum wage on the density in the range just above the minimum fits with papers that have estimated a small or zero effect of minimum wages on percentiles of the wage distribution above the minimum wage (e.g., Dickens and Manning (2004a,2004b), Steward(2012)). However, the hazard rate effects indicate that there are spill-over effects - but there are two kinds and they are offsetting. To the extent that we are interested in using variation in the minimum wage to learn about the functioning of labour markets, the result that there is a shape spillover effect as revealed in the hazard is of interest. To the extent we care about impacts on the dispersion of the actual wage distribution, though, the early conclusions that minimum wages have limited effects of this type (and, thus, reduced effects on wage inequality relative to a situation with more of a spill-over) continue to hold.

The truncation spillover effect can arise in the context of many models of the labour market to the extent that a minimum wage induces a reduction in employment. In addition, Butcher et al(2012) argue that a monopsony style model can also induce truncation spillover effects even in the absence of reductions in employment. In our context, we can represent employment effects by examining the unconditional (on employment) wage distribution. In such a distribution, we could represent non-workers as contributing to a mass point at a zero wage. We can then break the estimation down into two parts: the effect of the minimum wage on the height of the mass point at zero (the employment rate effect) and the effect on the wage distribution conditional on working. Brochu and Green(2013) present results from estimating the impact of minimum wages on the employment rate for male and female, 15 to 59 year olds who are non-students and whose highest level of education is high school graduation or less. The specification is a standard one, including province and year fixed effects as well as the real minimum wage. They find that a 10% increase in the minimum wage causes a statistically insignificant 0.5% decrease in the employment rate for this overall sample (though a 2.5% decline for teenagers alone). This implies that the truncation effects that affect our wage density estimates are unlikely to arise from disemployment effects and, instead, may be related to the types of readjustments made by firms across employed workers emphasized in Butcher et al(2012).

Figures 3 and 4 contain the hazard and density plots for males. The hazard plot shows the same overall pattern as for females but with a smaller spike at the minimum wage and smaller sized effects in other parts of the distribution as well (notice the difference in scales on the vertical axes in the female versus male distributions). This is simply a reflection of the fact that there is more mass in the underlying female wage distribution around the point where the minimum wage bites. But in contrast to the female results, the truncation spillover effects are not large enough to fully offset the shape spillover effects evident from the hazard, with the result that the male wage density with the minimum wage is lower than the density without the minimum wage in the range stretching up to about 3$ above the minimum wage. It is possible these different truncation spillover effects by gender relate to different sized disemployment effects or, as in the Butcher et al(2012) discussion, because the minimum wage has a different "bite" in each gender's wage distribution. We intend to examine this issue further in the next iteration of the paper.

The figures for each of the three subgroups (Joiners, Leavers, and Stayers) showing hazard rates and densities with and without the minimum wage are very similar to those for the full sample, just presented. In Figures 5 through 7, we show comparisons of the fitted hazard rates for our base person type female in each of our subgroups. Thus, Figure 5 plots the hazard rate with the minimum wage at \$8.35 for Joiners and Stayers. The plots show a similar lack of mass below the minimum wage for both groups but a much larger spike at the minimum wage for the Joiners and more mass for Joiners between the minimum wage and about 14$. That is, not surprisingly, new hires enter at lower wages than those continuing in employment. The new information here is the extent of the hiring at the minimum wage: between 15 and 20$ of low educated female workers aged 25 to 34 start work at the minimum wage.

Figure 6 contains similar plots but for Leavers and Stayers. Leavers and Stayers have similar sized spikes in their hazards at the minimum wage but Leavers are more likely to have wages between the minimum wage and about $14. Thus, Leavers tend to be lower earning workers but there is no particular extra tendency to let go workers earning at the minimum wage.

Figure 7, showing hazards for Leavers and Joiners, completes the circle. Here we can see, again, the Joiners are disproportionally concentrated at the minimum wage, but the differences in the hazards for the two groups are otherwise very similar. That is, Leavers and Joiners are equally likely to be low wage earners, setting aside the much higher probability of Joiners earning the minimum wage.

Figures 8 through 10 contain the same set of comparisons but for males. The patterns are the same as for females, but applied to parts of wage distributions with less mass, implying smaller effects in absolute value.

At the start of 2009, the minimum wage in Ontario was $8.75 (or $7.70 in 2002 dollars) and then rose to $9.50 ($8.35 in 2002 dollars) in March of that year. In Figures 11 through 13, we present fitted densities for our base type female at each of these minimum wages for our three subgroups. Figure 11 contains this plot for Joiners. We can again see the large spikes at the minimum wage, though this spike becomes much larger as we move to the new minimum wage. Above the new minimum wage, the two densities are extremely similar. This is not just an artifact of examining the densities - the hazard rates (not shown here for brevity) are also very similar above the new minimum wage. The same pattern of an almost complete (though not quite 100%) drop in the hazard between the old and new minimum wage, a substantial increase in the spike at the hazard, and similar densities above the minimum wage are evident in the plots for both the Leavers (Figure 12) and Stayers (Figure 13).

One natural question is whether the mass of workers between the old and new minimum wage moves to or above the new minimum wage. Stayers make an interesting group for examining this question since they, by definition, are the same set of people before and after a minimum wage change. In Figure 14, we plot the fitted CDFs under the two minimum wages for our base type female (i.e., the integration of the density in figure 14). The CDF plots imply that about 11% of these workers earn between the old and new minimum wage when the old minimum wage was in force. After the increase, about half of those workers end up with wages at or below the new minimum wage while the other half earn wages above the new minimum wage. The fact that some of these inter-minimum range workers are shifted above the new minimum wage does not fit with models with truncation spillovers only. It does, however, fit with models such as Teulings model or with models based on fairness comparisons.

7 Conclusion

At this point, our estimates do not allow for dynamic effects of minimum wage changes. Since our estimator is essentially a fixed effect estimator, our estimates reflect both the short term effects of a minimum wage change (e.g., the lay-offs that happen at the time of a minimum wage change) and longer term, equilibrium effects (the extent to which the type of workers laid off in a higher minimum wage equilibrium are different from those laid off in a lower minimum wage equilibrium). We are currently waiting for the results from a dynamic specification to be release. Thus, at this point, our conclusions are somewhat limited. The key conclusions at this point are as follows.

First, that there is clear evidence of a spike in the hazard at the minimum wage. The spike, as others have pointed out, rules out simple forms of competitive models and points toward models with frictions. Second, there is a reduction in the hazard above the minimum wage.The reduction in the hazard fits with models in which minimum wages induce increases in wages for workers above the minimum wage, without enough formerly sub-minimum wage workers promoted into that wage range to maintain the sizes of the hazard rates. Thus, the models of interest are ones in which minimum wages have complex effects in raising wages of non-minimum wage workers. This is supported by our investigation of a minimum wage change for Stayers. For that group (who face no disemployment or selection effects by definition), we observe that for workers earning between the old and new minimum wage about half end up at the new minimum wage and half get increases that push them above the new minimum wage. To the extent that these workers are moved up to just above the new minimum wage, observed decreases in the hazard in that range must understate the extent to which workers who were in that range before the minimum wage change obtained wage increases of their own. Third, once we move from hazard rates to densities, there is much less (or no) evidence of a spillover effect. In essence, the inflating truncation effect offsets the shape spillover effects in the hazard for this range. This explains the difference between our findings and those in earlier investigations. Whether the changes in the hazard or the lack of change in the density is of interest depends on the question being asked. Our results may not change outlooks on the impact of minimum wage changes on inequality but could be useful for thinking about different models of how wages are set.

8 References

Autor, David, Alan Manning, and Christopher Smith (2010). The Contribution of the Minimum Wage To US Wage Inequality over Three Decades: A Reassessment Centre for Economic Performance, Working Paper No. 1025.

Brochu, Pierre and David A. Green (2013). The Impact of Minimum Wages on Labour Market Transitions, The Economic Journal, vol. 123, no. 3. Pp. 1203-35.

Butcher, Tim, Richard Dickens and Alan Manning (2012). Minimum Wages and Wage Inequality: Some Theory and an Application to the UK, Centre for Economic Performance, Working Paper No. 1177.

Dickens, Richard and Alan Manning (2004). Has the National Minimum Wage Reduced UK Wage Inequality? Journal of the Royal Statistical Society, Series A, Vol 167, Part 4, pp. 613-626.

Dickens, Richard and Alan Manning (2004). Spikes and Spill-overs: The Impact of the National Minimum Wage on the Wage Distribution in a Low Wage Sector, Economic Journal, March, C95-c101.

Donald, Stephen G., David A. Green and Harry 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), pp. 609-33.

Flinn, Christopher J. (2006). Minimum Wage Effects on Labor Market Outcomes Under Search, Matching, and Endogenous Contact Rates, Econometrica, vol. 74, No. 4, pp. 1013-1062.

Fortin, Nicole M., and Thomas Lemieux (2015) Changes in Wage Inequality in Canada: An Interprovincial Perspective forthcoming in the Canadian Journal of Economics, May.

Green, David A. and Harry J. Paarsch (1996). "The Effect of the Minimum Wage on the Distribution of Teenage Wages," UBC Economics Working paper, 97-02.

Lee, David S. (1999). Wage Inequality in the United States during the 1980s: Rising Dispersion or Falling Minimum Wage, Quarterly Journal of Economics 114, pp. 977-1023.

Lemieux, Thomas (2011). "Minimum Wages and the Joint Distribution of Employment and Wages," UBC Economics.

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

Meyer, Robert H. and David A. Wise (1983). The Effects of the Minimum Wage on the Employment and Earnings of Youth, Journal of Labor Economics, Vol. 1, No. 1, pp. 66-100.

Ryu, Kuenkwan (1992). Group Duration Analysis of the Proportional Hazard Estimator: Minimum Chi-Squared Estimators and Specification Tests, UCLA Department of Economics.

Stewart, Mark 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, Vol. 175, Part 1, pp. 263-287.

Teulings, Coen N. (2000). Aggregation Bias in Elasticities of Substitution and the Minimum Wage Paradox, International Economic Review, Vol. 41, No. 2, pp. 359-398.

Table 1: Estimated Hazard Coefficients for Minimum Wage Variables, Males

	All	Joiners	Leavers	Stayers
At Minimum wage	1.159	1.111	1.024	1.151
	[0.405]**	[0.522]**	[0.488]**	[0.395]**
50c or below	-2.384	-2.860	-1.869	-2.420
	[0.363]***	[0.482]***	[0.426]***	[0.360]***
30c to 50c below	-1.447	-1.608	-1.339	-1.430
	[0.370]***	[0.479]***	[0.434]**	[0.373]***
10c to 30c below	-1.230	-1.423	-0.720	-1.237
	[0.328]***	[0.437]***	[0.427]*	[0.323]***
10c to 30c above	-0.155	-0.286	0.0700	-0.156
	[0.313]	[0.419]	[0.393]	[0.309]
30c to $1 above	-0.410	-0.581	-0.528	-0.392
	[0.383]	[0.503]	[0.479]	[0.374]
$1 to $2 above	-0.310	-0.463	-0.325	-0.302
	[0.263]	[0.341]	[0.310]	[0.262]
$2 to $3 above	-0.199	-0.327	-0.253	-0.190
	[0.102]*	[0.145]**	[0.144]*	[0.102]*
$3 to $4 above	0.0286	-0.0332	0.0288	0.0343
	[0.107]	[0.132]	[0.134]	[0.105]
Observations	2505779	1989081	1108234	2481710

Notes: Sample excludes those earning tips. Standard errors (clustered by province) in brackets. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 2: Estimated Hazard Coefficients for Minimum Wage Variables, Females

	All	Joiners	Leavers	Stayers
At Minimum wage	1.121	1.203	1.087	1.098
	[0.339]***	[0.452]**	[0.478]**	[0.331]***
50c or more below	-2.758	-3.127	-2.416	-2.783
	[0.297]***	[0.365]***	[0.424]***	[0.298]***
30c to 50c below	-1.786	-1.857	-1.359	-1.808
	[0.346]***	[0.403]***	[0.449]**	[0.342]***
10c to 30c below	-1.479	-1.612	-0.699	-1.528
	[0.284]***	[0.372]***	[0.393]*	[0.279]***
10c to 30c above	-0.0829	-0.0854	0.308	-0.0964
	[0.257]	[0.327]	[0.395]	[0.254]
30c to $1 above	-0.291	-0.346	-0.111	-0.301
	[0.300]	[0.413]	[0.408]	[0.295]
$1 to $2 above	-0.240	-0.276	-0.0590	-0.250
	[0.176]	[0.245]	[0.268]	[0.173]
$2 to $3 above	-0.120	-0.145	-0.0554	-0.120
	[0.0639]*	[0.0725]*	[0.100]	[0.0653]*
$3 to $4 above	0.0202	0.0420	0.0759	0.0171
	[0.0839]	[0.0986]	[0.115]	[0.0849]
Observations	2288017	1635961	916421	2265580

Notes: Sample excludes those earning tips. Standard errors (clustered by province) in brackets. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

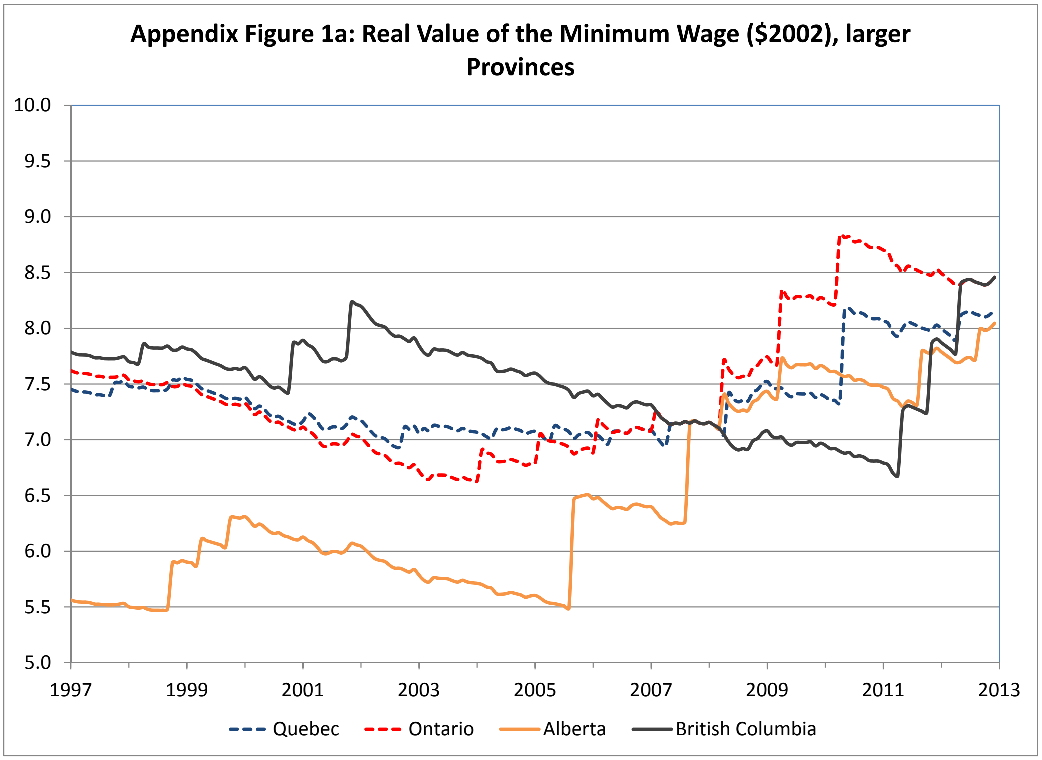
Appendix Figure 1a: Real Value of the Minimum Wage ($2002), larger Provinces
10.0
9.5
9.0
8.5
8.0
7.5
7.0
6.5
6.0
5.5
5.0
1997
1999
2001
2003
2005
2007
2009
2011
2013
Quebec
Ontario
Alberta
British Columbia

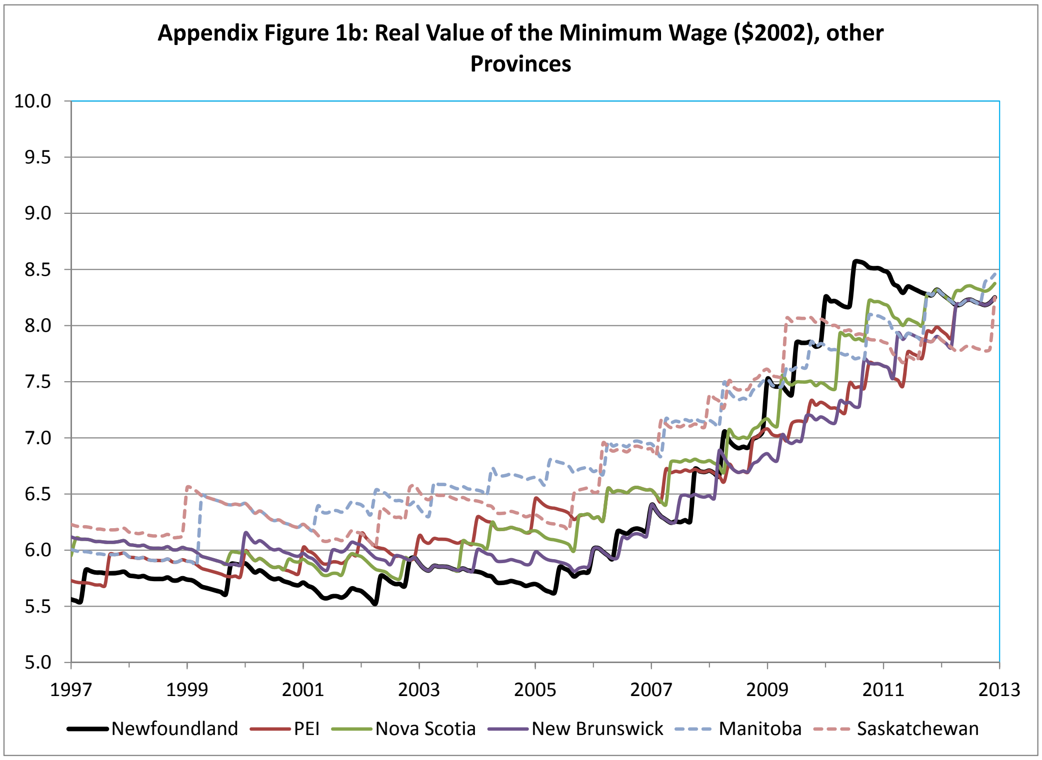
Appendix Figure 1b: Real Value of the Minimum Wage ($2002), other Provinces
10.0
9.5
9.0
8.5
8.0
7.5
7.0
6.5
6.0
5.5
5.0
1997
1999
2001
2003
2005
2007
2009
2011
2013
Newfoundland
PEI
Nova Scotia
New Brunswick
Manitoba
Saskatchewan

Figure 1:
Wage Hazard Function With and Without the Minimum Wage
All Earners, Females

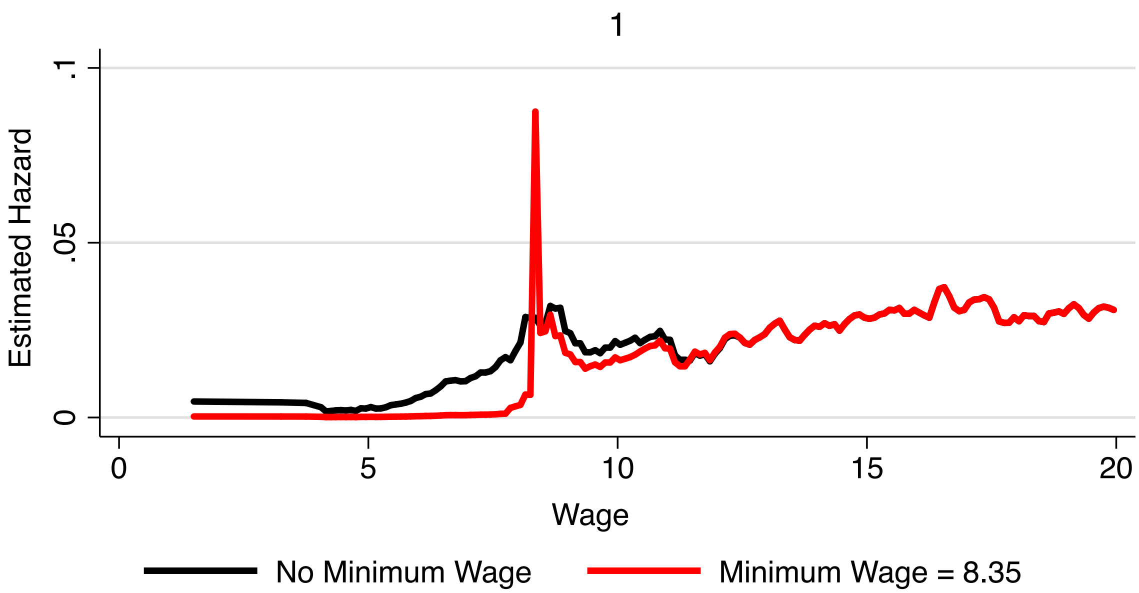

The selected type of worker is female, 25 years old, with less than High School, living in Ontario during November-2009.

Figure 2:
Wage Densities With and Without the Minimum Wage
All Earners, Females

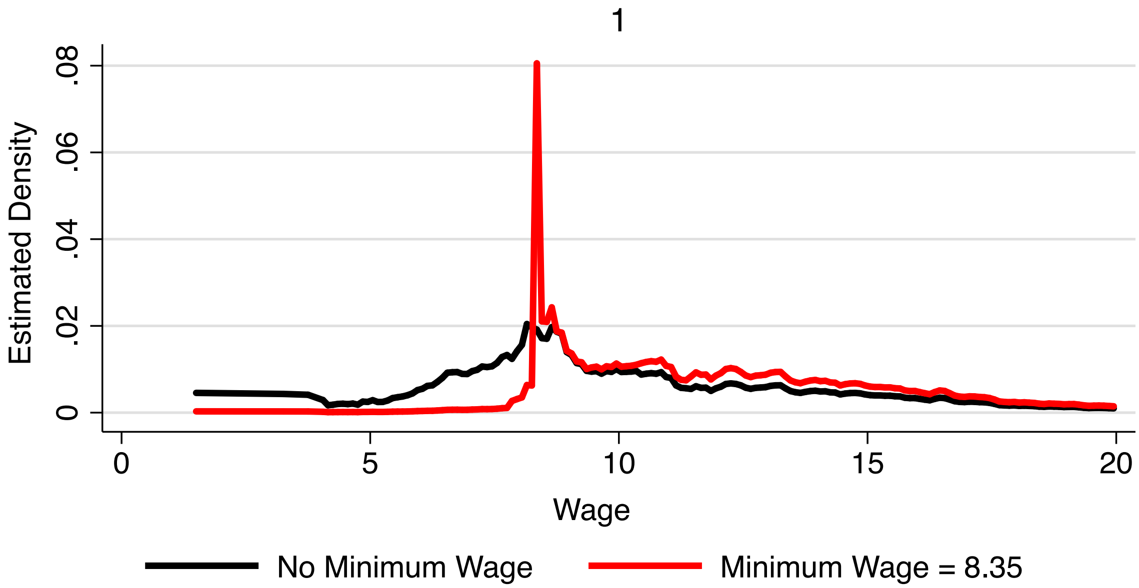

The selected type of worker is female, 25 years old, with less than High School, living in Ontario during November-2009.

Figure 3:
Wage Hazard Function With and Without the Minimum Wage
All Earners, Males

The selected type of worker is male, 25 years old, with less than High School, living in Ontario during November-2009.

Figure 4:
Wage Densities With and Without the Minimum Wage
All Earners, Males

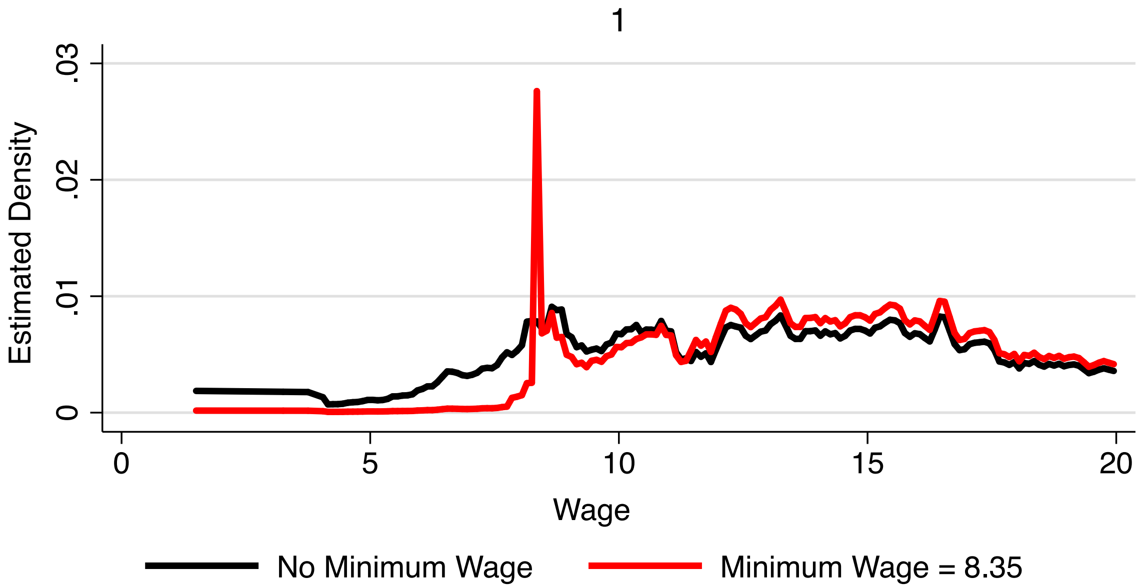

The selected type of worker is male, 25 years old, with less than High School, living in Ontario during November-2009.

Figure 5:
Wage Hazard Functions for New Hires and Stayers,
Females

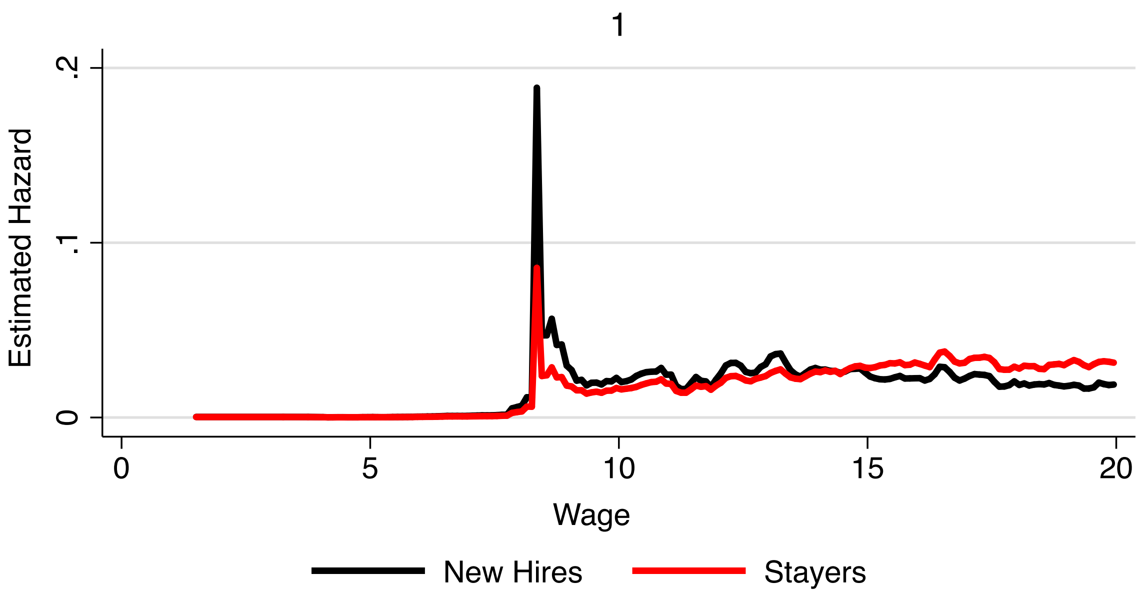

The selected type of worker is female, 25 years old, with less than High School, living in Ontario during November-2009. Minimum Wage=8.35

Figure 6:
Wage Hazard Functions for Leavers and Stayers,
Females

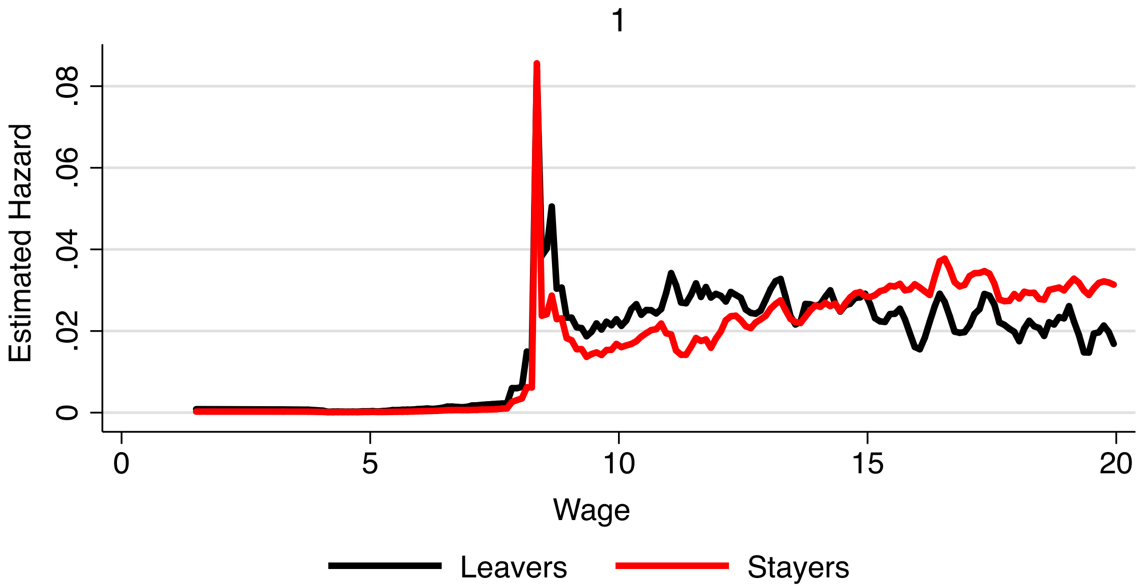

The selected type of worker is female, 25 years old, with less than High School, living in Ontario
during November-2009. Minimum Wage=8.35

Figure 7:
Wage Hazard Functions for New Hires and Leavers, Females

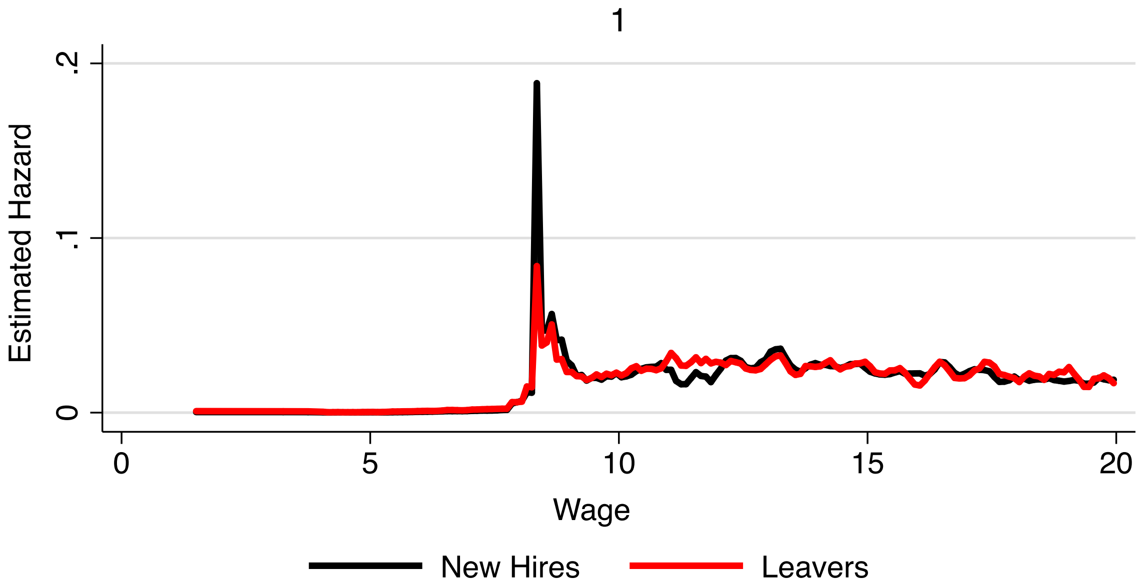

The selected type of worker is female, 25 years old, with less than High School, living in Ontario during November-2009. Minimum Wage=8.35

Figure 8:
Wage Hazard Functions for New Hires and Stayers, Males

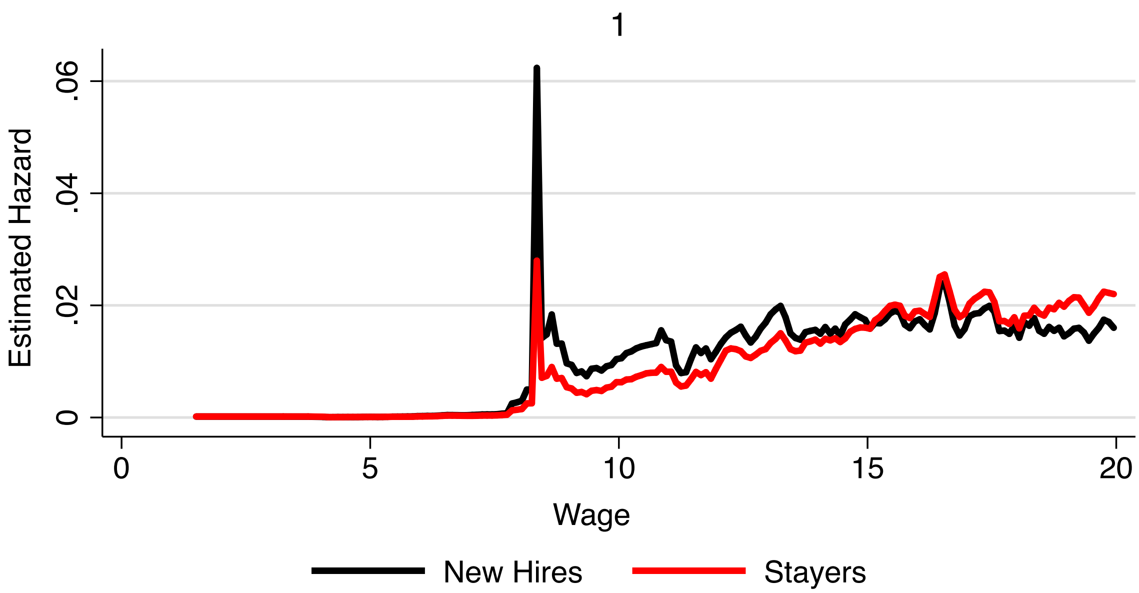

The selected type of worker is male, 25 years old, with less than High School, living in Ontario during November-2009. Minimum Wage=8.35

Figure 9:
Wage Hazard Functions for Leavers and Stayers,
Males

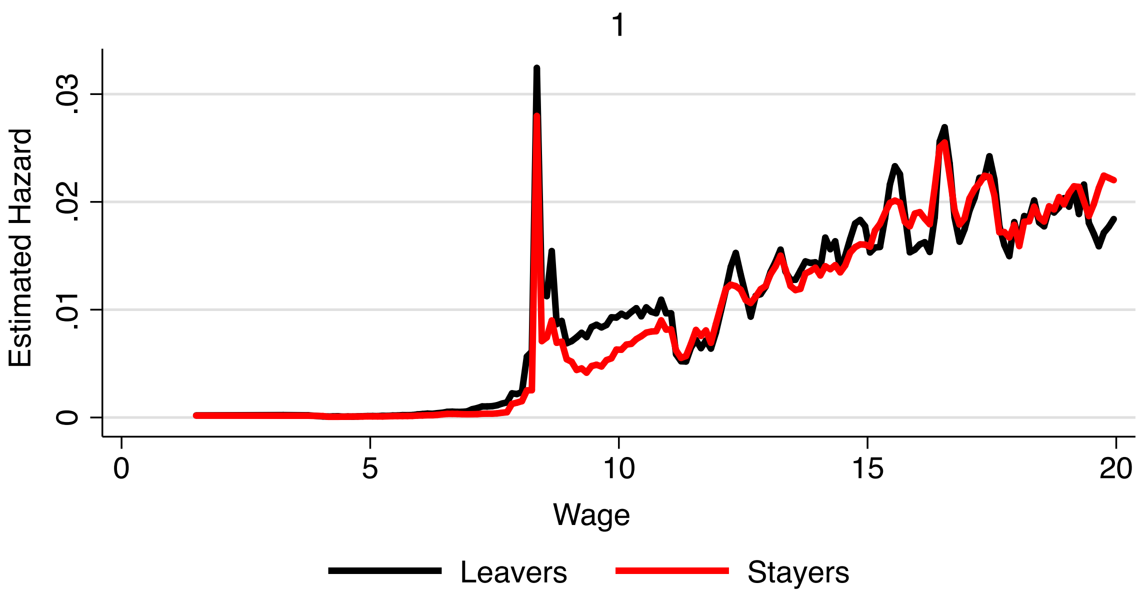

The selected type of worker is male, 25 years old, with less than High School, living in Ontario during November-2009. Minimum Wage=8.35

Figure 10:
Wage Hazard Functions for New Hires and Leavers, Males

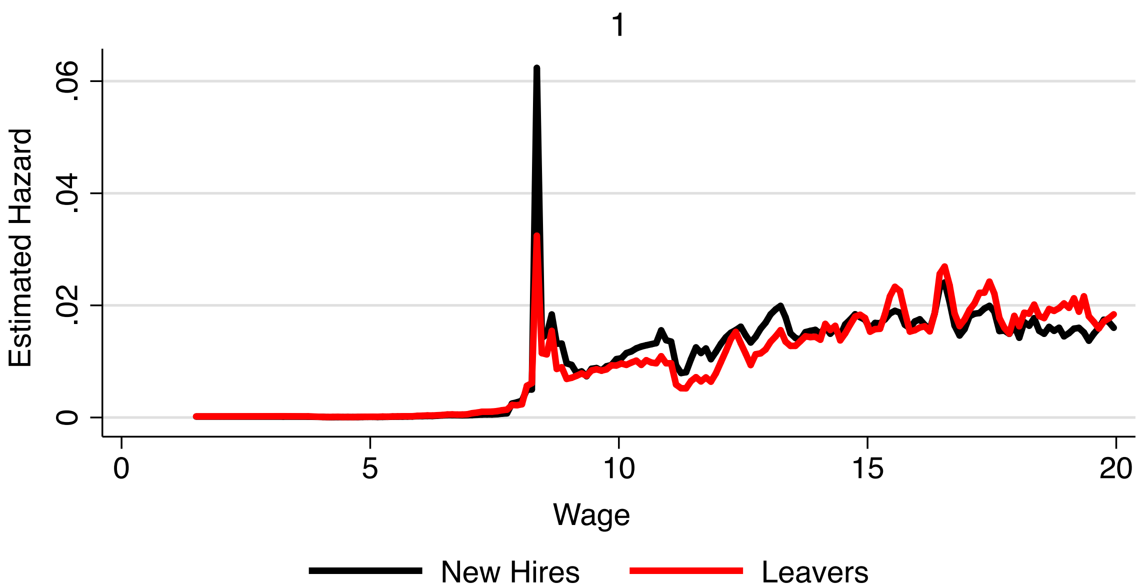

The selected type of worker is male, 25 years old, with less than High School, living in Ontario during November-2009. Minimum Wage=8.35

Figure 11:
Wage Densities for New Hires
Minimum Wages at 7.70 and 8.35, Females

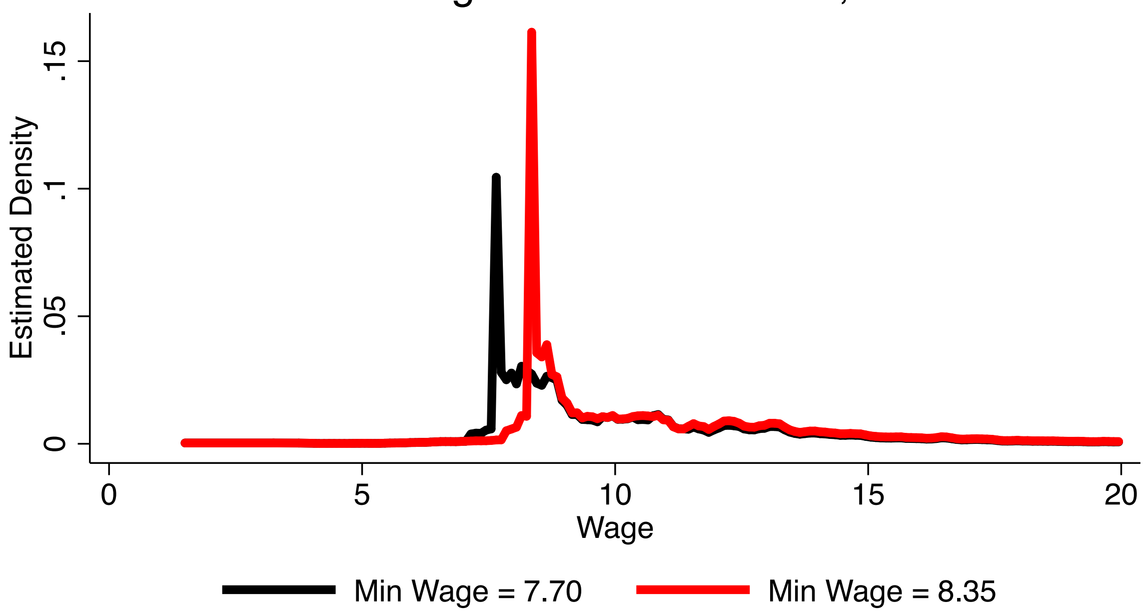

The selected type of worker is female, 25 years old, with less than High School, living in Ontario during November-2009.

Figure 12:
Wage Densities for Leavers
Minimum Wages at 7.70 and 8.35, Females

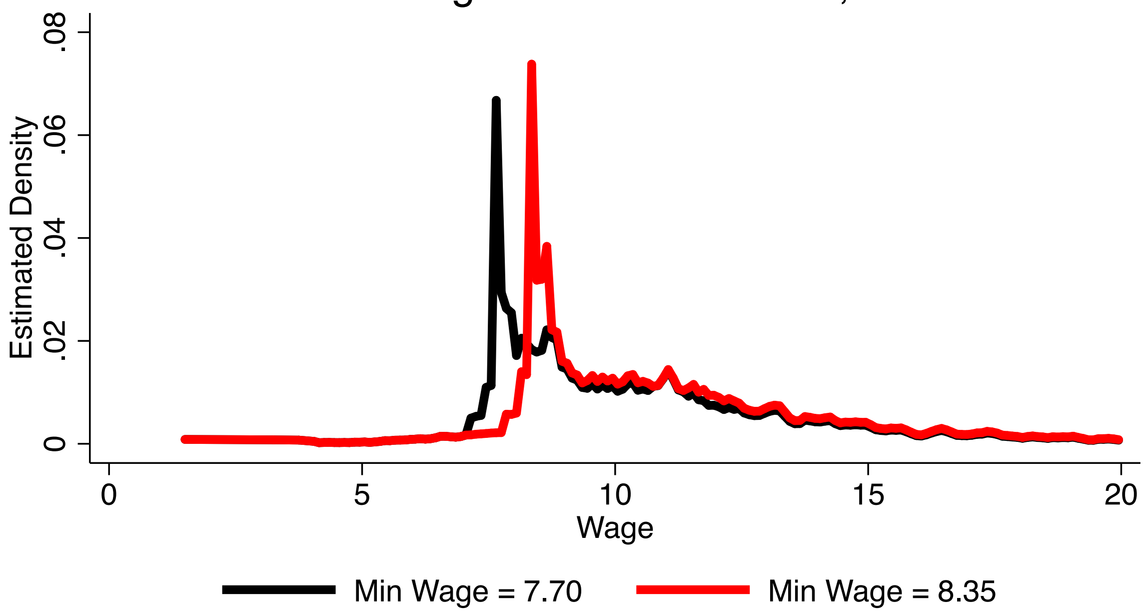

The selected type of worker is female, 25 years old, with less than High School, living in Ontario during November-2009.

Figure 13:
Wage Densities for Stayers
Minimum Wages at 7.70 and 8.35, Females

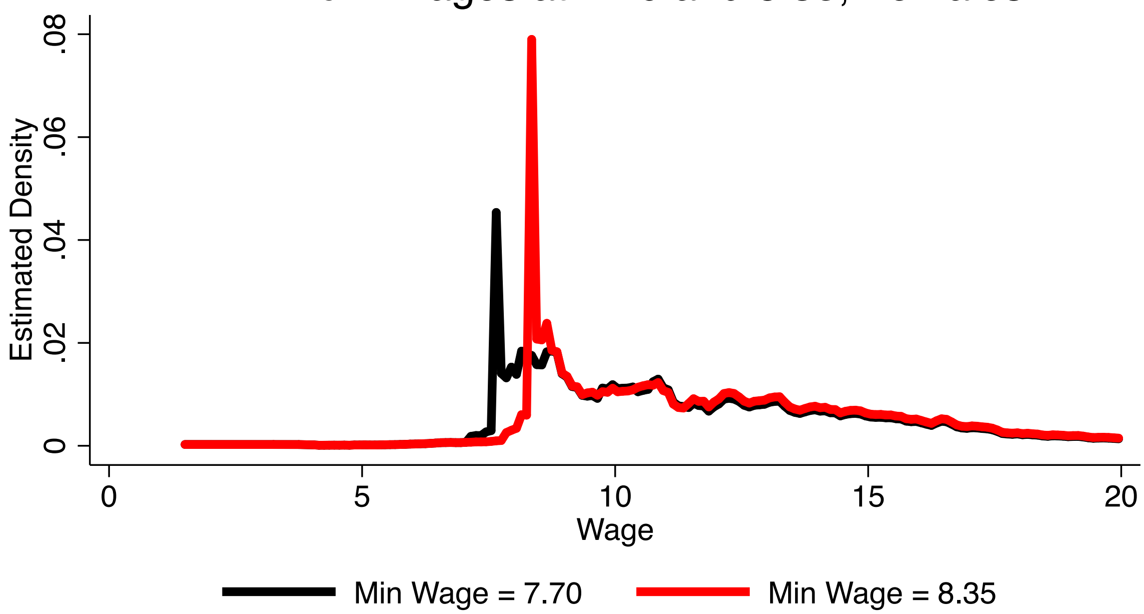

The selected type of worker is female, 25 years old, with less than High School, living in Ontario during November-2009.

Figure 14
Wage Distributions For Stayers
Minimum Wages at 7.70 and 8.35, Females

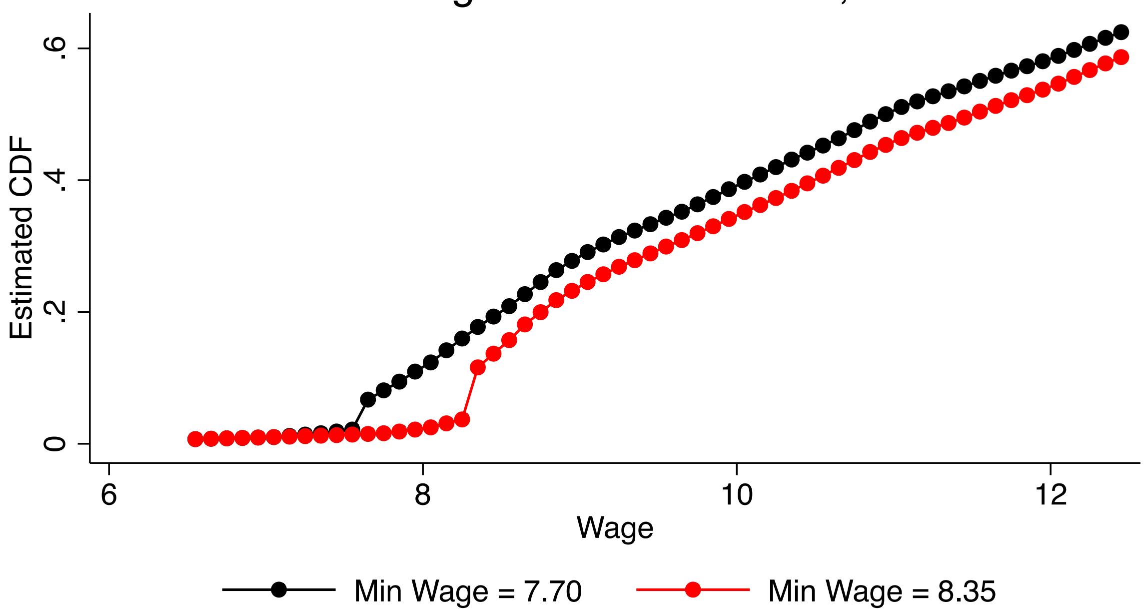

The selected type of worker is female, 25 years old, with less than High School, living in Ontario during November-2009.