

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

Pierre Brochu,

University of Ottawa

David A. Green,

University of British Columbia

Thomas Lemieux,

University of British Columbia

James Townsend

University of Winnipeg

January 2 2018

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 substantial spillover effects of the minimum wage, and significant differences across the three groups.

Acknowledgments: We would like to thank Iain Snoodly and Daniel Gao for expert research assistance, and SSHRC and the Bank of Canada Fellowship Award for financial support.

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, 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). Another advantage of the hazard setup is that it models the effect in the minimum wage in a proportional way, which is consistent with the view that the impact of the minimum wage depends on the fraction of workers affected. In an attempt to deal with the problem of measurement error in hourly wages (see, e.g., Autor, Manning, and Smith, 2016), we also explicitly model the probability that workers report being paid an integer hourly wage value (e.g. \$10 an hour). Consistent with existing evidence for the United States (Lee, 1999) and Canada (Fortin and Lemieux, 2015) we find evidence of substantial spillover effects up to about \$2 above the minimum wage.

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, $F(\cdot)$. We are interested in how $F(\cdot)$ 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 start a new job in $t+1$ (i.e., who either moved from non-employment in t to employment in $t+1$ or who are job to job switchers between t and $t+1$), 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:

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

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

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

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

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

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

$$\begin{aligned}
& F_{t+1}(w|m_1) - F_t(w|m_0) = & (3) \\
& +(Pr(J(t, t+1|E(t+1) = 1, m_1) - Pr(L(t, t+1|E(t) = 1, m_0)) \\
& \quad \cdot [F_t(w|L(t, t+1, m_0) - F_t(w|S(t, t+1) = 1, m_0)] \\
& + [F_{t+1}(w|J(t, t+1) = 1, m_1) - F_t(w|L(t, t+1) = 1, m_0)] \\
& \quad \cdot Pr(J(t, t+1|E(t+1) = 1, m_1) \\
& + (F_{t+1}(w|S(t, t+1) = 1, m_1) - F_t(w|S(t, t+1) = 1, m_0)) \\
& \quad \cdot Pr(S(t, t+1|E(t+1) = 1, m_1)
\end{aligned}$$

Thus, changes in the CDF can be broken into three components: 1) the difference in the 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 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).

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 of Joiners and Leavers will be equal, the wage distribution for Stayers 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

¹The 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,

$$\begin{aligned}
F_{t+s}(w|m_1) - F_t(w|m_0) &= (Pr(S(t+s-1, t+s)|E(t+s) = 1, m_1) \tag{4} \\
&\quad - Pr(S(t-1, t)|E(t) = 1, m_0)) \cdot F_t(w|S(t-1, t) = 1, m_0) \\
&+ (F_{t+s}(w|S(t+s-1, t+s) = 1, m_1) - F_t(w|S(t-1, t) = 1, m_0)) \cdot Pr(S(t+s-1, t+s)|E(t+s) = 1, m_1) \\
&+ Pr(J(t+s-1, t+s|E(t+s) = 1, m_1)) \cdot [F_{t+s}(w|J(t+s-1, t+s) = 1, m_1) - F_t(w|J(t-1, t) = 1, m_0)] \\
&+ (Pr(J(t+s-1, t+s|E(t+s) = 1, m_1) - Pr(J(t-1, t|E(t) = 1, m_0))) \cdot F_t(w|J(t-1, t, m_0)
\end{aligned}$$

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 equilibria, 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.

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 proportions 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 have 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 . That means that if we think about the change in the overall wage distribution, there will again be a truncation effect inflating the density above m_1 because of the layoffs, but in Teulings' simulations there is also a shape spillover effect in the form of the extra mass just above m_1 . This, in turn, implies that, in contrast to the simple model, the wage hazard rates above m_1 will change. In particular, the hazard rates just above m_1 will increase. Working with hazard rates would allow us to see that effect more clearly, differentiating it from the general inflation of the density associated with the increased truncation.

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). What happens to the Stayers' wage distribution is less clear. For workers who formerly earned just above m_1 , the increase in demand for their skills will result in wage increases. Thus, if we focused just on those workers, their wage hazard rates just above m_1 would fall since they are moving to higher wages. But, at the same time, workers whose wages were formerly just below m_1 will also get wage increases that will move them into the range just above the new minimum. The net effect is uncertain theoretically, though in Teulings (2000)'s simulations, the latter effect dominated, implying that a minimum wage induces extra mass just above m_1 .

It is interesting to think of the wage distribution changes experienced by each sub-

group 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. The key difference relative to the simple competitive model is that Butcher et al. (2012) assume that there are no aggregate employment effects: the workers who were formerly at firms with productivity below the minimum wage will be reallocated to other firms in accordance with the labour supply functions faced by each firm. That means that the inflation of the density above the minimum wage is not just a re-scaling but corresponds to an actual increase in workers at each above-minimum firm. How much that happens depends on how many of the formerly below minimum firm workers end up employed right at the minimum, i.e., how much the set of firms most affected by the minimum wage expand their employment.

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.

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.²

²As a side point, this model has similarities to the employment implications in Butcher et al(2012). In particular, the shutting down of firms with productivities below the new minimum wage will throw some workers into the unemployment pool. That makes it easier for vacancies to be filled, resulting in more vacancies and the hiring of some of the laid off workers. This will go part way to mitigating any negative employment effects, though not as far as in Butcher et al. (2012), where all the workers from firms that shut down are assumed to be re-employed.

At first glance, there is little to separate Butcher et al. (2012) from Flinn (2006) in terms of implications. But Aaronson and French (2007) present evidence that firms pass minimum wage increases through to consumer prices. They argue that this does not fit with monopsony models since those imply increases in employment which would, in turn, imply increased production and decreased consumer prices. However, that evidence could fit with a Flinn-type model since the minimum wage causes some firms to shut down while the ones that remain in business face a change in their bargained wage but not necessarily in their production.³

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) should increase. Thus, we could see declines in the hazard rate for Stayers just above m_1 in the short run.

The implications in terms of workers grouped by turnover status have interesting echos in the work of Aaronson, French and Sorkin (2013) who show examine employment of entering and exiting firms at the time of a minimum wage increase. Using US data on the restaurant industry, they show that there is an increase in both firm exits and firm entry at the time of a minimum wage increase. They explain this using a putty-clay model of technological choice in which firms can choose among different capital intensity options when they first open but cannot alter their production technology thereafter. This means that adjustments to a minimum wage increase must operate through the closing and entry of firms and that short run effects will be smaller than longer run effects. Mapping to our data, the workers in the closing firms would be our Leavers and they should come from firms with relatively low capital intensity and, therefore, low wages. Employment in the entering firms would show up as Joiners in our data, and they

³Aaronson and French(2007) argue their results fit with a competitive market model but have to resort to a spike in the skills distribution whereby all low skilled workers are paid the same wage in order to generate a spike in the wage distribution.

would tend to be higher paying firms. Our Stayers would show little or no effect of the minimum wage apart from moving workers formerly paid below m_1 to the spike. That is, there should be no effect on their hazard above m_1 .

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). 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 Kahneman et al. (1986) 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 Stewart (2012) 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 (2016) concludes that spillover effects are substantially smaller than 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 (2016) also use variation in state minimum wages to estimate spillover effects. This can be done starting in the late 1990s as an increasing amount of variation in the minimum wage is now happening at the state level. 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. Fortin and Lemieux find some evidence of spillover effects, though those don't tend to go above the 15th percentile of the wage distribution for all workers. This remains substantial as the minimum wage was only around the 5th percentile of the wage distribution during the period they study (1997-2013).

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 (2016).

One shortcoming of these studies is that they look at the impact of the minimum on various quantiles of the wage distribution without explicitly modelling how the minimum wage affects wages at a given distance from the minimum wage. By modelling the impact of the minimum wage on the probability (or hazard) of earning wages at precise points of the distribution, we are able to recover an underlying mapping between the minimum wage and the probability of earning wages right at the value of the minimum wage, as well as at different points above and below the minimum wage. A related advantage of our approach is that, as argued by Fortin, Lemieux, and Firpo (2011) and Chernozhukov, Fernandez-Val, and Melly (2013), it is generally easier to econometrically model the probability distribution than quantiles.

Another 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. We have witnessed a direct example of this in the re-scaling of the density above the minimum wage associated with the truncation of the wage distribution in the simple competitive model. 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 generally 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. The only partial exception to these statements of which we are aware is Neumark et al. (2004). They use matched CPS data to examine wage and employment effects of minimum wage changes conditioning on the initial wage at which a worker is observed. Their estimated wage effects are for the set of workers who stay on the same job between the two surveys and so match what we call wage effects for Stayers. Their employment effects show what part of the initial wage distribution the workers who lose their jobs because of a minimum wage come from and, thus, match our Leavers results. They do not have estimated effects for Joiners and they do not, as we will, examine effects in terms of hazard rates. They present their results in terms of initial wage ranges defined in terms of multiples of the minimum wage and show that wages increase by 80% of a minimum wage increase for

wages initially between the original minimum wage and 1.1 times that wage. Wage increases decline as one moves up the wage range. Unfortunately, the authors never present figures showing the ultimate effect on the wage distribution and it is difficult to map from these wage changes to distribution effects.⁴ 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 formerly 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, for example. 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 in-

⁴If we think about the up to 1.1 category moving up 80% of the minimum while the 1.1 to 1.2 category moves up 41% (as their results indicate) then there could be more moving into the 1.1 to 1.2 category than those moving out but it depends on how far each moves. So, the hazard could increase or decrease in that category.

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,

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

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.

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. This is useful in the case of the minimum wage where we expect the minimum wage impact to depend in a roughly proportional way on the the fraction of workers affected by the minimum wage. 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:

$$y_0 < y_1 < \dots < y_P \quad (7)$$

The survivor function then becomes

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

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 within a given segment is given by

$$h(y|x) = \exp(x\alpha_p) h_0(y). \quad (9)$$

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

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

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, or wage bins, ($P = 164$) but restricting $\alpha_p = \alpha_s$ within each of $s = 1, \dots, S$ non-overlapping sets of baseline segments. We will call the sets within which the covariate effects are constant "covariate segments" and will allow for 5 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 rich set of province and year controls (more on this below). 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.

We consider several specifications for province and year effects, starting with a tra-

ditional difference-in-differences approach where a set of province and year dummies are included. In light of the large differences in real wage growth across provinces over time (Fortin and Lemieux, 2015, Green, Morissette, and Sand, 2016), we also estimate models with province-specific linear and quadratic trends. The most flexible specification we consider includes a full set of province-year year dummies within each of the five covariate segments. Identification is still achieved because the minimum wage affects differently the probability/hazard of being in each wage segment. Thus, we can think of our identification strategy as a “triple-differences” design. In addition to the full set of interactions between province and year effects, we also include interactions between the wage level (value of the wage segment cutoffs) and province and year dummies. This allows the scale parameter (i.e. the variance of the latent wage distribution) to differ across provinces and year. With this most general specification we have a conventional “triple-differences” design where province-year, wage segment-province, and wage segment-year interactions are controlled for, and identification is obtained using variation in the minimum wage impact at the wage segment-province-year level.

Appendix Figure 1 illustrates how the model is identified using an example where the baseline hazard is constant, and spillover effects go up to \$2 above the minimum wage. Consider increasing the minimum wage from \$8 to \$9. Since the hazard rate above \$11 is unaffected by the change in the minimum wage, it implicitly acts as a control in the triple-differences design. Spillover effects between \$ 1 to \$ 2 above the minimum are identified by comparing the hazard in the \$10 to \$11 range under the \$ 9 (light red bars) and \$8 (dark red bars) minimum wages. Likewise, the difference between the hazard in the \$9 to \$10 range under the \$9 and \$8 minimum wage identifies the difference in spillover effects up to \$1 above and between to \$1 and \$2 above the minimum wage. As we keep moving to the left of the distribution, a similar argument can be used to understand how the spike at the minimum wage and the (large) negative effect on the probability of earning less than the minimum are also identified using a recursive argument.

Two further points about identification are of interest. First, Stewart (2012) showed that using high percentiles (e.g., the 50th percentile) as controls could be problematic because they appear to be following different trends from the lower percentiles in his UK data. Our deals with this by allowing 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 (\$2 or \$3 above the minimum wage in most specification we present below). Second, we could, in principle, allow for the minimum wage to affect the hazard rate higher up in

the wage distribution. As discussed in the results section below, we have tried such specifications and generally found no statistically significant or economically substantial effects beyond \$2 above the minimum. Imposing the restriction of no effects beyond a certain point generally yields more precise estimates given the recursive nature of the estimation.

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. (2012) 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 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.⁵ At time, 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 mainly focus on the general adult minimum wage for each province.⁶ To exploit within-year variation in the minimum without unduly increasing the computational burden, we focus on quarterly frequencies. In particular, we use the highest minimum wage that was prevailing during a given quarter.

⁵There 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.

⁶We 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.

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, using 2002 as our base year. Appendix Figures 2a and 2b show the real minimum wage patterns by province and year. Importantly, the minimum wage shows considerable variation over time within each of the provinces.

The second source of data is the Canadian Labour Force Survey (LFS) master files, and in some cases (see below) the public version of these files.⁷ 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. A key variable for our purposes 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 1976.⁸

We restrict our LFS sample to individuals aged 15 to 59. 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 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.⁹ However, we are constrained in some of our definitions by our need to use the incoming sample members

⁷The master files were accessed on site at the Manitoba and Carleton, Ottawa, Outaouais local Research Data Centre (COOL RDC).

⁸See Brochu (2006) for a detailed discussion of the limitations of other North American data sets.

⁹A 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 either non-employed or employed with only 1 month of job tenure 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. (2016)'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. (2016), 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 samples in an appendix (not available in this version of the paper).

We also explicitly model the probability of reporting an integer wage value to allow for measurement error in the form of rounding off the true value of the wage. This form of measurement error is likely important as around 30% of workers report an integer value of the wage, which sounds implausibly high. We model this phenomenon by creating a dummy variable equal to 1 when an integer value is included in the wage segment, and 0 otherwise. In some specifications we also include an additional dummy for \$10, as well as interactions between the integer dummy and being at the minimum wage.¹⁰

¹⁰Given that the different effects are multiplied by each other under the proportional hazard specification, the relative size of the minimum wage spike is expected to be smaller when the minimum wage is itself an integer value.

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 5 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 about 20% of the wages within our wage range (i.e., the wages between \$0 and \$20) for the sample of all earners pooled across years. In practice, we get close to this figure by using a first covariate 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.

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 quarter. 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 29 cents above m ; 30 cents to 49 cents above m ; and 50 cents bands above that until reaching \$2.00 (and more in some specifications) 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. 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 as high as 745, depending on specifications.

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.¹¹

¹¹If 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

6 Results

In Table 1a and 1b, 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. This first set of results is for all workers (continuers, leavers and joiners combined) and based on public use files of the LFS where rotation groups and individual identifiers are not available. But although we cannot identifier leavers in the public use files, it is still possible to estimate the models for all workers combined.¹² Furthermore, an important advantage of working with public use files is that disclosure analysis doesn't need to be performed on the results, and more powerful computers can be used for the estimation than those currently available in the RDC. However, a shortcoming of the public use files is that the lack of information about the rotation group means that we cannot limit the analysis to "fresh wages" in month 1. We deal with this issue in the estimation by adding interactions between the minimum wage variables and a dummy for whether the last increase in the minimum wage took place over the last six month. Thus, the minimum wage effects reported in that case represent the medium-run effect of the minimum wage that arise at least 6 months after the minimum wage increase took place.

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

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.

¹²As discussed above, we define job leavers as those employed in month 1 but either non-employed or employed with only 1 month of job tenure in month 2. This cannot be done in the public use files where the information for longitudinally linking individuals is not available.

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 1a and 1b and then multiplying them by the hazard determined by a combination of the baseline hazard and the effects of the other covariates. Given that, it is more convenient to present results in the form of figures, but there are some immediate points that can be made examining these coefficients. The tables also indicate how different specification choices affect the estimated minimum wage effects.

The first three columns of Table 1 show the impact of using increasingly general specifications for province and year effects. While the spike at the minimum wage is always large and significant, we only get significant spillover effects right above (10 to 30 cents) the minimum wage in a traditional difference-in-differences specification where only province and year dummies are included (column 1). Spillover effects become significant up to \$1.50 above the minimum wage when either quadratic provincial trends (column 2) or a full set of province-year dummies (column 3) are included. Note that as the minimum wage always lies in the first covariate segment (less than \$10 in dollars of 2002), we only include the full set of province-year dummies in that segment in column 3 to reduce the computational burden.

We then add the dummy variable indicating whether there is an integer dollar value in the wage bin in columns 4 to 6. For both men and women, doing so makes spillovers between \$1.50 and \$2.00 above the minimum wage positive and significant, and very substantially improves the fit of the model. This suggests that, if anything, measurement error in the form of rounding off wages tends to understate the importance of spillover effects. Column 7 then shows that spillover effects between \$2 and \$3 are small and not statistically significant, and that including those makes the other estimated minimum wage effects more imprecise. We have also estimated models with spillover effects going up to \$5 above the minimum wage. The estimated effects are reported in Figures 1a and 1b. For women and especially men, the estimated effects above \$2 tend to be much smaller (and often not significant, or only marginally significant) relative to spillover effects closer to the minimum wage. Thus, our preferred specification will limit spillover

effects to \$2 above the minimum wage. The last column of Table 1 (column 8) shows that adding a full set of province-year dummies in the second covariate segment has little impact on the results.

To help interpret the results, we translate the estimated minimum wage effects in percentage terms in Figures 2a and 2b by taking the exponential of the estimated coefficients in the hazard models reported in Figure 1a and 1b. We then overlay (red line) to these figures the cumulative distribution of wages averaged over all provinces and years, and recentered around the minimum wage for the sake of comparability across provinces and year. Finally we “shut down” the minimum wage effects to show the distribution (dotted red line) that would prevail in absence of any minimum wage effects.¹³ The results indicate that, for women, spillover effects go up to the 20th percentile of the wage distribution. In the case of men, spillovers appear to be more modest as they only go up to the 10th or 15th percentile of the distribution. Interestingly, this difference is entirely due to the fact the wage distribution of men is shifted to the right (fewer workers at low wages) relative to women. By contrast, the estimated minimum effect plotted using the blue bars are quite similar for men and women (as they are in Table 1a and 1b). The similarity of the estimates for men and women can be viewed as a (successful) specification test of our proportional hazard specification.

Although we cannot precisely identify leavers, joiners, and continuers in the public use files, Table 2 shows that the minimum wage effect tends to be larger for recently hired workers with less than one year of tenure. Table 3a and 3b show (preliminary) results for leavers, joiners, and continuers based on the master files of the LFS. The results indicate that, if anything, minimum wage effects tend to be larger for both leavers and joiners. In the case of leavers, it would be particularly interesting to know whether there is an abnormal concentration of workers between the current and an upcoming increase in the minimum wage. We are in the process of estimating such specifications but results are not available yet.

7 Conclusion

At this point, our estimates do not allow for dynamic effects of minimum wage changes that we just discussed. 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

¹³For the purpose of this counterfactual exercise, we set spillover effects above \$2 to zero.

(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 released. Thus, at this point, our conclusions are somewhat limited. The key conclusions at this point are nonetheless 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, we find substantial spillover effects that go to about \$2 above the minimum wage. This corresponds, on average, to about the 20th percentile of the wage distribution for women, and the 10th or 15th percentile for men. So although these effects do not reach out quite as high in the wage distribution relative to Lee (1999), our estimated spillovers appear substantially larger than those found by Autor et al. (2016).

Furthermore, in our proportional hazard model, the impact the minimum wage depends in a proportional way on the fraction of workers affected by the minimum wage. This suggests that upcoming minimum wage increase in Canada could have a large impact on the wage distribution. The minimum wage is set to reach \$15 in both Alberta and Ontario in less than a year from now. In the case of Ontario, \$15 currently represents the 25th percentile of the distribution. Once factoring in our estimated spillover effects, this suggests that the future impact of the minimum wage on the distribution may turn out to be substantially larger than what has been the case during our analysis period (1997-2016).

8 References

Aaronson, Daniel and Eric French(2007). "Product Market Evidence on the Employment Effects of the Minimum Wage," *Journal of Labor Economics*, Vol. 25, No. 1, pp. 167-200.

Aaronson, Daniel, Eric French and Isaac Sorkin (2013). "Firm Dynamics and the Minimum Wage: A Putty-Clay Approach," *Federal Reserve Bank of Chicago Working Paper No. 2013-26*.

Autor, David, Alan Manning, and Christopher Smith (2016). *The Contribution of the Minimum Wage To US Wage Inequality over Three Decades: A Reassessment*, *American Economic Journal: Applied Economics*, vol. 8, no. 1, pp. 58-99.

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*.

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

Chernozhukov, Victor, Ivan Fernandez-Val, and Blaise Melly (2013). *Inference on counterfactual distributions*, *Econometrica*, vol. 81, No. 6, pp. 2205-2268.

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.

Dube, Arindrajit, Laura Giuliano, and Jonathan Leonard (2015). "Fairness and Frictions: The Impact of Unequal Pay Raises on Quit Behavior," University of Massachusetts, Amherst.

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., Thomas Lemieux, and Sergio Firpo. (2011). Decomposition methods in economics. *Handbook of Labor Economics*, vol. 4A, 1-102.

Fortin, Nicole M., and Thomas Lemieux (2015) Changes in Wage Inequality in Canada: An Interprovincial Perspective forthcoming in the *Canadian Journal of Economics*, Vol. 48, no. 2, 682-713.

Green, David A., René Morissette and Benjamin M. Sand (2016) Geographic Spillovers of Booms: The Effects of Canada's Resource Boom on Canada-US Differences in Wages, UBC Economics mimeo.

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.

Kahneman, D., J.L. Knetsch, and R. Thaler(1986). "Fairness as a Constraint on Profit Seeking: Entitlements in the Market," *The American Economic Review*, Vol. 76(4), pp. 728-41.

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 mimeo.

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 1a: Estimated minimum wage effects for women, LFS public use files 1997-2016

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Minimum wage effects:								
More than 50¢ below	-2.123 (0.413)	-1.702 (0.283)	-1.711 (0.317)	-1.590 (0.405)	-1.713 (0.373)	-1.717 (0.345)	-1.601 (0.278)	-1.717 (0.376)
30¢ to 50¢ below	-1.133 (0.341)	-0.824 (0.233)	-0.832 (0.256)	-0.590 (0.371)	-0.918 (0.465)	-0.971 (0.427)	-0.864 (0.372)	-0.922 (0.467)
10¢ to 30¢ below	-0.442 (0.235)	-0.164 (0.118)	-0.173 (0.137)	-0.192 (0.164)	-0.761 (0.171)	-0.823 (0.126)	-0.719 (0.069)	-0.765 (0.174)
At minimum wage	1.697 (0.238)	1.966 (0.129)	1.959 (0.144)	1.837 (0.304)	1.849 (0.350)	1.816 (0.287)	1.917 (0.215)	1.846 (0.352)
10¢ to 30¢ above	0.539 (0.260)	0.799 (0.147)	0.793 (0.167)	0.790 (0.200)	0.776 (0.221)	0.689 (0.181)	0.787 (0.124)	0.773 (0.225)
30¢ to 50¢ above	0.200 (0.165)	0.429 (0.106)	0.425 (0.103)	0.599 (0.126)	0.616 (0.133)	0.550 (0.149)	0.643 (0.144)	0.613 (0.136)
50¢ to \$1 above	0.125 (0.117)	0.316 (0.074)	0.312 (0.073)	0.393 (0.075)	0.364 (0.079)	0.294 (0.078)	0.379 (0.089)	0.361 (0.082)
\$1 to \$1.50 above	0.107 (0.078)	0.256 (0.025)	0.255 (0.026)	0.215 (0.026)	0.195 (0.031)	0.138 (0.073)	0.213 (0.114)	0.193 (0.033)
\$1.50 to \$2.00 above	-0.055 (0.113)	0.047 (0.130)	0.049 (0.126)	0.203 (0.057)	0.233 (0.076)	0.196 (0.096)	0.257 (0.119)	0.233 (0.076)
\$2.00 to \$2.50 above							0.064 (0.062)	
\$2.50 to \$3.00 above							-0.008 (0.035)	
Integer wage in bin				1.816 (0.029)	1.818 (0.029)	1.824 (0.029)	1.824 (0.029)	1.818 (0.029)
Log pseudolikelihood (/1000)	-370900	-368615	-368398	-249299	-247923	-246243	-246228	-247867
Province trends	no	yes	yes	yes	yes	yes	yes	yes
Segments w/ prov-year dummies			1	1	1	1	1	1&2
Interaction w/ first 6 months	no	no	no	no	yes	yes	yes	yes
Prov-wage & year-wage effects	no	no	no	no	no	yes	yes	no

Table 1b: Estimated minimum wage effects for men, LFS public use files 1997-2016

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Minimum wage effects:								
More than 50¢ below	-1.755 (0.440)	-1.426 (0.295)	-1.473 (0.300)	-1.323 (0.417)	-1.369 (0.389)	-1.295 (0.361)	-1.335 (0.284)	-1.378 (0.392)
30¢ to 50¢ below	-0.753 (0.380)	-0.523 (0.270)	-0.562 (0.272)	-0.265 (0.427)	-0.360 (0.519)	-0.317 (0.494)	-0.356 (0.415)	-0.367 (0.522)
10¢ to 30¢ below	-0.145 (0.245)	0.060 (0.121)	0.023 (0.130)	-0.021 (0.155)	-0.481 (0.159)	-0.452 (0.127)	-0.491 (0.078)	-0.488 (0.161)
At minimum wage	1.845 (0.232)	2.046 (0.120)	2.012 (0.128)	1.898 (0.308)	1.945 (0.368)	1.996 (0.322)	1.957 (0.217)	1.938 (0.370)
10¢ to 30¢ above	0.557 (0.212)	0.753 (0.096)	0.722 (0.102)	0.718 (0.145)	0.713 (0.154)	0.710 (0.138)	0.672 (0.133)	0.706 (0.157)
30¢ to 50¢ above	0.247 (0.157)	0.424 (0.103)	0.397 (0.110)	0.614 (0.137)	0.657 (0.142)	0.674 (0.147)	0.637 (0.153)	0.651 (0.144)
50¢ to \$1 above	0.155 (0.112)	0.305 (0.080)	0.283 (0.085)	0.384 (0.078)	0.389 (0.079)	0.398 (0.076)	0.363 (0.103)	0.384 (0.081)
\$1 to \$1.50 above	0.122 (0.086)	0.243 (0.035)	0.228 (0.039)	0.178 (0.040)	0.182 (0.045)	0.189 (0.075)	0.156 (0.133)	0.178 (0.046)
\$1.50 to \$2.00 above	-0.101 (0.127)	-0.014 (0.144)	-0.021 (0.145)	0.172 (0.067)	0.229 (0.084)	0.240 (0.101)	0.209 (0.138)	0.227 (0.084)
\$2.00 to \$2.50 above							-0.025 (0.080)	
\$2.50 to \$3.00 above							-0.066 (0.043)	
Integer wage in bin				2.081 (0.041)	2.084 (0.041)	2.089 (0.041)	2.088 (0.041)	2.084 (0.041)
Log pseudolikelihood (/1000)	-375599	-373603	-373360	-236036	-235391	-234173	-234161	-235317
Province trends	no	yes	yes	yes	yes	yes	yes	yes
Segments w/ prov-year dummies			1	1	1	1	1	1&2
Interaction w/ first 6 months	no	no	no	no	yes	yes	yes	yes
Prov-wage & year-wage effects	no	no	no	no	no	yes	yes	no

Table 2: Estimates effects of the minimum wage by job tenure, LFS public use files 1997-2016

	Women		Men	
	Tenure<=1 yr	Tenure>1 yr	Tenure<=1 yr	Tenure>1 yr
	(1)	(2)	(3)	(4)
Minimum wage effects:				
More than 50¢ below	-1.835 (0.416)	-1.565 (0.304)	-1.464 (0.442)	-1.101 (0.284)
30¢ to 50¢ below	-0.871 (0.533)	-1.012 (0.343)	-0.213 (0.601)	-0.404 (0.381)
10¢ to 30¢ below	-0.804 (0.174)	-0.846 (0.118)	-0.473 (0.133)	-0.465 (0.128)
At minimum wage	2.125 (0.326)	1.517 (0.291)	2.203 (0.360)	1.744 (0.315)
10¢ to 30¢ above	0.862 (0.226)	0.531 (0.179)	0.778 (0.159)	0.610 (0.127)
30¢ to 50¢ above	0.710 (0.215)	0.445 (0.116)	0.763 (0.169)	0.598 (0.137)
50¢ to \$1 above	0.404 (0.122)	0.224 (0.059)	0.453 (0.080)	0.353 (0.077)
\$1 to \$1.50 above	0.214 (0.113)	0.094 (0.049)	0.208 (0.083)	0.172 (0.066)
\$1.50 to \$2.00 above	0.264 (0.122)	0.164 (0.078)	0.284 (0.102)	0.216 (0.098)

Table 3a: Estimated minimum wage effects for women, LFS master files 1997-2016

	All (1)	Leavers (2)	Joiners (3)	Stayers (4)	Tenure<1 yr (5)	Tenure>1 yr (6)
Minimum wage effects:						
More than 50¢ below	-1.907 (0.187)	-1.612 (0.182)	-1.815 (0.183)	-1.892 (0.200)	-1.869 (0.204)	-1.798 (0.174)
30¢ to 50¢ below	-0.956 (0.296)	-0.644 (0.316)	-0.498 (0.238)	-0.995 (0.303)	-0.726 (0.333)	-1.037 (0.249)
10¢ to 30¢ below	-0.777 (0.065)	-0.425 (0.089)	-0.588 (0.159)	-0.790 (0.077)	-0.637 (0.050)	-0.836 (0.078)
At minimum wage	1.862 (0.122)	2.203 (0.125)	2.459 (0.189)	1.787 (0.136)	2.227 (0.125)	1.572 (0.148)
10¢ to 30¢ above	0.678 (0.071)	0.934 (0.122)	1.049 (0.201)	0.637 (0.072)	0.920 (0.121)	0.507 (0.048)
30¢ to 50¢ above	0.603 (0.153)	0.885 (0.203)	0.952 (0.252)	0.575 (0.148)	0.821 (0.189)	0.499 (0.113)
50¢ to \$1 above	0.313 (0.099)	0.468 (0.137)	0.529 (0.180)	0.299 (0.096)	0.436 (0.136)	0.264 (0.067)
\$1 to \$1.50 above	0.192 (0.131)	0.327 (0.162)	0.412 (0.170)	0.174 (0.132)	0.292 (0.151)	0.148 (0.119)
\$1.50 to \$2.00 above	0.223 (0.124)	0.371 (0.155)	0.484 (0.142)	0.201 (0.122)	0.302 (0.152)	0.194 (0.101)
\$2.00 to \$2.50 above	0.084 (0.085)	0.079 (0.121)	0.180 (0.136)	0.077 (0.082)	0.112 (0.105)	0.074 (0.079)
\$2.50 to \$3.00 above	-0.007 (0.048)	-0.032 (0.070)	-0.003 (0.114)	-0.007 (0.043)	-0.030 (0.087)	0.003 (0.034)

Note: All models include five covariate segments, province-specific trends, a full set of province-year controls in the first covariate segments, province-wage bin and year-wage bin controls, and interactions between the minimum wage function and whether the a minimum wage increase took place in the last 6 months.

Table 3b: Estimated minimum wage effects for men, LFS master files 1997-2016

	All (1)	Leavers (2)	Joiners (3)	Stayers (4)	Tenure<1 yr (5)	Tenure>1 yr (6)
Minimum wage effects:						
More than 50¢ below	-1.686 (0.184)	-1.740 (0.305)	-1.716 (0.297)	-1.662 (0.176)	-1.984 (0.228)	-1.313 (0.156)
30¢ to 50¢ below	-0.534 (0.318)	-0.610 (0.359)	-0.410 (0.373)	-0.536 (0.315)	-0.512 (0.392)	-0.480 (0.244)
10¢ to 30¢ below	-0.551 (0.057)	-0.532 (0.104)	-0.597 (0.106)	-0.537 (0.058)	-0.616 (0.090)	-0.466 (0.060)
At minimum wage	1.867 (0.114)	1.951 (0.149)	2.126 (0.191)	1.824 (0.116)	1.980 (0.110)	1.745 (0.163)
10¢ to 30¢ above	0.522 (0.165)	0.539 (0.206)	0.524 (0.158)	0.520 (0.168)	0.493 (0.211)	0.548 (0.118)
30¢ to 50¢ above	0.533 (0.161)	0.523 (0.256)	0.594 (0.182)	0.535 (0.151)	0.538 (0.201)	0.576 (0.102)
50¢ to \$1 above	0.235 (0.116)	0.225 (0.159)	0.267 (0.155)	0.237 (0.109)	0.193 (0.159)	0.304 (0.063)
\$1 to \$1.50 above	0.097 (0.150)	0.050 (0.178)	0.080 (0.122)	0.098 (0.155)	0.025 (0.190)	0.167 (0.113)
\$1.50 to \$2.00 above	0.102 (0.143)	0.128 (0.187)	0.174 (0.161)	0.098 (0.137)	0.076 (0.165)	0.148 (0.111)
\$2.00 to \$2.50 above	-0.065 (0.100)	-0.104 (0.104)	-0.164 (0.105)	-0.048 (0.098)	-0.105 (0.124)	-0.030 (0.082)
\$2.50 to \$3.00 above	-0.079 (0.064)	-0.083 (0.049)	-0.178 (0.107)	-0.062 (0.059)	-0.149 (0.083)	-0.024 (0.048)

Note: All models include five covariate segments, province-specific trends, a full set of province-year controls in the first covariate segments, province-wage bin and year-wage bin controls, and interactions between the minimum wage function and whether the a minimum wage increase took place in the last 6 months.

Figure 1a: Effect of the minimum wage on wage bin probabilities for women (logs)




Figure 1b: Effect of the minimum wage on wage bin probabilities for men (logs)

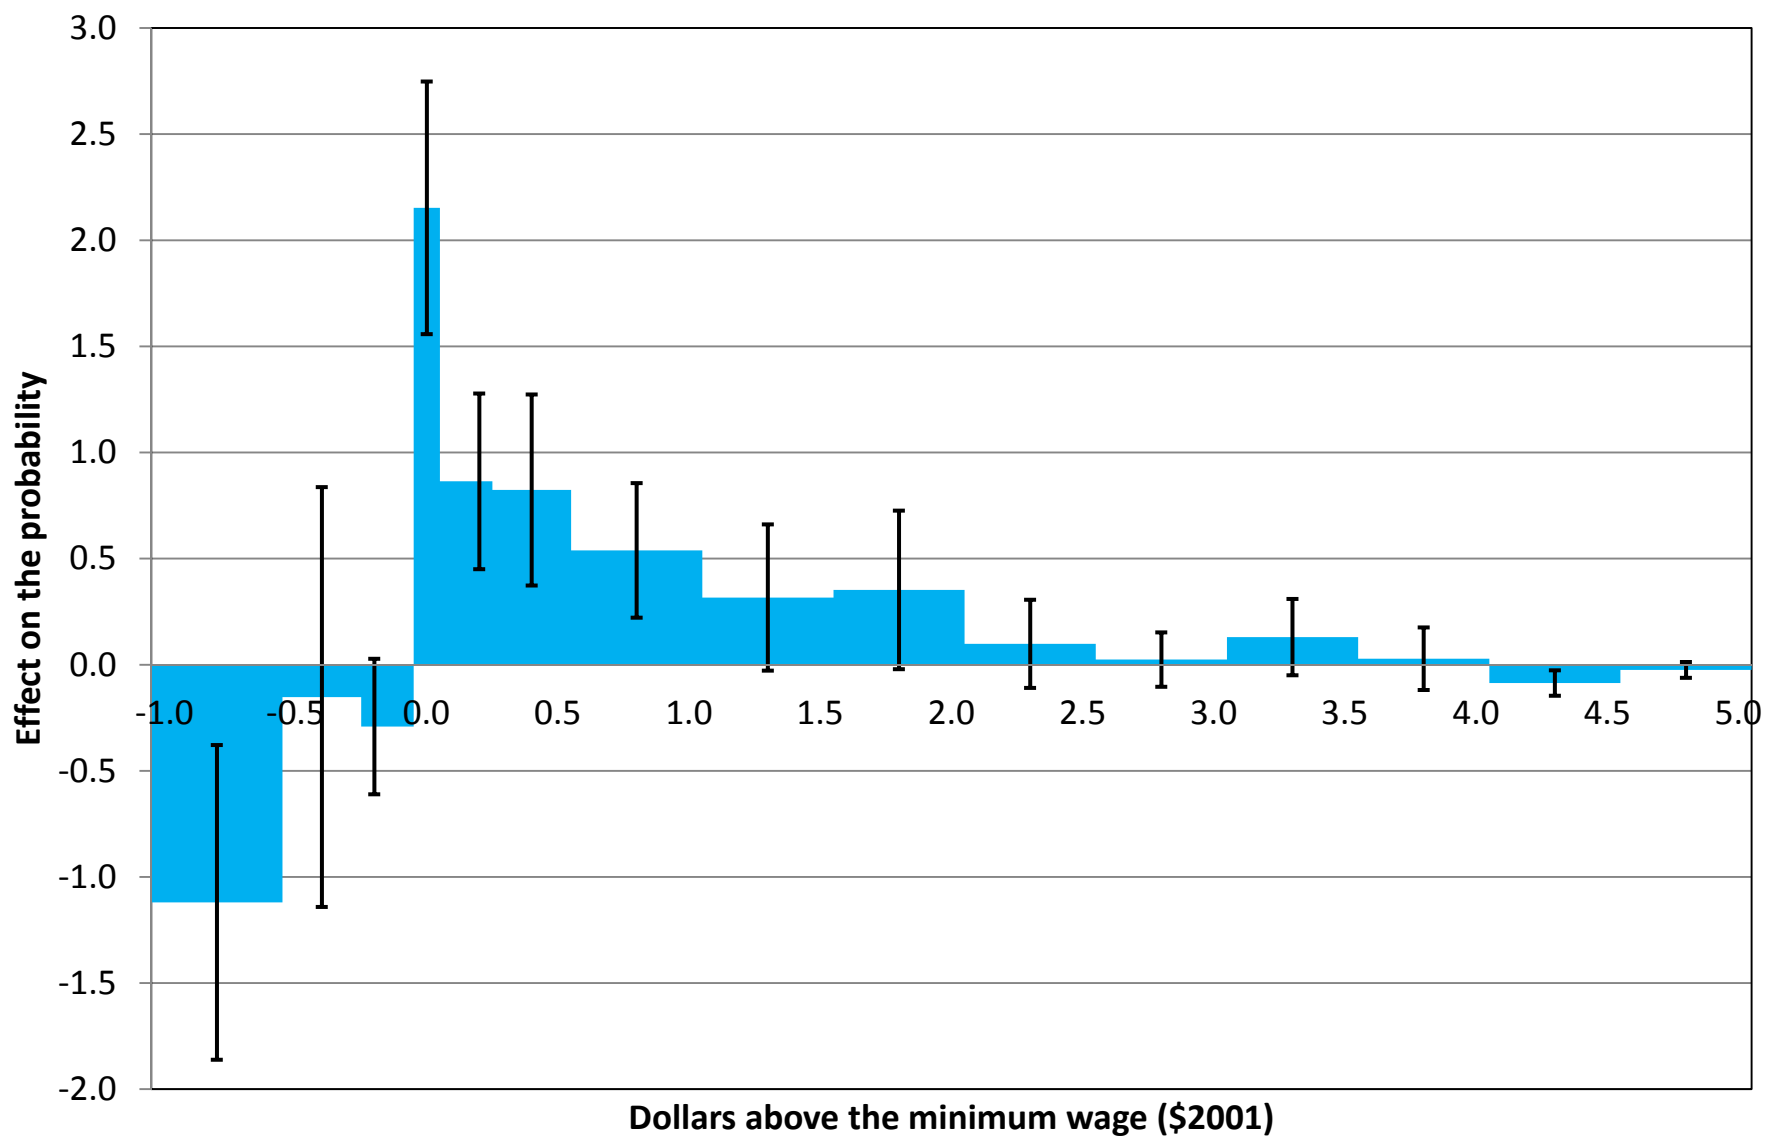


Figure 2a: Effect of the minimum wage on wage bin probabilities for women (% terms)

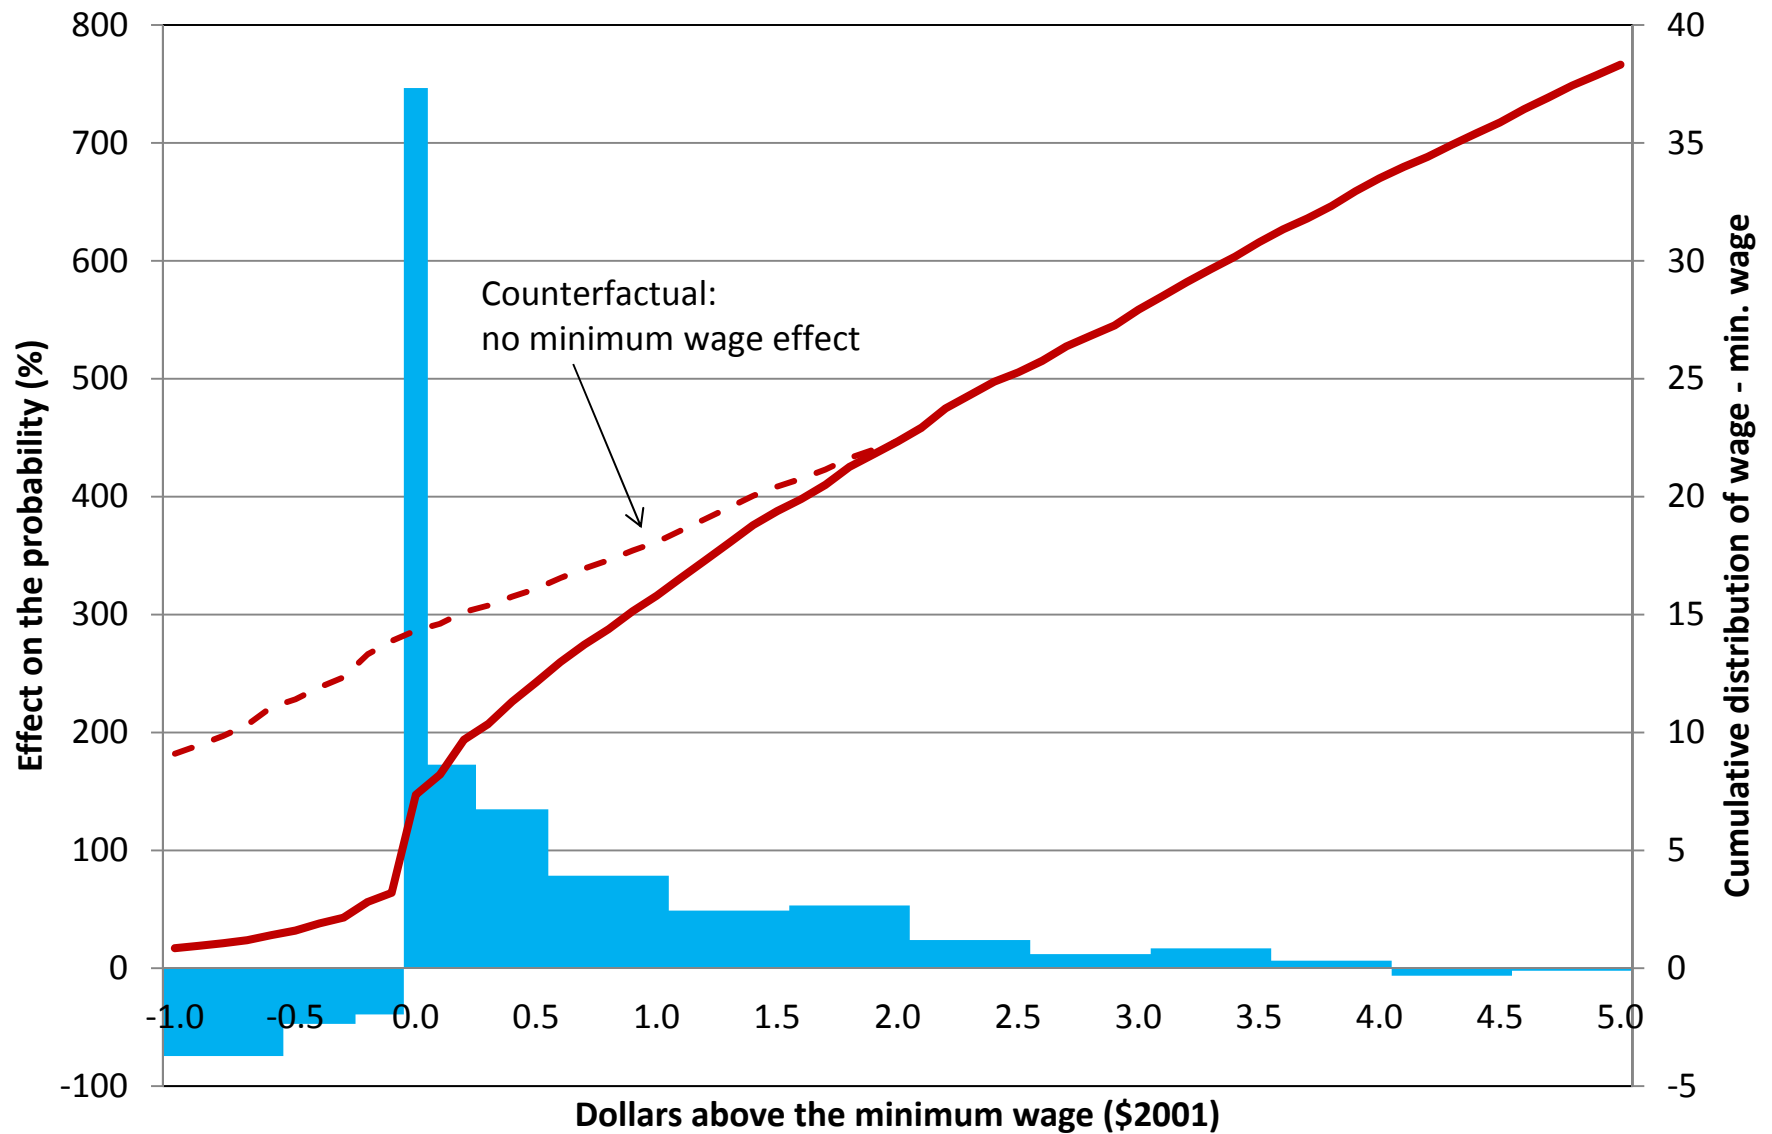
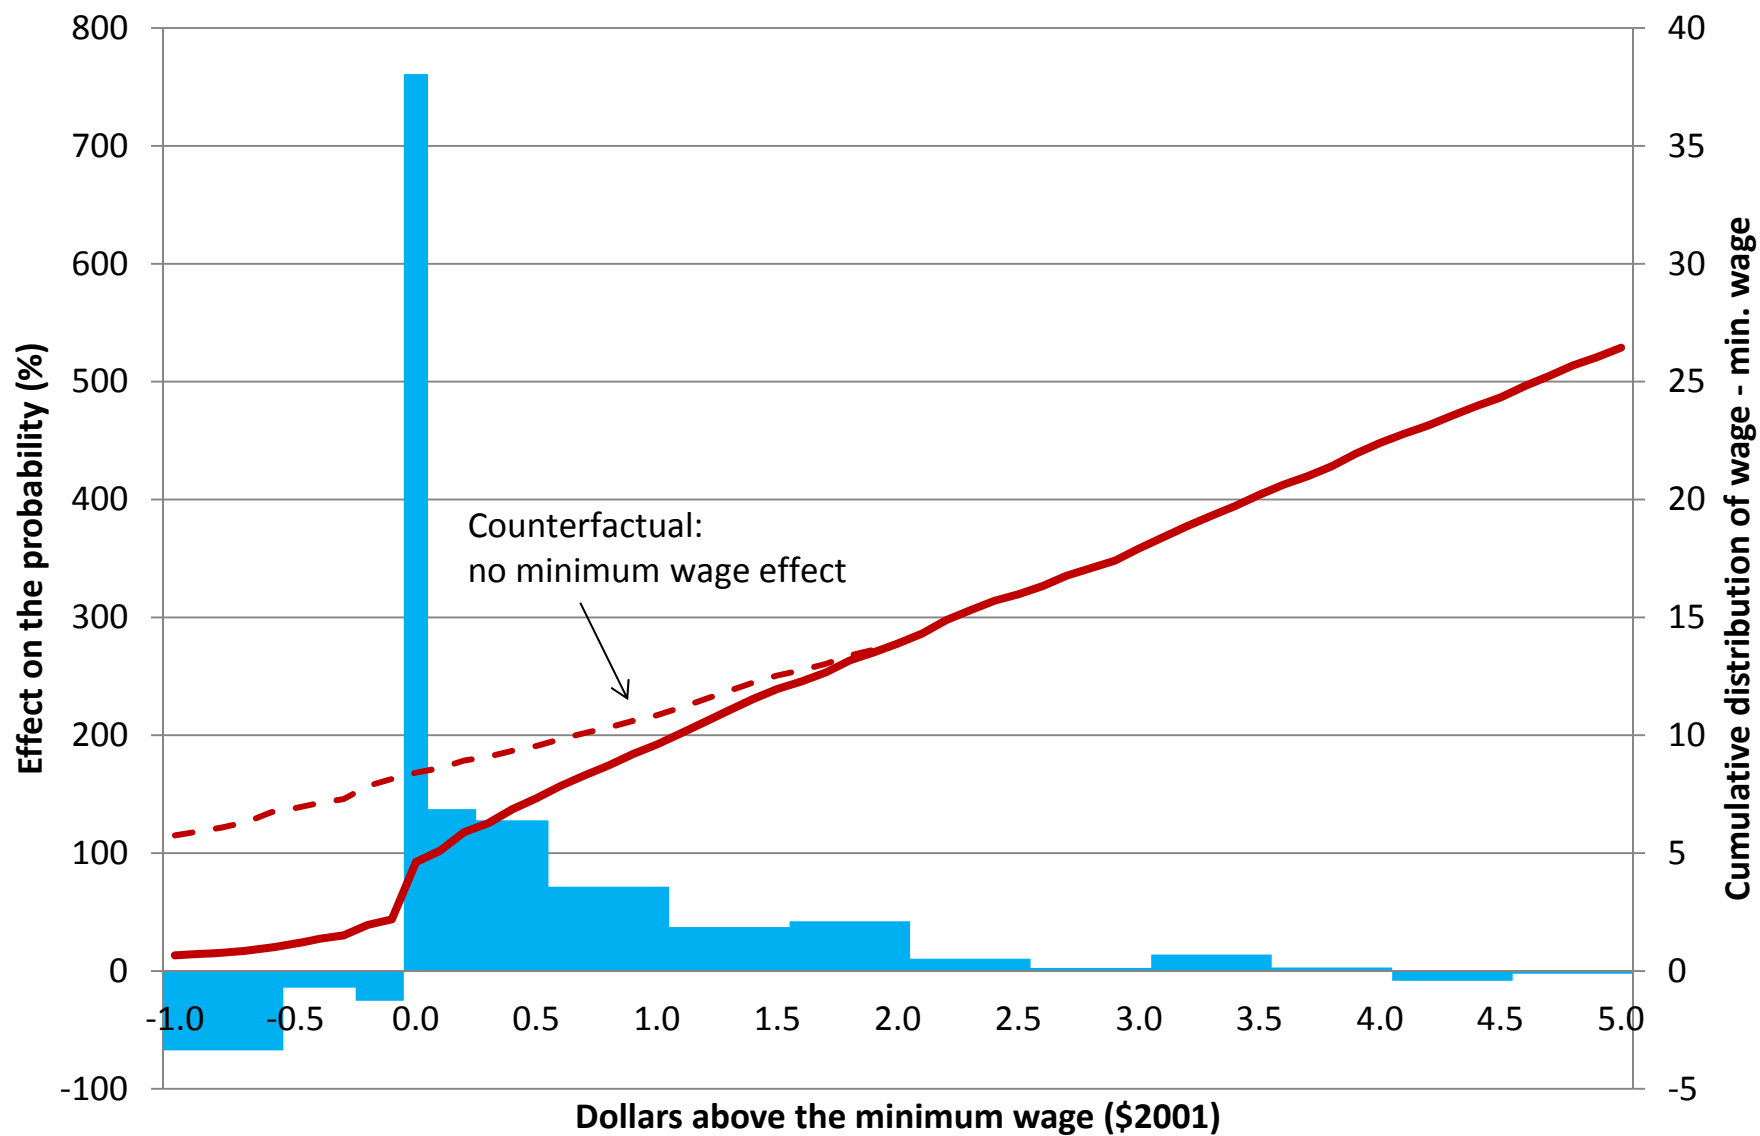
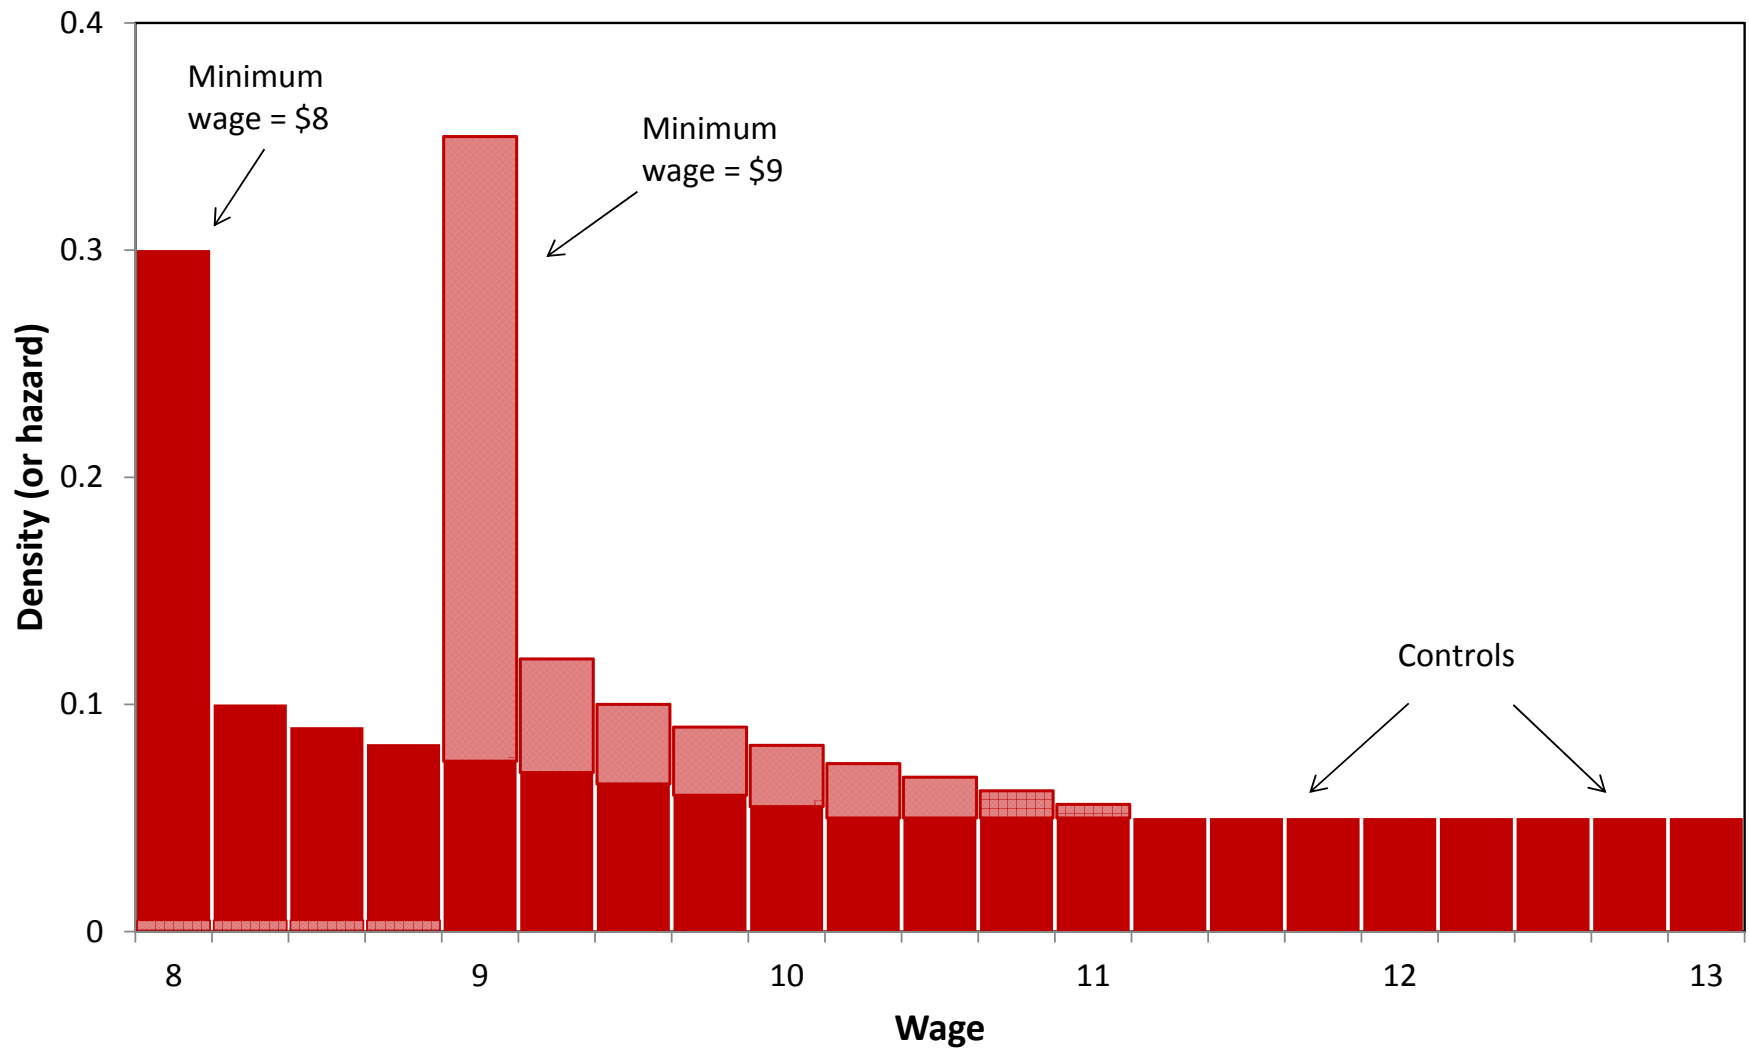


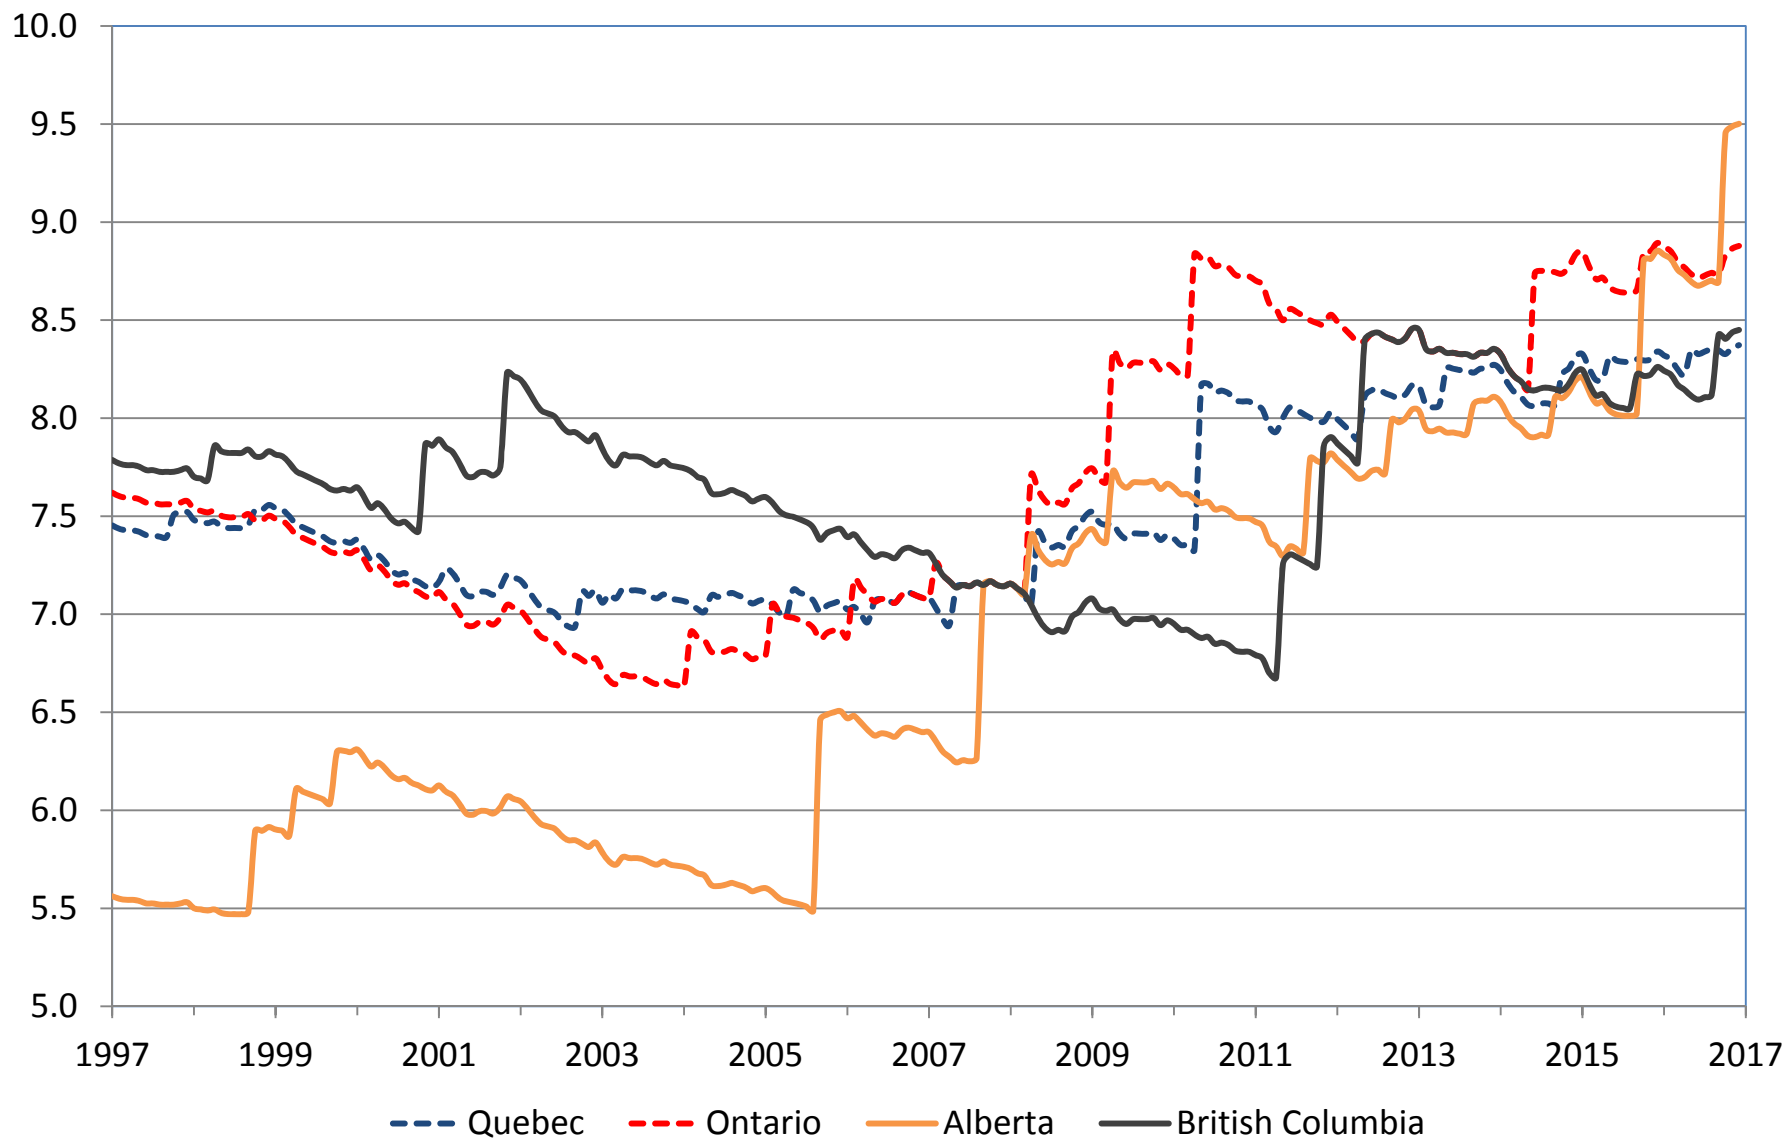
Figure 2b: Effect of the minimum wage on wage bin probabilities for men (% terms)



Appendix Figure 1: Triple-differences identification strategy



Appendix Figure 2a: Real Value of the Minimum Wage (\$2002), Larger Provinces



Appendix Figure 2b: Real Value of the Minimum Wage (\$2002), other Provinces

