

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

April 2023

Abstract

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

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

1 Introduction

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

Previous research on the minimum wage has provided fruitful ground for understanding which models best describe the functioning of the low-wage labour market. While a standard neoclassical model predicts that higher minimum wages should reduce employment, Card and Krueger (1995) conclude that a model in which employers have monopsony power and set their own wages best accounts for the lack of minimum wage employment effects they document in their research (see also Manning (2003)). Likewise, the presence of a spike at the minimum wage is incompatible with a standard neoclassical model where workers are paid their marginal product, which is smoothly distributed over the workforce. Again, monopsony power –or other sources of labour market imperfections– can readily account for the existence of the minimum wage spike.¹ However, an important limitation of standard monopsony models –or dynamic monopsony models with search frictions– is that they are not rich enough to account for the full set of wage adjustments that take place after a minimum wage increase. In particular, these models have little to say about wage dynamics among workers who remain on their existing job. Yet, we show that these wage adjustments largely shape the way in which the entire wage distribution evolves in response to a change in the minimum wage.

¹See Dickens et al. (2012) and Haanwinckel (2020) for examples of models with monopsony power that generate a spike at the minimum wage, and Engbom and Moser (2022) for an extension of the job ladder model that yields a spike among workers who are not actively searching for jobs.

Our main contribution in this paper consists in providing a new set of data patterns that are useful for examining the implications of key models of wage setting and the labour market. More specifically, using rich Canadian data consisting of a large set of short panels (and covering a period with over 150 changes in nominal minimum wages), we are able to decompose the effects of a change in the minimum wage on the wage distribution into elements related to job turnover and wage changes for continuing workers. An increase in the minimum wage can impact the wage distribution through three channels. The first is a selection channel linked to changes in the skill composition of workers (e.g. layoffs among low skilled workers and increased hiring of higher skilled workers) or of the firms employing these workers (e.g. shut-down of lower productivity firms, with demand shifting in favour of higher productivity firms). The second is through alterations in firm wage policies or in equilibrium wages induced by substitution effects. The third is a mechanical channel pointed out by Meyer and Wise (1983), that if there is a change in the proportion of people who work then the cumulative distribution function (CDF) of workers' wages will need to be rescaled in order to integrate to 1. This can complicate attempts to understand the impact of minimum wage changes on the shape of the CDF. We are able to examine the relative importance and the nature of these different channels.

At the heart of our investigation is a hazard based estimator presented in Donald et al. (2000) that allows us to examine the impact of the minimum wage on the shape of the wage distribution while controlling for potentially confounding factors such as education and age. Employing the hazard rates corresponding to a wage distribution (e.g., the probability that an observed wage is \$10 conditional on it being at least \$10) as the basis of our estimation has several advantages. First, if a minimum wage increase leads to reduced employment then the wage density will have to be inflated by 1 minus the reduction in the probability of employment in order to make it continue to integrate to 1. This is the third minimum wage impact channel mentioned earlier, and it is one that can make it difficult to discern shape changes in the wage distribution (a point raised by Flinn (2006) and Stewart (2012)). The hazard does not face this difficulty since the re-scaling factor affects both the numerator (the density) and the denominator (1 minus the CDF) and cancels out. Using the hazard function also allows us to tap into the rich specifications that have been developed in the hazard literature (e.g., Meyer (1990)). We employ a specification with a non-parametric estimator for the underlying shape of the hazard function and a very flexible mechanism for introducing covariate effects that allows them to have differing effects across the range of wages. Finally, the estimator is straightforward to implement since we do not need to impose integration constraints that

would be necessary for consistency in directly estimating densities, and we provide a new, Generalized Linear Estimator approach. As long as the hazard rates are non-negative, we can integrate them to form consistent, associated density and CDF estimates.

We implement our estimator, first, for the overall wage distribution. This allows us to replicate the results from papers such as Cengiz et al. (2019) and Gopalan et al. (2021). We identify the effects of minimum wage increases on the wage distribution by examining relative changes in different parts of the distribution before and after minimum wage changes across different provinces. Our estimator relies on a triple differences identification strategy as it allows us to control for idiosyncratic trends in different parts of the wage distribution separately by province. We then implement the estimator for Job Leavers, Job Joiners (i.e., new hires), and Job Stayers, allowing for minimum wages to have immediate and longer term effects. Mirroring the results in earlier papers, our estimates show a substantial spike at the minimum wage and spillover effects of minimum wage increases reaching up to approximately \$2 above the minimum wage.

The estimated effect of the minimum wage on the CDF for the three groups described earlier yield several key results. First, we find no effect of minimum wage changes on the wage distribution for Job Leavers, either in the short or long run.² This result clearly does not fit well with neoclassical models in which workers with abilities below the new minimum wage are simply laid off or more sophisticated ones with substitution from low skilled to high skilled workers. It also does not fit with versions of monopsony or bargaining models in which firms with productivity below the new minimum wage shut down. Thus, this result (plus the existence of a spike at the minimum wage) rejects standard neoclassical models and is potentially problematic for other models as well.

The second main result is that for Job Stayers, we find that the spike and spillover effects of a minimum wage are simply shifted right to the new minimum wage at the time of a minimum wage increase. That is, the effects of a minimum wage for continuing workers maintain a fixed pattern that is anchored on the minimum wage but is independent of its level. This finding is difficult to reconcile with most models typically invoked to understand minimum wage effects. Monopsony and search and bargaining models predict that workers whose wages were between the old and new minimum wage should all be moved to the new minimum wage. This generates a spike but, typically, no

²This potentially fits with results implying that minimum wages have small impacts on employment (e.g., Cengiz et al. (2019)), though null effects for employment rates do not necessarily imply no changes in layoff rates. Brochu and Green (2013) show that null effects for the employment rate actually reflect matching reductions in layoff and hiring rates. Alternatively, if higher skilled wage workers are substituted for lower skilled workers, we would see increased layoffs and hiring.

spillover effects among Job Stayers. Burdett and Mortensen (BM) type job ladder models provide a natural explanation for spillover effects as firm set wages to maintain their relative position in the wage distribution (Burdett and Mortensen (1998)). However, BM type models do not generate a spike at the minimum wage, which is a key feature of our data. As we discuss in more detail later in the paper, these limitations of monopsony or job ladder models are due to strong assumptions typical invoked to make these models tractable. More general versions of these models could yield both a spike and spillover effects. A more fundamental limitation of this class of models is that they are not well suited for understanding the source of spillover effects among Job Stayers. The findings for Job Stayers are instead consistent with a class of models based on internal incentives (promotions or tournaments), efficiency wages, or fairness issues where it is important for firms to preserve their internal wage structure.

On balance, we conclude that the evidence is consistent with a more general model where entry wages are set according to a job ladder, while wage growth for Job Stayers is driven by the kind of internal considerations discussed above. Such a model can be viewed as an extension of the job ladder model of Burdett and Coles (2003) where firms offer contracts with entry wages that later grow with job tenure. The difference here is that when firms are forced to raise wages at the bottom in response to a minimum wage, they also raise the whole wage profile to maintain the relative position of workers in the internal wage distribution.³

The paper also contributes to the growing literature on the impact of minimum wages on the wage distribution. The literature starts with Grossman (1983) and Meyer and Wise (1983), with important early contributions in DiNardo et al. (1996) and Card and Krueger (1995). These papers use US data and much of the existing literature is focused on either the United States (e.g., Lee (1999), Teulings (2000), Neumark et al. (2004), Flinn (2006), Autor et al. (2016), Cengiz et al. (2019)) or the United Kingdom (e.g., Manning (2003), Dickens and Manning (2004a), Dickens and Manning (2004b), Dickens et al. (2012), Stewart (2012)). But there is also a small but important literature examining Brazil, where the minimum wage appears to have played a role in substantially reducing inequality (Jales (2017), Engbom and Moser (2022), Alvarez et al. (2018), Haanwinckel (2020)) and new papers studying the introduction of the minimum wage in Germany

³Dube et al. (2019) study the case of a firm that did indeed adjust its whole wage schedule in response to the US federal minimum wage increase of 1996-97. Interestingly, the way the changes were implemented (using a step function) lead to uneven wage growth for otherwise similar workers. Consistent with the importance of fairness issues, workers who received smaller “spillover adjustments” were much more likely to quit their job.

(Dustmann et al. (2022)). For Canada, the only existing paper of which we are aware is Fortin and Lemieux (2015), who show important effects of minimum wage increases in reducing inequality. There is some disagreement in the literature on the extent of spillovers onto above-minimum-wage wages, ranging from large spillovers (Lee (1999)) to quite small ones (Autor et al. (2016)), though our reading of the most credible recent evidence for the United States is for spillovers up to about \$2 above the minimum wage. We find similar results for Canada.

Much of the existing literature estimates the impact of minimum wage changes on different quantiles of the wage distribution using either country level variation (especially in the United Kingdom and Brazilian studies) or sub-country federal units (for the later US studies and Canada). Estimates of effects on quantiles are useful for studying overall inequality effects but offer only a rough picture of the impact on the shape of the wage distribution. For example, a spike in the distribution will be located at different quantiles in different jurisdictions and so its existence may be hard to discern in quantile effects. Given our interest in decomposing changes in the shape of the wage distribution, it is important that we use an estimator that allows us to identify those shape effects directly.

Two papers share similarities with ours in broad methodology, though not in focus. Cengiz et al. (2019) employ an estimation approach that has similarities to ours, consisting of comparing employment changes in narrow wage bins just above and below the minimum wage. They show that minimum wage increases induce reductions in employment in wage bins below a new minimum wage that are exactly offset by wage increases at and just above the new minimum wage. In contrast to our investigation, their focus is mainly on employment effects. They do not show effects for the CDF for Leavers (one of our key results) and their estimates for what they call incumbents (our Stayers) are grouped in a way that makes it impossible to see our second main result. Gopalan et al. (2021) use rich administrative data and a matched counties along state boundaries estimation approach to identify wage and employment effects of minimum wage changes along the wage range. Their approach allows them to demonstrate the existence of spillover and spike effects but does not identify the shifts in the shape of the CDF that is our focus. They also show that layoff and quit rate effects are zero in all wage bins, which fits with our first result that the CDF of wages for Job Leavers does not change with a minimum wage increase. As a general statement, both papers show changes in the overall wage CDF that we replicate for Canada. Our contribution is to examine the component parts of the overall changes and to use those to differentiate between different models of the labour market.

The paper proceeds as follows. We start by explaining in Section 2 how the overall wage distribution is connected to those for Leavers, Joiners, and Stayers. In Section 3, we describe our estimator and its advantages. In Section 4, we describe our data and the minimum wage variation we use. We also present a preliminary event study in order to establish the core patterns. We want to make sure that readers see those patterns in case there are any concerns that our highly non-linear estimator is imposing any of our main results. Section 5 discusses implementation details for our estimator. The results for the all workers and the three separate subgroups (Leavers, Joiners, and Stayers) are presented in Section 6. We then discuss the implications of our key findings for models of the labour market in Section 7, focusing on the extent to which our empirical results do or do not match the predictions of these models. Section 8 contains conclusions.

2 Decomposing Changes in the Wage Distribution

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

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

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

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

quit, or have their firm shut down, and they could move immediately into a new job or be non-employed in $t+1$.

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

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

where $N_{J,t+1}$ is the number of Joiners in $t+1$ who were either in a different job or were not working in period t and $F_{J,t+1}(w|m_0)$ is their CDF of wages. As with the Leavers, there are no restrictions on where the Joiners came from. They could have been employed or nonemployed in period t and could have separated from a period t job for any reason.

Combining equations (1) and (2) highlights the forces that cause a change in the CDF relative to the previous period:

$$\begin{aligned} F_{t+1}(w|m_0) - F_t(w|m_0) &= \frac{N_{J,t+1}}{E_{t+1}} \cdot [F_{J,t+1}(w|m_0) - F_{L,t}(w|m_0)] \\ &\quad + \frac{N_S}{E_{t+1}} \cdot \Delta F_{S,t}(w|m_0) \\ &\quad + \left[\frac{N_S}{E_{t+1}} - \frac{N_S}{E_t} \right] \cdot F_{S,t}(w|m_0) + \left[\frac{N_{J,t+1}}{E_{t+1}} - \frac{N_{L,t}}{E_{t+1}} \right] \cdot F_{L,t}(w|m_0) \end{aligned} \quad (3)$$

where $\Delta F_{S,t}(w|m_0) = F_{S,t+1}(w|m_0) - F_{S,t}(w|m_0)$. The expression in equation (3) reveals three channels through which a change in the CDF of wages can occur. The first line shows that $F_{t+1}(w|m_0)$ will differ from $F_t(w|m_0)$ to the extent that the CDFs for Leavers in period t and Joiners in period $t+1$ are different. This is a standard selection effect, capturing forces such as a tendency for firms to lay off less skilled workers and replace them with more skilled workers. The second line points to changes in the shape of the wage distribution for Stayers as a second channel. This would arise either from policy changes other than the minimum wage, from changes in firm wage practices, or from macroeconomic changes. The third line captures compositional effects linked to changes in employment and worker flows in and out of employment. It is equal to zero in a steady-state environment where the flows in (Joiners) and out (Leavers) of employment are equal and constant over time.

Our main interest, of course, is in the impacts of changes in the minimum wage, and those changes can impact the wage CDF through any or all of the three components of equation (3), which correspond to the three channels discussed in the introduction. Previous papers mainly focus on the overall change in the CDF, showing that it involves

both shifts in the location of a spike at the minimum wage and spillover effects above the minimum wage. But we see equation (3) as revealing that relating those overall shifts to economic forces is complicated. Policy and economic changes can operate through any of the three channels in ways that could be either reinforcing or offsetting.

Consider an increase in the minimum wage from m_0 to m_1 that takes place between t and $t+1$. Using the notation $m_0 \rightarrow m_1$ to indicate the minimum wage increase, the decomposition for the change in the CDF in equation (3) remains the same except that m_0 has to be replaced with $m_0 \rightarrow m_1$. This denotes that a minimum wage increase will potentially change the components on the right hand side of the equation. Our estimation approach essentially consists of comparing changes in the components when there is a minimum wage increase with changes when there is no increase. For example, the density for Leavers, $F_{L,t}(w|m_0 \rightarrow m_1)$, would be quite different from the one when the minimum wage remains unchanged ($F_{L,t}(w|m_0)$) if the minimum wage increase leads to systematic layoffs of low-wage workers.

As we will discuss in Section 7, the implications of minimum wage changes in different models are most easily seen in their impacts in the selection (first line in equation (3)) and shape (second line in equation (3)) channels. For those reasons, we will focus on estimating the impact of minimum wage changes on: $[F_{J,t+1}(w|m_0 \rightarrow m_1) - F_{L,t}(w|m_0 \rightarrow m_1)]$ - the selection effect (as well as the individual components, $F_{J,t+1}(w|m_0 \rightarrow m_1)$ and $F_{L,t}(w|m_0 \rightarrow m_1)$); and the shape component ($\Delta F_{S,t}(w|m_0 \rightarrow m_1)$). In contrast, the composition effect term on the third line of equation (3) is unlikely to play much of a role since the existing evidence suggests that the minimum wage has, at best, a modest effect on employment. Small employment effects imply that the terms in square brackets are close to zero, and that the contribution of the third line of equation (3) is negligible.⁴ Furthermore, we explain in the next section that our hazard based estimator controls for mechanical effects linked to the systematic layoffs of workers earning below the new minimum wage. In the extreme case where all these workers are laid off (Meyer and Wise (1983)), $F_{L,t}(w|m_0 \rightarrow m_1)$ essentially represents the truncated segment at the bottom of the distribution, and the term in square brackets is negative due to a large layoff rate $\frac{N_{L,t}}{E_{t+1}}$. While the mechanical effect is inconsistent with small employment effects, it is re-assuring to know that our estimation procedure is not affected by this truncation problem.

⁴Brochu and Green (2013) show that in Canadian data, a lack of change in the employment rate due to a minimum wage increase conceals offsetting decreases in hiring and layoff rates. Under that scenario, the third line in equation (3) would remain small as the second term in the square bracket would still be close to zero.

Since the framework points to ongoing shifts in the CDF even in the absence of a minimum wage change, we will examine the minimum wage impacts using a difference-in-differences type approach. In the next section, we present an estimator for the CDFs that allows for considerable flexibility in the impact of other forces on them and is also helpful in mitigating the impact of the third, mechanical channel.

3 Hazard Based Estimator

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

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

$$f(y|x) = h(y|x)S(y|x), \quad (4)$$

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,

$$S(y|x) = \exp\left(-\int_{y_0}^y h(u|x)du\right) \quad (5)$$

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 (5), one can retrieve estimates of $F(y|x) = 1 - S(y|x)$. This approach allows us to make use of hazard function estimators developed in the duration model literature that: i) permit covariates to be introduced easily; and ii) are flexible, in the sense that they impose a minimum of restrictions on the shape of the hazard function for any value of the covariate vector. In contrast to standard quantile estimators, the resulting estimates of $F(y|x)$ are also guaranteed to be consistent in the sense that the distribution function always lies between 0 and 1.

One key advantage of focusing on the hazard rate is that it provides a convenient way of re-scaling the observed wage density to control for possible employment effects of

the minimum wage. That is, it eliminates the effects of the third (mechanical) channel, allowing us to concentrate on the effects of the selection and shape channels. To see this point, consider an example in which the minimum wage simply truncates the density without otherwise changing its shape. In this case, if the true wage density is given by $f(y)$ then the observed wage density is given by $f^*(y) = \frac{f(y)}{1-F(m)}$, where $F(\cdot)$ is the CDF for the wage and m is the minimum wage. The denominator in this expression is needed to ensure that the observed density integrates to 1, and it implies an inflating of the underlying density that will increase as m rises and the employment rate declines. Stewart (2012) calls this effect a truncation spillover, distinguishing it from a shape spillover in which the actual shape of the density changes above m . As we will see, the various models we are interested in have differing implications for shape spillovers that we would like to distinguish from truncation spillovers. Unlike the density, the observed hazard rate at y given by $h^*(y) = \frac{f^*(y)}{1-F^*(y)}$ is unaffected by the truncation of the wage distribution since the numerator, $\frac{f(y)}{1-F(m)}$, and the denominator, $1 - F^*(y) = \frac{1-F(y)}{1-F(m)}$ (the observed survivor function), are both divided by $1 - F(m)$. Thus, in our example in which minimum wages only serve to truncate the wage distribution, the hazard rate would remain unaffected. The same logic holds in our more general situation described in the previous section, as we show in Appendix A.⁵

A typical approach to including covariate effects 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 wage range into P segments defined by thresholds:

$$y_0 < y_1 < \dots < y_P. \tag{6}$$

The hazard function within a given segment p is given by

⁵One might wonder about the relevance of this result if the minimum wage has no impact on the employment rate since, in the simple example, the re-scaling occurs because of employment effects. Given that there is still some debate on the existence and size of employment effects, we believe it is better to use an estimator that is robust to their existence.

$$h(y|x) = \exp(x\alpha_p)h_0(y), \quad (7)$$

where α_p denotes the covariate effects on the hazard within the segment. The survivor function then becomes

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

where $p(y)$ denotes the first threshold above y , that is $p(y)$ is such that $y_{p(y)-1} < y < y_{p(y)}$. Since the effect of x is constant within each segment (y_{p-1}, y_p) , $p = 1, \dots, P$, it follows, as in Meyer (1990), that

$$\int_{y_{p-1}}^{y_p} \exp(x\alpha_p) h_0(u) du = \exp(x\alpha_p) \gamma_p, \quad (9)$$

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. 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.

When estimating minimum wage effects, we incorporate 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. Within the context of this estimator, minimum wages are analogous to time varying covariates in the standard duration literature. More specifically, 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, which is similar in nature to the approach taken in Cengiz et al. (2019) in their estimation of employment effects within wage bins. 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. Thus, our specification for the hazard in wage interval, p , is given by,

$$h(p) = \exp(D_{mp}\beta)\exp(x\alpha_{s(p)})\gamma_p \quad (10)$$

where, D_{mp} is the vector of minimum wage related dummy variables, taking values of 1 in the relevant wage interval. The estimated coefficient on the minimum wage variable itself corresponds to how much the baseline gets shifted with the presence of the minimum wage and provides an estimate of the “spike” in the hazard (and, in consequence, the density) at the minimum wage. The coefficients on the other variables in D_{mp} have a similar interpretation and, using those estimates, we can map out the effect of the minimum wage on the shape of the wage distribution.

We estimate the models 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 have investigated a variety of specifications for province and year effects. In our preferred specification, we treat the first covariate segment (where all of the minimum wages in our sample period lie) differently from the other four segments. In that first

segment, we include a complete set of province by year effects. In the other segments, we include separate sets of province effects, year effects, and a quadratic trend in time in order to control for the large differences in real wage growth across provinces over time (Fortin and Lemieux (2015), Green et al. (2019)).⁶ Recall that the effects of these variables are allowed to vary by covariate segment. To further our approximation, we also include interactions between the average wage in a bin and each of the full set of province and year effects. Because the bins are uniform in width (10 cents), this amounts to an interaction of the bin number with the province and year effects. We include these interactions in all the covariate segments, including the first, and allow their associated coefficients to differ by covariate segment. Thus, we employ a specification with considerable flexibility, with the shape of the distribution being allowed to vary by province and year but also, to a substantial extent, by province-year cells. This is important because Stewart (2012) showed that using high percentiles (e.g., the 50th percentile) as controls in an identification strategy for spill-over effects of this form could be problematic because the higher percentiles appear to be following different trends from the lower percentiles in his UK data. Cengiz et al. (2019), similarly, show that employment shocks can affect different parts of the wage distribution differently and in ways that confound estimation of minimum wage effects in their US data. Our specification addresses these issues by allowing for different trends in different parts of the wage range corresponding to the covariate segments. Moreover, as we will see, our identification is based just on movements in the wage range directly above the top of our minimum wage effect range (\$2 above the minimum wage in our main specification).

Given the complete set of province-by-year dummies in the first covariate segment, our most general specification is similar to a conventional “triple-differences” design. We demonstrate the identification strategy in Figure 1 where we plot imaginary estimates in a specification with baseline wage intervals that are 25 cents wide, a flat baseline hazard, the minimum wage allowed to have effects on the hazard rate up to 2 dollars above the minimum wage, and all of the wage intervals shown in the figure assumed to be within the first covariate segment. Thus, there is a common year by province effect helping to determine the height of hazard rates in all the wage intervals to a common degree. We plot the hazard rate in each wage interval for the cases of an \$8 minimum wage

⁶Our estimates change little when we include the complete set of province by year effects in the second covariate segment as well as the first. Indeed, allowing for province specific time trends is sufficient to capture the important variation within province over time, with results changing little when we replace those trends with the complete set of province by year interactions in the first or first and second covariate segments.

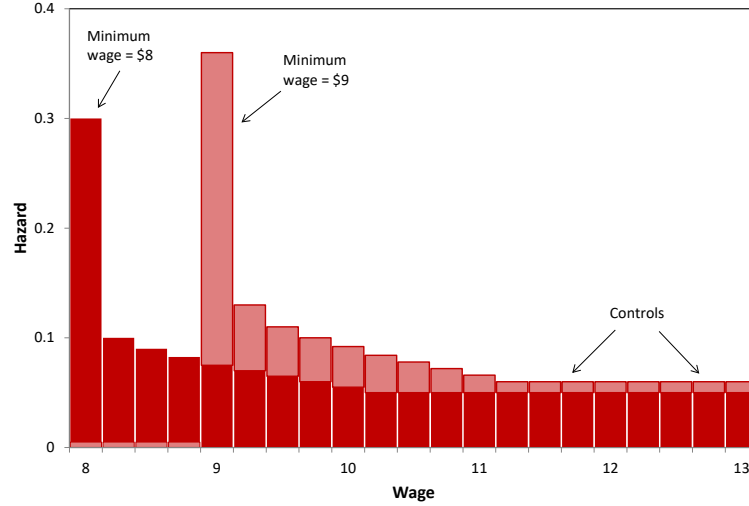
(dark bars) and a \$9 minimum wage (lighter bars) under an assumed pattern in which the minimum wage induces a spike (considerable extra mass) in the wage interval that contains the minimum wage and a spillover pattern with declining impacts as we move further above the minimum wage.

In this simple example, the province-by-year baseline hazards can be identified in the “control” zone above the upper limit of the spillover effects. The spillover effects can then be estimated by comparing the difference in hazards slightly above the minimum wage to those in the “control” region to the right. Consider, for example, the \$10 and \$11 wage range where there are some spillover effects in the province-year combination with the \$9 minimum wage (spillover between \$1 and \$2 above the minimum wage), but no spillover effects in the province-year combination with the \$8 minimum wage. The spillover effects can then be estimated by comparing the difference in hazards in the \$10 to \$11 to the difference in the control zone above \$11. Using a recursive approach, we can next estimate the spillover effects up to \$1 above the minimum by comparing the difference in hazards in the \$9 to \$10 wage range to those in the \$10 to \$11 wage range. As we keep moving to the left in 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. In essence, our approach enables us to fully control for province-year effects in the baseline level of the hazard while exploiting shape differences to estimate the spike and spillover effects.⁷

Several points follow from this design. First, a complete triple difference specification in our context would include a complete set of wage bin by province effects and wage bin by year effects. Because we use a large number of wage bins (164), including all of these effects is computationally infeasible. Our flexible province by year by covariate segment effects approach is an approximation to this. Second, because of the recursive nature of the identification, the parts of the minimum wage effects that are nearest to the upper cut-off (\$2 above the minimum wage in our example) tend to be more precisely estimated. Third, we can vary the range over which the spillovers are allowed to occur. In specifications not reported here, we allowed spillover effects as high as \$5 above the minimum wage. We have generally not found statistically significant or economically substantial effects above \$2 above the minimum wage and we use that as the limit of

⁷Note that our use of the hazard is important in this identification approach. If truncation spillover effects are included, as they would be if we worked with density functions, there would be no region of the wage range that is unaffected by the minimum wage since the truncation spillover effect inflates the entire density.

Figure 1: Triple-differences identification strategy



the spillover effects in our main estimates. We report on specifications that allow for a higher range for spillovers in Appendix x. **Do we have estimates with higher cut-offs?** Given the first point, this allows for more precise estimates of minimum wage effects in the region where they seem to actually arise. Fourth, unlike examinations of quantile movements, we obtain estimates of the impact of the minimum wage on specific parts of the wage distribution defined relative to where the minimum wage is located. In comparison, estimates of movements in, say, the 10th percentile of the wage distribution will reflect a combination of shifts in the spike at the minimum wage if the minimum wage is near the 10th percentile for some jurisdictions and of spillover effects if the 10th percentile falls above the minimum wage in other jurisdictions. This is related to the point made in Firpo et al. (2011) and Chernozhukov et al. (2013) that it is generally easier to econometrically model the probability distribution than quantiles.

4 Data and Minimum Wage Variation

4.1 Wage Data

Our wage data comes from the Canadian Labour Force Survey (LFS) master files.⁸ The LFS is a large Canadian household survey interviewing approximately 50,000 households per month that is similar in nature to the US CPS or UK LFS. We restrict our LFS sample to individuals aged 15 to 60. The self-employed are not included in the analysis as wage information is not available for these workers in the LFS. We further exclude full-time students because working is not their main activity. The LFS began to include a question on wages each month starting in November 1996. For our main estimates, we work with the data at the monthly level, with our data period running from January 1997 through December 2016. We construct education dummy variables for four groups: high school drop outs; high school graduates plus those with some but not completed post-secondary; people with a post-secondary diploma or certificate; and people with a university degree. We also use a set of age dummy variables corresponding to: 15 to 19; 20 to 24; 25 to 34; 35 to 54; and 55 to 60. Our wage variable is the hourly wage deflated to 2002 dollars using the consumer price index (CPI). Hourly wages are reported directly for workers paid by the hour but constructed using reported hours of work for those paid by the week, month or year. We have estimated our specifications restricting our sample to workers paid by the hour and did not find substantial differences in our estimated minimum wage effects.

As we describe in Section 2, minimum wage effects for groups of workers defined by movements into and out of jobs are central to our discussion. To construct those groups, we take advantage of the rotating panel design of the LFS. Individuals remain in the sample for six consecutive months, and every month one-sixth of the panel is replaced. As such, one can link consecutive months of the LFS thereby creating six-month mini panels.⁹ Combining this with the consistent LFS job tenure question which indicates for how long workers have been with their current employer, we construct three sub-groups:¹⁰

⁸The master files were accessed on site at the Manitoba and Carleton, Ottawa, Outaouais local Research Data Centre (COOL RDC). The research and analysis are based on data from Statistics Canada and the opinions expressed do not represent the views of Statistics Canada.

⁹A detailed description of how the data was linked can be found in Appendix B.

¹⁰The LFS asks when individuals started working for their current employer. Statistics Canada uses this information to construct a job tenure variable indicating the number of months workers have been with their employer. What distinguishes the LFS from other Canadian data sets, and US data sets, is that the job starting data question (with no change in wording) has been asked every month since 1976. See Brochu (2013) for a detailed discussion of the limitations of other North American data sets.

1) Job Leavers - workers who are observed on a job in the reference week, but who are not in that job two months later, whether because they have left employment or moved to a different job; 2) Job Joiners - workers who are on a job in the reference week, but have been on that job for two months or less; 3) Job Stayers - workers who have been on the same job for at least three months.

Note that we look two months instead of one month ahead to identify Job Leavers due to the timing of minimum wage changes that happen on the 15th of the month in a few cases (the minimum wage is raised on the first day of the month in most cases). Since the LFS reference week is typically defined to include the 15th of the month, it is unclear how a minimum wage increase on the 15th day of a month would affect reference week wages. To avoid this complication, we use a two-month window to insure a minimum wage change clearly happens between the two reference weeks. The same rationale is used to classify Job Joiners and Job Stayers using a wider time window than just month-to-month changes in job status.

Another challenge with the LFS is that respondents are only asked about their wage on a job at the first interview or if they take up a new job. Thus, if they do not change jobs, their recorded wage becomes “stale” in their subsequent months in the survey. For that reason, we restrict our attention to the incoming sample members each month. This means that we cannot construct panel estimates of, for example, the change in the wage for a Stayer at the time of the minimum wage change. Instead, we estimate the net effect of the minimum wage on the Stayer wage distribution, focusing on workers in the first month of their participation in the LFS in order to get “fresh” wage observations. To do this, we take advantage of the mini-panel aspect of the LFS and the job tenure question to compare a statistically similar group of Job Stayers before and after a change in the minimum wage. Looking forward, we identify Job Stayers who are still with their employer two months later, as indicated by a two-month increase in job tenure. Looking backward, focusing on workers with at least three months of job tenure insures these workers were on the same job two months ago. In Section 5.2, we discuss how this information can be combined to compute the impact of a minimum wage change on Job Stayers holding the composition of this group of workers constant before and after the minimum wage change.

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. We use the all earners sample for our main estimates but show in Appendix B that our key results remain the same with either of the other two samples. **do we have these results - for different samples - in a form we can put in an appendix?** 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. 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.¹¹ 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.

4.2 Minimum Wages

We use provincial minimum wage data for the period covered by our wage data: 1997-2016. The minimum wage is set at the provincial level in Canada, which introduces considerable variation in the minimum wage.¹² At times, some provinces have implemented lower rates for special classes of workers (e.g. students in Ontario). Most empirical studies indicate that firms do not take advantage of these special categories (e.g., Card and Krueger (1995)). Thus, we will focus on the general adult minimum wage for each province.¹³ As mentioned above, in a few cases provinces increase their minimum wage on the 15th day, as opposed to the 1st day, of the month. For the sake of consistency, we define the monthly minimum wage as the one prevailing on the first day of the month. We convert both minimum wages and our LFS wage measure to real terms using the monthly CPI (using 2002 as the base year).


Figures 2 and 3 show the evolution of real minimum wages by province for our sample period. Figure 2 shows the patterns for the four largest provinces. It is clear from this figure that there is considerable variation within provinces over time and since the

¹¹Dube et al. (2020) show that bunching at integer values is pervasive in wage data, though not all of the bunching is due to measurement error. Some bunching also occurs in high-quality administrative data.

¹²There 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.

¹³We 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.

provincial trends often move in different directions, this variation will be available even after controlling for general time effects. In the period between 2001 and 2009, for example, Ontario and Quebec showed flat trends in real terms while the BC real minimum wage was strongly declining and the Alberta minimum wage was strongly increasing. The patterns in Figure 3 for the six smaller provinces are much more similar across provinces, though even here there are some clear trend differences (e.g., Newfoundland between 2005 and 2011). Altogether, there are 157 nominal minimum wage changes in our time period. Many of these are small - the median nominal change is 3.8% - but 47 of the changes are 5% or larger and the largest is an 18% change in Alberta in 2006.




4.3 Event Study

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

$$y_{qpt} = \eta_p + \psi_t + \sum_{\tau=-8}^8 D_{\tau pt} \theta_{\tau} + \epsilon_{qpt} \quad (11)$$

Figure 3: Real Value of the Minimum Wage (\$2002), other Provinces




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

We define the relevant events as changes in the nominal minimum wage that are 5% or greater in size, treating periods with smaller or no nominal minimum wage changes as the control periods.¹⁴ We allow for the minimum wage changes to have effects starting as early as 8 quarters before the actual change to test for possible pre-trends linked to anticipatory changes or endogenous changes in the minimum wage that may be more likely to take place when the local economy is doing well and real wages are growing fast. We also look at wage effects as late as 8 quarters after that minimum wage change and leave out the dummy for period $\tau = -1$ so that the θ_τ 's correspond to effects relative to

¹⁴Adding a set of dummies to control for small changes in the minimum wage has little impact on the results and slightly reduces the precision of the main estimates.


this pre-event benchmark. We provide more details of our implementation of the event study in Appendix C.

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



For the 10th percentile (Figure 5), there is again little in the way of a pre-trend. With the exception of the quarter when the minimum wage takes place, the estimated effects are all positive and statistically significant after the increase in the minimum wage. Although the estimates are smaller than for the 5th percentile, there is again a clear trend in the size of the effects over time. By the 15th percentile (Figure 6),


Figure 5: Event study for women, log 10th percentile



the post-event effects are smaller still, and only become statistically significant in the fourth quarter after the minimum wage increase. Estimation for the 20th percentile (not shown here) reveals no significant post-event effects. For men (again, not shown here, for brevity) the estimated effects are smaller (often less than half those for women) and are not statistically significant past the 10th percentile.

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

¹⁵By direct effect we mean that when a wage percentile –say the 5th percentile– is equal to the minimum wage, there is a one-for-one relationship between the minimum wage and the wage percentile even in absence of spillover effects.



5 Implementation Details

5.1 Model Specification and Estimation

In implementing our estimator with the LFS data, we top-code wages at \$20 per hour since the computing effort required to estimate the hazard beyond that point is not useful when estimating minimum wage effects.¹⁶ We employ 164 baseline segments (wage bins) 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.

¹⁶Approximately 35 percent of workers earn below \$20 an hour in our estimation sample (recall that the wage is expressed in real dollars of 2002). Note that all observations, including those above \$20 an hour, are used to compute the hazard rates. As such, the estimation procedure accounts for the fact that a fraction of observations “survive” above the \$20 threshold.

As discussed earlier, we include two types of covariates. The first correspond to individual characteristics and time and province effects. In particular, we include dummy variables corresponding to 5 age groups and 4 education groups.¹⁷ We run the models separately for males and females and restrict our sample to be between 15 and 60 years of age. As described in Section 3, we include flexible sets of province and time effects that allow the wage density to change shape in complicated ways. The specification insures that the identification of minimum wage effects relies on the triple-differences design discussed in Section 3. As described earlier, we also include a dummy variable indicating whether a particular bin corresponds to an integer in the nominal wage data.¹⁸

The second set of covariates correspond to the minimum wage effects - at the minimum wage, below it, and for a range of wages above it. As described in Section 3, we include 9 dummy variables capturing the minimum wage effects, including one at the minimum wage itself. Also as described in that Section, we allow for effects up to \$2 above the minimum wage.

Combining the baseline hazard coefficients, the regular covariates, various types of time and province effects, and the minimum wage variables, the total number of estimated coefficients is as high as 745, depending on the specification. 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 month 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 reduce the computational burden and work with the GLM command in Stata.¹⁹

¹⁷The age groups are: 15-19, 20-24, 25-34, 35-54, and 55-60. The education groups are: high school drop-outs; high school grads plus people with some but not completed post-secondary education; a post-secondary degree or certificate; and a university degree (BA or more).

¹⁸We initially conducted extensive checks on the robustness of our results to different specifications using a public use version of the LFS. We found that the size of the spike and the extent and nature of the estimated spillovers are quite stable across specifications as long as we include two elements: the indicator variable for integer wage values, and province specific trends. Further complexity in the form of, for example, the complete interactions of province and year effects for the first covariate segment, did not materially alter results as long as those two elements were included.

¹⁹If there were wage observations in all of our cells (i.e., if the hazard rate were non-zero in each

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

As we discussed in Section 2, we are interested in both the impact of minimum wage changes on the overall wage distribution and on the wage distributions for job Leavers, Joiners, and Stayers. Moreover, we are interested in the dynamics of minimum wage impacts for these component groups, that is, how the immediate impacts differ from longer term impacts.

We estimate those dynamics by adding a set of dummy variables D_{mp}^c capturing changes in the minimum wages to the main hazard model. In particular, in the case of Leavers, we are interested in whether the wage CDF for Leavers who separate from their job at the time of a minimum wage increase is different from the CDF of Leavers at other times. Recall from Section 2 that these two CDFs correspond to $F_{L,t}(w|m_0 \rightarrow m_1)$ and $F_{L,t}(w|m_0)$, respectively. Given that Leavers are defined as workers in a month, t , who are not on the same job 2 months later, we focus on changes in the minimum wage happening over this two month window. Adding the new dummy variables D_{mp}^c to the base model in equation (10) yields:

$$h^L(p) = \exp(D_{mp}^b\beta + D_{mp}^c\delta)\exp(x\alpha_{s(p)})\gamma_p, \quad (12)$$

The set of dummy variables for the minimum wage in month t (before any increase), D_{mp}^b , is the set introduced in Section 3. We have added a b superscript to differentiate them from the minimum wage change variables. The set of dummy variables D_{mp}^c capturing upcoming changes in the minimum wage contains three groups of variables. The first group consists of dummy variables indicating if the wage interval p is below the old minimum wage - the one in place in month t - (50 cents or more below, 30 to 50 cents below, and 10 to 30 cents below). The second group of dummies allows for a flexible impact in the range going from the old to the new minimum wage. Three dummy variables are used for indicating if a given wage interval p is at the old minimum wage, at the new minimum wage, or in the range between the old and new minimum wages. The final group of dummies captures wage ranges that are above the new value of the minimum wage - the one that will be in place in month $t+2$. These are the same five

cell) then this estimator could be implemented with simple OLS. Because this is not the case here, 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. 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.

dummies going up to two dollars above the new minimum wage that are used to capture spillover effects.

Including the set of dummies D_{mp}^b and D_{mp}^c allows for a very flexible impact of an upcoming minimum wage change on the Leavers' wage distribution. The coefficients on the D_{mp}^c vector captures a change in the CDF of wages on the previous job for those observed to separate at the time of a minimum wage change relative to the pre-separation CDF for job Leavers in periods with no upcoming minimum wage change. If firms lay off some of the workers earning wages below the new minimum wage in response to a minimum wage increase, we should see a sharp increase in the hazard in this wage range. For example, the coefficient on the dummy indicating wages in between the old and new minimum wage should be positive and significant. In contrast, the estimated effects of D_{mp}^b capture the "steady-state" effects of the minimum wage in periods where the minimum wage is stable over time, and all the elements of D_{mp}^c are equal to zero. Recent vs. steady-state effects may differ if, for example, firms do not adjust the entire wage structure right away and wait a few months before giving wage increases to workers earning just above the new minimum wage.

We utilize a similar specification for Joiners - in this case, designed to capture short term impacts of minimum wage changes on who firms hire. For example, if a large group of workers get laid off by low-productivity firms in response to a minimum wage increase, many of these workers may get re-hired by higher productivity firms at or just above the new minimum. This would generate an abnormal concentration of newly hired workers around the minimum wage just after the increase that should vanish as the labour market returns to a steady-state equilibrium with a stable minimum wage. To estimate these effects, we compare Joiners in a period, t , when there has been a minimum wage increase in the previous 2 months (recalling that Joiners are defined as workers with job tenure under 2 months) with Joiners in provinces and months where there has been no recent minimum wage increase. These effects are captured using the following hazard specification:

$$h^J(p) = \exp(D_{mp}^b\beta + D_{mp}^c\delta)\exp(x\alpha_{s(p)})\gamma_p, \quad (13)$$

where D_{mp}^c includes the same variables as in the Leavers specification but with each dummy defined based on the differences between the old minimum wage (in place 2 months prior) and the new, current minimum wage. In this case, the set of dummies, D_{mp}^b is based on the value of the minimum prior to the change. The β coefficients again show the steady state effects of a minimum wage while δ reveals changes in those effects right after a minimum wage increase.

The specification for Stayers is complicated, to some extent, by the nature of our data. With true panel data recording wages in each month, we would follow Job Stayers and see whether their CDF changed in different ways at the time of a minimum wage change relative to other periods. As we described earlier, the LFS does not re-ask about the wage on a job for Job Stayers and, so, it cannot be used in this way. Instead, we estimate the impact of a minimum wage change on the CDF for Stayers using statistical arguments. If we think about a minimum wage increase happening between months t and $t+2$, we can define what we call “Forward” and “Backward” Stayers. Forward Stayers are workers who enter the LFS in month t and who have job tenure of at least 2 months when they are observed in month $t+2$. Backward Stayers are workers who enter the LFS in month $t+2$ and who have job tenure of at least 2 months at that time. Importantly, the Backward Stayers at time $t+2$ are statistically identical (drawn from the same population of Stayers) to the Forward Stayers at time t .

In a simple case with just two time periods before and after a minimum wage increase, we could estimate the impact of the minimum wage by comparing the CDF of Forward and Backward Stayers. Based on the statistical equivalence between these two groups, doing so is equivalent to following a given group of Job Stayers over time with true panel data. In a case like ours with many changes in the minimum wage over time and space, one could instead estimate equation (10) for Forward and Backward Stayers, and use the estimated β 's to compute the impact of minimum wage increases on the wage distribution of Job Stayers.

One shortcoming of this approach is that it is only valid when the wage distribution of job Leavers immediately switches to its new steady state right after the minimum wage increase. To allow for richer dynamic adjustments, we instead estimate the hazard specification introduced for Leavers (equation (12)) for Forward Stayers, and the hazard specification introduced for Joiners (equation (13)) for Backward Stayers. In both cases, the estimated β vectors should still show the impact of minimum wages in steady state. In the case of Forward Stayers, the δ coefficient vector from the equation (12) specification shows any anticipatory effects in terms of wage changes just before a minimum wage increase for workers who will ultimately stay with the firm. In contrast, the δ vector for Backward Stayers shows the impact of a recent minimum wage increase on the Stayers' CDF, measured relative to the steady state. If the steady state β vectors are equal and there are no anticipatory effects, the δ estimates for Backward Stayers captures the full short-run effect of the minimum wage change on the Stayers' CDF.

We present estimates for the Forward And Backward Stayers estimation in Appendix

D, showing both that the β estimates are very similar for both and that the anticipatory effects of δ are nearly all neither statistically nor economically substantially different from zero.²⁰ Given that, we focus on the Backward Stayers estimates in the main body of the paper, interpreting the β estimates as the steady state effects of minimum wages and the δ estimates as the immediate impacts of a minimum wage increase for a set of workers who do not change jobs as a result of the minimum wage change.

6 Estimation Results

6.1 Results for all Workers

Table 1 reports the estimated “steady-state” minimum wage coefficients (i.e., the coefficients corresponding to periods when there is no recent or upcoming minimum wage change) for all workers and for our three subgroups, broken down by gender. For all groups except Job Leavers (column 2), we report the β estimated using equation (13). In the case of Job Leavers we report the estimation coefficients β based on equation (12). Although the set of dummies, D_{mp}^c , are included in all the models, we report the coefficients on these variables in a subsequent set of tables.

We begin by considering the estimates for all workers, pooled together. Based on the discussion in Section 2, the “all workers” specification provides estimates of a reduced form impact of the minimum wage on the unconditional wage distribution. The estimated effects likely combine the direct impact of the minimum wage on the wage distribution, and the indirect impact linked to changes in the composition of employed workers. Looking at the results for all female workers pooled together in the first column of Table 1, we see strong reductions in the hazard rate below the minimum wage, with the effects being larger the farther below the minimum wage a wage is. There is a large, statistically significant spike in the distribution at the minimum wage and spillover effects up to \$2 above the minimum wage. The spillover effects decline in size as we move farther above the minimum wage but are all statistically significant at at least the 10% level up to \$2 above the minimum wage.²¹ These patterns are similar to those in Cengiz

²⁰The only exception to this statement is that the estimate of δ at the upcoming minimum wage is statistically significantly positive. In other words, it appears that firms start moving up some workers to this new minimum wage shortly before it comes into effect.

²¹Our robustness investigations show that spillover effects beyond \$2 above the minimum wage are small and not statistically significant for both males and females.

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

Wage Range	All Workers		Stayers		Leavers		Joiners	
	Females	Males	Females	Males	Females	Males	Females	Males
Over \$0.50	-2.07**	-1.5**	-2.04**	-1.52**	-1.61**	-1.41**	-2.22**	-1.26**
below	(0.29)	(0.35)	(0.3)	(0.32)	(0.32)	(0.28)	(0.32)	(0.54)
\$0.30 to 0.50	-1.03*	-0.43	-1.07**	-0.48	-0.67	-0.37	-0.71	-0.05
below	(0.41)	(0.47)	(0.41)	(0.45)	(0.37)	(0.39)	(0.41)	(0.58)
\$0.10 to 0.30	-0.83**	-0.41**	-0.85**	-0.43**	-0.31**	-0.22	-0.72**	-0.23
below	(0.11)	(0.11)	(0.11)	(0.1)	(0.13)	(0.2)	(0.25)	(0.17)
At min wage	1.67**	1.92**	1.60**	1.84**	2.06**	2.09**	2.12**	2.45**
	(0.22)	(0.27)	(0.23)	(0.26)	(0.23)	(0.25)	(0.27)	(0.33)
\$0.10 to 0.30	0.59**	0.74**	0.56**	0.69**	0.86**	0.79**	0.83**	1.03**
above	(0.14)	(0.11)	(0.14)	(0.11)	(0.17)	(0.1)	(0.23)	(0.13)
\$0.30 to 0.50	0.46**	0.63**	0.45**	0.6**	0.73**	0.67**	0.61**	0.87**
above	(0.11)	(0.11)	(0.11)	(0.11)	(0.15)	(0.11)	(0.17)	(0.16)
\$0.50 to 1	0.25**	0.35**	0.24**	0.33**	0.43**	0.39**	0.38**	0.53**
above	(0.06)	(0.06)	(0.06)	(0.05)	(0.07)	(0.06)	(0.14)	(0.08)
\$1 to 1.50	0.1	0.15*	0.09	0.13*	0.25**	0.15**	0.22*	0.3**
above	(0.06)	(0.06)	(0.06)	(0.06)	(0.08)	(0.04)	(0.1)	(0.08)
\$1.50 to 2.00	0.14	0.18*	0.13	0.15*	0.27*	0.19*	0.27*	0.36**
above	(0.07)	(0.07)	(0.07)	(0.07)	(0.1)	(0.07)	(0.11)	(0.08)
Num Obs	701137	685396	660231	640589	59894	63596	40906	44807


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

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

et al. (2019) and Gopalan et al. (2021) for the United States even though we focus on a different country and a different estimator. In Figure 7, we plot the exponentiated coefficients along with the estimated CDF for 20-24 year old high school graduates living in the province of Ontario with and without the minimum wage effects enacted. While we do not view the CDF without minimum wage effects as a true reflection of what would occur if there were truly no minimum wage in Canada, we present the non-minimum wage CDF as a benchmark against which to measure our estimated effects. Those effects are sizeable. The counterfactual CDF without the minimum wage is at 0.103 just below the minimum wage while the CDF including minimum wage effects is at 0.032 at that point. The CDF with the minimum wage more than doubles at the minimum wage point, rising to 0.074. Spillovers take the CDF higher and the two lines converge at \$2 above the minimum wage. The exponentiated minimum wage effect coefficients show that the minimum wage generates over a 400% increase in the wage hazard rate at the minimum wage and the spillover effects in the region above the minimum wage correspond to an 80% increase in the hazard rate just above the minimum wage that gradually declines as we move up in the wage distribution.

The results for men (column 2 of Table 1 and Figure 8) are remarkably similar to those for women. Recall that our estimated minimum wage effects represent proportional shifts in the underlying hazard function. Thus, the similarity of the estimated male and female minimum wage effects imply that estimates such as those in DiNardo et al. (1996) showing larger effects of the minimum wage on the female wage distribution arise because the female wage distribution has more mass in the region near the minimum wage rather than because minimum wages have a larger effect on the female labour market, per se. This is an intriguing result that suggests that employers do not assign female low wage workers to the minimum wage more frequently than they do low wage male workers, that is, that any discrimination does not take this specific form.

The remaining columns of Table 1 contain the estimated coefficients for the “steady state” effects of minimum wages on the CDFs for Stayers, Leavers and Joiners. Plots of the CDF without minimum wage effects (not shown here) reveal, not surprisingly, that both the Leavers and Joiners wage distributions have more weight at lower wages than the distribution for Stayers. The results in Table 1 show, in addition, that both the Leavers and Joiners distributions also have larger spikes at the minimum wage and more mass just above the minimum wage than for the Stayers distribution. This seems reasonable because it indicates that when the distribution is compressed with a minimum wage and workers whose wages would otherwise be below the minimum wage are now



at or just above the minimum wage, layoffs will be more likely in this compressed area. But it is equally the case that new hires are likely to happen in this bottom region of the distribution and, in fact, the comparable columns in Table 1 contain quite similar estimates for the effects at and above the minimum wage for Leavers and Joiners.

6.2 Dynamic Effects for Stayers, Joiners, and Leavers

As we discuss in Section 2, changes in the CDF for all workers in response to a minimum wage increase are a complicated function of the responses for Stayers, Leavers, and Joiners. Looking at these individual component responses helps distinguish the selection and shape change channels, which play different roles in different models. The responses to minimum wage changes, as summarized by the δ coefficients, are reported in Table 2.

6.2.1 Stayers: Evidence on Shape Changes


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

Table 2: Minimum Wage Effects on the Hazard
with Recent or Upcoming Minimum Wage Change

Wage Range	All Workers		Stayers		Leavers		Joiners	
	Females	Males	Females	Males	Females	Males	Females	Males
Over \$0.50 below old min wage	0.1 (0.07)	-0.19* (0.09)	0.09 (0.06)	-0.15 (0.08)	0.15 (0.11)	-0.25 (0.15)	0.11 (0.23)	-0.39* (0.16)
\$0.30 to 0.50 below old min wage	-0.7** (0.11)	-0.67** (0.06)	-0.67** (0.08)	-0.57** (0.07)	0.15 (0.13)	-0.1 (0.1)	-0.88** (0.27)	-1.23** (0.17)
\$0.10 to 0.30 below old min wage	-0.24 (0.29)	-0.14 (0.17)	-0.25 (0.3)	-0.18 (0.23)	0.01 (0.25)	0.29 (0.31)	-0.13 (0.32)	0.14 (0.18)
At old min wage	-1.41** (0.15)	-1.37** (0.09)	-1.33** (0.14)	-1.3** (0.08)	-0.1 (0.11)	-0.07 (0.11)	-1.94** (0.18)	-1.66** (0.19)
Between old and new min wage	-0.74** (0.1)	-0.66** (0.05)	-0.74** (0.1)	-0.62** (0.05)	-0.01 (0.11)	0.05 (0.19)	-0.74** (0.16)	-0.81** (0.11)
At new min wage	0.87** (0.07)	1.12** (0.12)	0.82** (0.08)	1.05** (0.12)	0.3* (0.15)	0.5** (0.17)	1.18 (0.11)	1.44 (0.14)
\$0.10 to 0.30 above new min wage	0.46** (0.07)	0.52** (0.06)	0.42** (0.07)	0.52** (0.06)	0.11 (0.1)	0.34* (0.14)	0.63** (0.12)	0.59** (0.08)
\$0.30 to 0.50 above new min wage	0.14* (0.05)	0.15* (0.07)	0.15** (0.04)	0.14 (0.08)	0.08 (0.09)	0.09 (0.05)	0.06 (0.09)	0.15 (0.09)
\$0.50 to 1 above new min wage	0.09* (0.03)	0.06 (0.04)	0.08* (0.04)	0.05 (0.04)	0.18** (0.05)	0.14* (0.05)	0.12 (0.07)	0.11 (0.06)
\$1 to 1.50 above new min wage	-0.04 (0.03)	-0.1 (0.04)	-0.02 (0.03)	-0.09 (0.05)	-0.05 (0.08)	-0.01 (0.09)	-0.27** (0.08)	-0.17 (0.11)
\$1.50 to 2.00 above new min wage	-0.07 (0.03)	0 (0.03)	-0.05 (0.03)	-0.03 (0.03)	-0.05 (0.03)	0.02 (0.05)	-0.3** (0.08)	0.12 (0.12)
Num Obs	701137	685396	660231	640589	59894	63596	40906	44807

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

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




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

In the case of Stayers, we graphically illustrate the impact of a 50 cent minimum wage increase in Figures 9a (for women) and 9b (for men). The effect of the minimum wage on the wage bin probabilities before a minimum wage increase is based on the steady-state coefficients β estimated for the “forward” Stayers. We then add the δ coefficients applied to a 50 cent change to obtain the second set of effects reported in the figure. For example, in the case of women we add the -1.33 change in the probability of being at the old minimum wage to the 1.60 steady-state effect of being at the minimum wage. We also plot the error bars to indicate whether the changes are statistically significant. For example, the standard error on the change at the old minimum wage is 0.139, yielding an error bar of ± 0.273 . The steady-state estimate of 1.60 falls well outside of the error bars, indicating that the 50 cents minimum wage increase significantly reduces the probability of being at the old minimum wage.


Interestingly, the results for both women and men indicate there is still a small abnormal concentration of observations at the old minimum wage after a minimum wage

increase, suggesting some inertia in how firms adjust their wages. This may be particularly possible for workers who are paid at the weekly or lower frequency rather than hourly since their compliance with minimum wages may be more difficult to assess. Combining the lingering spike at the old minimum wage with the evidence of more limited spillover effects (reaching only \$1 above the new minimum wage) suggests some delays in the way firms adjust their wage structure in response to an increase in the minimum wage.



Comparing the hazard effects before versus after a minimum wage increase suggests some similarity in shape. To evaluate this further, in Figure 10a, we replot the bars from Figure 9a but shift the “after” plot so that the new minimum wage is normalized to zero as in the “before” plot. This allows us to see whether the shapes of the hazards in the regions above the minimum wages are similar or different. The figure shows that the shape effects of a minimum wage are very similar at and above the minimum wage before and after a minimum wage change. The same similarity is evident in Figure 10b, for men. We don’t mean to imply from these figures that the shapes are identical. The exact size of the “after” effects depends on the size of the minimum wage change since that can change what elements of β each element of δ are added to. But there is a clear message

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




that directly after a minimum wage increase, the spike and spillover effects re-establish themselves in the same order of magnitude but anchored on the new, higher minimum wage. We do not, for example, observe a situation in which the spike becomes very large and there are limited spillovers because everyone who used to earn wages between the old and new minimum wage are moved right to the new minimum wage.

6.2.2 Leavers: Evidence on Selection

Assuming that firms don't alter wage policies in advance of minimum wage increases, the impact of the minimum wage on the job Leaver CDF identifies selection effects.²² In essence, our estimator compares the pre-separation wage distribution for Leavers in periods with a minimum wage increase to the same distribution for Leavers in other periods. We interpret the difference as showing us the impact of minimum wages on worker selection out of jobs. For example, if firms respond to a minimum wage increase by laying off their lowest ability/lowest wage workers then we should see extra mass at the low end

²²As noted earlier, when we estimate the "forward" stayer distribution, we find that $\delta=0$, i.e., that there is no evidence of anticipatory effects in wage setting before a minimum wage increase, at least for Job Stayers. We assume that conclusion carries over to the distribution for Job Leavers.

Figure 10a: Effect of the minimum wage on wage bin hazards before and after a 50 cents increase in the minimum wage (Stayers, Re-centred):
Women




of the wage distribution of Leavers in minimum wage increase periods relative to other periods. In terms of our estimator, this would show up as an increase in the spike at the old minimum wage and in the mass between the old and new minimum wages. The Leavers columns in Table 2, however, show small and statistically insignificant coefficient estimates for both of these effects. In fact, the estimated effects of the minimum wage increase on the Leavers' pre-separation distribution are nearly all statistically insignificant for both men and women. The only exceptions being a positive and statistically significant effect at the new minimum wage.

Many of the theoretical claims about the effect of a minimum wage increase in terms of worker separations focus on either weak firms shutting down or ongoing firms laying off their lowest wage workers. In Figures 11a and 11b, we present the effect of the minimum wage on the hazard rate before and after a \$0.50 increase in the minimum wage (similar to Figures 9a and 9b for Stayers) for workers who were laid off.²³ The similarity between

²³We define lay-offs as job separations for which the worker gave as a reason for losing their previous job as one of: end of seasonal job; end of temporary or casual job; company moved or out of business; business conditions; or dismissal or other reasons. Thus, layoffs include separations due to business closings, the end of temporary jobs, and being let go by an ongoing firm.

Figure 10b: Effect of the minimum wage on wage bin hazards before and after a 50 cents increase in the minimum wage (Stayers, Re-centred): Men




the No Minimum Wage change plot and the plot corresponding to an upcoming increase in the minimum wage in Figure 11a for women is remarkable, and only slightly less so in the plot for men in Figure 11b. We conclude that there is essentially no selection effect of a minimum wage increase in terms of the pre-separation wages for workers losing their jobs due to a layoff or firm closure.

In Figure 12, we present the same plot for workers who quit their previous job.²⁴ Because of small numbers of male job quitters, the plot for them has very large standard errors and is somewhat erratic. To provide more readable results (and because our results for men and women are very similar), we present a plot for men and women combined. Even then, the error bars in the figure are large, showing few statistically significant effects. The main significant effects are at the new minimum wage and just above it.

²⁴Quitters are job separators who give the reason for job loss as one of: Own illness or disability; caring for own children; pregnancy; other personal or family responsibilities; going to school; dissatisfied; retired.

Figure 11a: Effect of the minimum wage on Job Leavers (Layoffs) with and without an upcoming 50 cents increase in the minimum wage: Women




6.2.3 Joiners

For completeness, we also present the results for new job Joiners. Again, we present a plot of the impacts of a minimum wage on the hazard rate before and after a \$0.50 minimum wage increase (in Figure 13). To simplify the exposition, we present the results for men and women combined. The gender specific plots are very similar. The figure shows how a minimum wage increase affects wages for those who were hired in the previous 3 months and how that changes just after a minimum wage increase. As such, our estimated effects for Joiners reflect a combination of changes in who is hired (selection) and in how wage policies change (what we call wage policy shape changes). That implies that the Joiners plots have limited value in helping determine among models.

The Joiners plot shows a pattern that is quite similar to that for Stayers. In the wake of a minimum wage increase, there are substantial decreases in mass at the old minimum wage and between the old and new minimum wage. At the same time, a new spike grows in at the new minimum wage and the spillover effects just above it re-emerge. In fact, when we replicate Figures 10a and 10b for Joiners - re-centering the before and after plots on the same minimum wage - we again see that the minimum wage increase did

Figure 11b: Effect of the minimum wage on Job Leavers (Layoffs) with and without an upcoming 50 cents increase in the minimum wage: Men



not alter the shape of the hazard rate effects at and above the minimum wage.²⁵


In Figure 14, we plot the minimum wage effects on the hazard just after a \$0.50 increase in the minimum wage for Stayers and Joiners on the same figure. The effects are quite similar in the decrease in mass below the new minimum wage. But at and just above the new minimum wage, there is more mass for Joiners than Stayers. That is, new hires are particularly likely to be hired near and just above the minimum wage. Since both Joiners and Stayers show no real change in the shape of the minimum wage effect after versus before a minimum wage change, this is not a statement about the impact of a recent minimum wage increase but about the impact of minimum wages on Joiners even in steady state. Firms use the minimum wage more for new hires and then move workers to higher wages when they stay with the firm.

6.3 Results by Job Tenure

Differences between Joiners and Stayers raise the interesting possibility that the impact of a minimum wage could vary by job tenure. We have seen in the main results for men and women that the size of the jump in the wage CDF at the minimum wage can vary both

²⁵We do not show those plots in order to simplify the exposition.

Figure 12: Effect of the minimum wage on Job Leavers (Quits) with and without an upcoming 50 cents increase in the minimum wage:
Men and Women




because of the size of the baseline hazard (which reflects how far into the underlying wage distribution the minimum wage cuts) and the proportionality factor multiplying that hazard that is our estimated spike effect. We allow both factors to differ by job tenure by estimating our empirical specification separately for Stayers with under 1 year of job tenure, 1 to 5 years of job tenure, and over 5 years of job tenure. In order to highlight the tenure related differences, we plot the estimated effects for workers with under 1 year of job tenure and workers with over 5 years of tenure in Figures 15 (for women) and 16 (for men).

The results for females are striking. The plot for workers with under 1 year of tenure, naturally, looks like the one for Joiners, with a substantial increase in the hazard at the minimum wage and spillovers reaching at least a dollar above the minimum wage. But for workers with over 5 years of tenure, the effects below the minimum wage are similar (that is, there is a substantial reduction in mass in that region) but the effect at the minimum wage is half that for new workers and there are no spillover effects. For males, in Figure 16, long tenure workers also have a substantial reduction in the spike effect and reductions in spillovers, though the spillovers do not evaporate for them.

These estimated effects applied to the baseline hazard for each job tenure group yield very different minimum wage impacts. For our base group in the estimation (20-24year

Figure 13: Effect of the minimum wage on Job Joiners before and after a 50 centa increase in the minimum wage:
Men and Women




old high school graduates living in the province of Ontario), the CDF incorporating minimum wage impacts jumps from .027 to .049 when moving from \$9.90 to \$10.00 with a minimum wage set to \$10.00 for women with under 1 year of job tenure, and from 0.007 to 0.01 for women with over 5 years of job tenure. The difference reflects the greater bite of the minimum wage for new employees (the CDF value at \$10.00 with minimum wage effects set to zero is 0.16 for women with under 1 year of job tenure but only .034 for women with over 5 years of job tenure) as well as the larger estimated spike effects for them.

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

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

Figure 14: Effect of the minimum wage on Stayers and Joiners just after a 50 cent increase in the minimum wage:
Men and Women



of the labour market.


As a preview of our main findings, we first argue that, as in the minimum wage employment effect literature (e.g. Card and Krueger (1995)), our empirical results are incompatible with the predictions of standard neoclassical models. Introducing labour market imperfections (monopsony power or bargaining) helps account for some, but not all, of our findings on the importance of spillovers and a spike at the minimum wage. We conclude that firm pay policies linked to fairness or internal labour markets play a central role in how the minimum wage shapes the overall wage distribution.

7.1 Summary of Key Patterns

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

1. **Spike at the minimum wage.** We find a statistically significant and large increase in the hazard rate at the minimum wage, which fits with observations in many previous papers pointing to a spike in the density at the minimum wage.
2. **Spillovers up to \$2 above the minimum wage.** Like Cengiz et al. (2019) and Gopalan et al. (2021), we estimate statistically significant increases in mass for the

Figure 15: Effect of the minimum wage on Job Stayers by Job Tenure :
Women



hazard (and, so, for the density) in the range up to \$2 above the minimum wage.

3. No impact of an increase in the minimum wage on selection out of jobs.


Examining the impact of a minimum wage increase on the shape of the wage distribution for Job Leavers identifies who is selected out of jobs because of the increase. We cannot reject a null hypothesis that selection of Job Leavers does not change at the time of a minimum wage increase.

3a Quits and Layoffs The lack of impact of an upcoming change in the minimum wage is particularly evident for workers who are laid off (which, in our data, includes workers who are fired, let go, or affected by firm closures). For workers who quit their job at the time of a new minimum wage increase, in contrast, we estimate extra mass at and just above the new minimum wage.

4. The spike and spillover patterns for Stayers are re-established immediately.

We estimate that the spike at the minimum wage and the extra mass above it that exist before a minimum wage change are immediately re-established with the same orders of magnitude after a change but now anchored on the new minimum wage. There is no evidence of a large increase in spike such as would occur if all the workers who formerly earned wages between the old and new minimum wages

Figure 16: Effect of the minimum wage on Job Stayers by Job Tenure :
Men



are simply moved to the new minimum wage.

5. **Long tenure workers show smaller minimum wage effects.** Both the size of the spike at the minimum wage and spillover effects are smaller for longer tenure workers. In fact, for women, workers with over 5 years of job tenure have a spike effect that is half that for workers with under 1 year of job tenure and show no spillover effects. To the extent that estimated spillover effects reflect some kind of benchmarking of low end wages on the minimum wage, that benchmarking does not exist for longer job tenure women.

7.2 Economic Models

7.2.1 Neoclassical Models

We start by considering the simplest neoclassical model: one in which the wage equals the value of the marginal product of a worker. The observed distribution of wages in this model comes from a fixed distribution of productivities across workers. The introduction of a minimum wage should then price all workers with productivity below the minimum wage out of the market, resulting in a truncated distribution of wages. There is no reason

for a spike in the wage distribution at the minimum wage nor any reason for spillovers above the minimum wage.²⁶ In the Leavers distribution, we should see an extra mass of layoffs between the old and new minimum wages at the time of a minimum wage increase. Instead, our estimated patterns effectively reject this model. This is a well-known result that might almost be elevated to the status of a folk theorem: the spike in the wage distribution at the minimum wage implies a rejection of the simple neoclassical model (Card and Krueger (1995)).

Teulings (2000) develops a more realistic and complete neoclassical model in which workers are characterized by skills, with a large number of types of skills. Firms combine skills in production, taking into account the substitutabilities among skills and their relative wages. In this situation, a minimum wage will again truncate the wage distribution (dis-employing workers with skills whose marginal product falls below the minimum wage) but substitutabilities implies this is not the end of the story. As firms substitute away from skills with initial values of marginal product below the minimum wage, they increase their demand for more skilled workers.²⁷ The resulting increase in the wages for those above-minimum wage workers will cause a shift in relative demand for the workers who are most like them, including the workers whose former wages were just below the minimum wage. That means some of the workers who would be laid off in the simplest model will be employed at a wage just above the minimum. As a result, the wage distribution will display a spillover pattern with more mass just above the minimum wage and declining as we move up the wage range (toward workers who are increasingly less substitutable for the workers who were laid off). This is similar to what we estimate for the overall wage distribution. But as with the simplest neoclassical model, this model does not generate a spike at the minimum wage.²⁸ It also implies that there should be extra mass in the Leavers distribution at or near the old minimum wage, though the pattern will be less clear-cut than in the simple case. Again, the patterns that we estimate

²⁶We might create a model in which there just happens to be a clustering of productivity values at the minimum wage but we would then have to explain why that clustering moves when the minimum wage increases.

²⁷In Teulings' model, substitutability follows a simple pattern in which skills fall along a single dimension and skills that are closest to each other in value along that dimension are the most substitutable. It follows that demand increases the most for workers with previous wages just above the minimum wage since those workers are the most substitutable for the workers who were just laid off.

²⁸Brown (1988) makes the related point that if firms paid workers with varying abilities the minimum wage (in the period just after an increase), the spike would eventually unravel as firms compete for the more able workers by paying them just above the minimum wage in order to keep them from being poached. In contrast, we find that the minimum wage spike is permanent and is remarkably similar in the short and long runs for Stayers.

indicate a rejection of the model.

Haanwinckel (2020) generalizes Teulings’ model by introducing monopsony power and rich heterogeneity in firm production functions. As in other monopsony models we discuss next, this richer model yields a spike at the minimum wage. But as in Teulings’ model, the key force behind spillovers remains dis-employment effects among low productivity workers that increase the demand for slightly more productive workers. This prediction is at odds with our findings for Leavers.

7.2.2 Monopsony and Bargaining Models

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

$$w(p, w^*) = (1 - \alpha)w^* + \alpha p, \quad (14)$$

where p is the value of the worker marginal product in the firm; w^* is the worker’s next best alternative; and α is the share of surplus in the match going to the worker. Specifically, this is the form of the wage equation in search and bargaining models, where workers and vacancies have difficulty finding each other and the wage is bargained to divide a match specific surplus after they meet. (e.g., Flinn (2006)). It is also the form in what we might call classic monopsony models in which firms face their own upward sloping labour supply curve and post a wage before meeting workers (see Manning (2003)). This includes what Manning (2021) calls New Classical Monopsony models, which provide a behavioural foundation for the firm-specific labour supply curve (Card et al. (2018), Lamadon et al. (2022)). In the bargaining models, α is a parameter reflecting worker bargaining power while in classical monopsony models, it is the inverse elasticity of the supply of labour to the firm. Whether w^* is worker specific or varies at the level of the market or a type of worker varies depending on the model.²⁹

In contrast to neoclassical models, these models, when there is heterogeneity in p , can generate a spike at the minimum wage. To see this, consider the productivity threshold p^m , where $w(p^m, w^*) = m$, and m is the minimum wage. To make the exposition straightforward, assume that w^* is the same for all workers. p^m is then the productivity

²⁹Lamadon et al. (2022) derive a specification that also includes a compensating differential term and tax structure parameters, but the general form of the markdown on productivity related to worker preferences over heterogeneous firm attributes continues to hold.

at which the wage determination process would generate a wage of m . For all firms with productivity above m (so that it is still profitable for them to operate) and below p^m , the minimum wage forces them to increase the wage they pay but they are still getting some positive surplus. Thus, all these matches will continue and will pay exactly the minimum wage. Some of the models can also generate a spillover effect on above-minimum wages. In classical monopsony models, this occurs when the minimum wage affects the elasticity of labour supply facing firms (i.e., α) and does so more for firms paying just above the minimum wage than for firms paying higher wages (Manning (2003)).³⁰ In bargaining models this may happen if a minimum wage improves the outside option of workers (implying an increase in w^*)(Flinn (2006)).

Importantly, the spike and spillovers may arise in monopsony or bargaining models even when there are no minimum wage dis-employment effects. When the productivity of all firms is above m , we should see a spike at m and spillover effects for the reasons we just discussed. But employment remains the same as without a minimum wage since it is still profitable for all firms to operate. Thus, an important advantage of this class over neoclassical models “a la” Teulings is that dis-employment effects among low-skill workers –and the corresponding extra mass in the Leavers distribution at or near the old minimum wage– are not essential for generating wage spillovers.

However, as the minimum wage keeps increasing, firms with productivity p below a new minimum wage m' will shut down. Because those will be the lowest productivity jobs and the lowest productivity jobs will have been paying the old minimum wage, we should see an increase in size of the spike in our Leavers specification. This does not fit with what we observe. Furthermore, jobs that used to pay between the old and new minimum wage should, for the most part, now move to paying the new minimum wage, implying an increase in the size of the spike for Stayers. We also do not observe this. Instead, we see a re-establishment of the spike and spillover pattern anchored on the new minimum wage. The increased minimum wage will induce spillover rises in wages above the new minimum if the increase improved workers’ outside option values in a bargaining model, but those increases cannot be enough to re-establish the former pattern.³¹

³⁰In the Card et al. (2018) version of New Classical Monopsony models, there are no spillovers since neither w^* nor α vary across workers and firms. This is due to the convenient but simplifying assumption that workers choose among all possible jobs based on preferences that follow an extreme value distribution. In a more realistic version of the model where only employers in a certain segment of the labour market (e.g. around the minimum wage) compete for a given group of workers, employers that used to pay just above the minimum wage will need to compete more aggressively and raise wages when their close competitors increase their wages to comply with the minimum wage. This version of the model yields spillover effects, as in Manning (2003).

³¹For this to happen, the value of the outside option would have to increase enough to shift the wages

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

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

7.2.3 Models with Fairness Considerations and Internal Labour Markets

The idea that fairness considerations could explain minimum wage spillovers was first proposed by Grossman (1983). The core concept is that workers provide more effort if

of the workers formerly paid just above m to be just above m' , but since the outside option is a weighted average of wages paid in the economy, the increase, Δm , which applies to just part of the distribution, cannot increase the value of the whole weighted average by Δm .

they believe they are being treated fairly and their concept of fairness is based, in part, on how their wage relates to those of other workers in their firm. Akerlof and Yellen (1990) present a model based in the idea that the employment relationship can, at times, be viewed as a gift exchange in which workers reward firms who pay above the bare minimum with above minimum effort. The fairness standard against which workers compare the offered wage when determining their effort is written as a weighted average of the going wage in the market and the wages of others in their firm. A minimum wage then serves to move that standard for the lowest paid workers and then, as the wage of the lowest paid moves up, workers just above them require a wage increase to re-establish what they see as a fair wage premium. Since it is the comparison with the lowest paid workers that determines effort, firms will need to re-establish relative pay rankings relative to the minimum wage when it increases, which is what we observe. A spike would arise in this model if there is a set of workers whose ability is sufficiently low that, even including comparisons to others, their fair wage benchmark would be below the minimum wage. These workers would provide full effort when paid the minimum wage and, so, will all be paid at that level. This model could also explain the excess quits for workers who were paid at the level of the “new” minimum wage after a minimum wage increase. Those workers would have been paid more than the old minimum wage before the change and if their wage was not raised when the minimum wage increased then they could see that as unfair, inducing them to quit. This fits with Dube et al. (2019), who find that differential pay increases that were mechanically linked to the minimum wage at a large retail firm in the US caused workers with lower wage increases than identical colleagues to quit at a higher rate.³² The importance of fairness considerations in the context of the minimum wage is also supported by experimental studies such as Falk et al. (2006) which find that workers’ effort is related to the difference between their wage and the minimum wage. Because firms recognize this, in equilibrium, wage premia paid in order to induce effort will be paid relative to the minimum wage. This would generate the kind of pattern with a spillover pattern that moves up with the minimum wage that we estimate.

Internal labour markets provide another potential mechanism for explaining our results. In this broad class of models, the internal wage structure is used as a way of incentivizing workers to provide more effort and to stay with the firm. Consider, for example, the classic Lazear and Rosen (1981) paper on tournaments where firms set up a contest in order to extract optimal effort from workers. Worker output is determined

³²It also fits with evidence on effort and quit intentions in higher paid settings (e.g., Card et al. (2012) and Cullen et al. (2021)).

by the sum of their effort and a random shock, and output but not effort is observable by the firm and other workers. For a pair of workers, the firm offers a contract in which the person with the higher observed output gets a wage, W_1 , and the other person gets a wage, W_2 , with $W_1 > W_2$. The probability of workers winning the higher wage (which we can think of getting a promotion in the context of minimum wage jobs) is a function of their effort. The key result in these models is that the effort level depends only on the spread $W_1 - W_2$, not on the levels, W_1 and W_2 . But workers will only sign up with the firm if the expected value of the tournament minus the cost of the effort they anticipate expending exceeds the value of their outside option. Where the levels of W_1 and W_2 are set then depends on assumptions about the market structure. If the firm is a monopsonist, it will set $W_1 + W_2$ to take all of the surplus. If we assume that $W_1 > m$ while $W_2 < m$ then imposing the minimum wage, m , will transfer some of the surplus to the worker. As in the classic monopsony models discussed earlier, the jobs will continue because the firm is still earning positive profits (even if they are smaller than without a minimum wage). That is, we would, again, observe a spike at the minimum wage. As the minimum wage increases, the spread between the lower paid wage and the wage paid after promotion to the tournament winner will be re-established in order to continue to extract maximum effort. Thus, this simple form of the model can explain the result that the spike and spillovers are anchored on the minimum wage and are reproduced when the minimum wage increases. It also provides a potential explanation for why the spike and spillover patterns are lower for high tenure workers since the tournament will already have been run for them and they are either at their higher wage or will have left the firm if they lost. Similar conclusions hold in more realistic models with, for example, multiple workers competing for a limited set of promotions that pay a higher wage (Malcomson (1984)).

A drawback of these models is that they don't explain why we observe a spillover effect for Joiners. If there is variation in firm productivity and/or worker outside option values then tournaments at some firms will have a W_2 that is above the minimum wage, but there is then no reason for it to be affected by the minimum wage. In contrast, a Burdett-Mortensen type job ladder model can readily account for spillover effects in entry wages for Joiners as firms compete to maintain their position in the wage distribution.

In the end, none of the standard models of the labour market can fully account for the minimum wage effects we estimate. A "hybrid model" combining the key insights of models with internal considerations and imperfect labour markets is a more promising approach. Our results are consistent with monopsonistic firms competing with each other

when trying to hire workers, but internal considerations (fairness or internal labour markets) driving the evolution of wages within firms in response to minimum wage changes. Unlike a conventional job ladder model where workers are constantly searching for jobs, in a hybrid model where it is costly to change jobs, firms have incentives to set up an internal wage structure that incentivizes and helps retain workers. Such a hybrid model could also be viewed as an extension of Burdett and Coles (2003) where firms compete over contracts that offer both an entry wage and a tenure profile.

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

8 Conclusion

In recent years, our understanding of the effects of minimum wages has been substantially increased through the application of more sophisticated estimators to better data. In the US context, Cengiz et al. (2019) and Gopalan et al. (2021) have provided convincing estimates showing that minimum wage induce a sizeable spike in the wage density and spillover effects on wages up to about \$2 above the minimum wage. As they show, these results have important implications for both the employment effects of minimum wages and their impact on inequality. Reaching back at least to the 1980s, the variation induced by minimum wages has also been used to differentiate among different possible models of the labour market (Card and Krueger (1995), Manning (2003)). Structural modelling of the labour market in papers such as Engbom and Moser (2022) and Haanwinckel (2020) have taken the implications of minimum wages for the shape of the overall wage distribution seriously.

Our contribution in this paper stems from our use of panel data to decompose the effects of a minimum wage change on the wage distribution into a selection effect (in the form of changes in who separates from jobs as measured by where they were in the wage distribution before separation) and a wage distribution shape effect (measured by changes in the shape of the wage distribution for people who remain on their job at the time of a minimum wage increase). We argue that what happens with these different components

can help further differentiate among competing models of the labour market. Estimating these components amounts to estimating the effects of minimum wage changes on the whole shape of the wage distribution for different groups of workers. We accomplish this with a hazard function based estimator that allows for very flexible effects of minimum wages on the shape of the wage distribution estimated while controlling for the impacts of flexible functions of time, location and personal characteristics on the shape of the distribution. Using this estimator, we recreate the findings of a sizeable spike at the minimum wage and spillover effects up to about \$2 above the minimum wage in the Canadian context.

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

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

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

incorporate considerations explored in earlier generations of models on efficiency wages and job ladders.

References

- Akerlof, G. A. and J. Yellen (1990). The fair wage-effort hypothesis and unemployment. *Quarterly Journal of Economics* 105(2), 255–283.
- Alvarez, J., F. Benguria, N. Engbom, and C. Moser (2018). Firms and the decline of earnings inequality in Brazil. *American Economic Journal: Macroeconomics* 10(1), 149–189.
- Autor, D., A. Manning, and C. Smith (2016). The contribution of the minimum wage to US wage inequality over three decades: A reassessment. *American Economic Journal: Applied Economics* 8(1), 58–99.
- Brochu, P. (2013). The sources of the new Canadian job stability patterns. *Canadian Journal of Economics* 46(2), 412–440.
- Brochu, P. (2021). A researchers guide to the labour force survey: Its evolution and the choice of public versus master files. *Canadian Public Policy* 47(3), 335–357.
- Brochu, P. and D. A. Green (2013). The impact of minimum wages on labour market transitions. *The Economic Journal* 123(3), 1203–1235.
- Brown, C. (1988). Minimum wage laws: Are they overrated? *Journal of Economic Perspective* 2(3), 133–145.
- Burdett, K. and M. Coles (2003). Equilibrium wage-tenure packages. *Econometrica* 71(5), 1377–1404.
- Burdett, K. and D. T. Mortensen (1998). Wage differentials, employer size, and unemployment. *International Economic Review* 39(2), 257–273.
- Card, D., J. Heining, A. R. Cardoso, and P. Kline (2018). Firms and labor market inequality: Evidence and some theory. *Journal of Labor Economics* 36(S1), S13–S70.
- Card, D. and A. B. Krueger (1995). *Myth and Measurement: The New Economics of the Minimum Wage*. Princeton University Press.

- Card, D., A. Mas, E. Moretti, and E. Saez (2012). Inequality at work: The effect of peer salaries on job satisfaction. *American Economic Review* 102(6), 2981–3003.
- Cengiz, D., A. Dube, A. Lidner, and B. Zipperer (2019). The effect of minimum wages on low wage jobs. *Quarterly Journal of Economics* 134(3), 1405–1454.
- Chernozhukov, V., I. Fernandez-Val, and B. Melly (2013). Inference on counterfactual distributions. *Econometrica* 81(6), 2205–2268.
- Cullen, Z., S. Li, and R. Perez-Truglia (2021). What’s my employee worth? the effects of salary benchmarking. Technical Report 30570, National Bureau of Economic Research.
- Dickens, R. and A. Manning (2004a). Has the national minimum wage reduced UK wage inequality? *Journal of the Royal Statistical Society, Series A* 167(1), 613–626.
- Dickens, R. and A. Manning (2004b). Spikes and spill-overs: The impact of the national minimum wage on the wage distribution in a low wage sector. *The Economic Journal* 114(494), 95–101.
- Dickens, R., A. Manning, and T. Butcher (2012). Minimum wages and wage inequality: Some theory and an application to the UK. Technical Report 1177, Centre for Economic Performance.
- DiNardo, J., N. M. Fortin, and T. Lemieux (1996). Labor market institutions and the distribution of wages, 1976-1992: A semiparametric approach. *Econometrica* 64(5), 1001–1044.
- Donald, S. G., D. A. Green, and H. J. Paarsch (2000). Differences in wage distributions between Canada and the United States: An application of a flexible estimator of distribution functions in the presence of covariates. *Review of Economic Studies* 67(4), 609–633.
- Dube, A., L. Guiliano, and J. Leonard (2019). Fairness and frictions: The impact of unequal pay raises on quit behavior. *American Economic Review* 109(2), 609–633.
- Dube, A., A. Manning, and S. Naidu (2020). Monopsony and employer mis-optimization explain why wages bunch at round numbers. Technical Report 24991, National Bureau of Economic Research.

- Dustmann, C., A. Lindner, U. Schoenberg, M. Umkehrer, and P. von Berge (2022). Reallocation effects of the minimum wage. *Quarterly Journal Labor Economics* 137(1), 267–328.
- Engbom, N. and C. Moser (2022). Earnings inequality and the minimum wage: Evidence from Brazil. *American Economic Review* 112(12), 3803–3847.
- Falk, A., E. Fehr, and C. Zehnder (2006). Fairness perceptions and reservation wages - the behavioral effects of minimum wage laws. *Quarterly Journal of Economics* 121(4), 1347–1381.
- Firpo, S., N. M. Fortin, and T. Lemieux (2011). Unconditional quantile regressions. *Econometrica* 77, 953–973.
- Flinn, C. J. (2006). Minimum wage effects on labor market outcomes under search, matching, and endogenous contact rates. *Econometrica* 74(4), 1013–1062.
- Fortin, N. M. and T. Lemieux (2015). Changes in wage inequality in Canada: An interprovincial perspective. *Canadian Journal of Economics* 48(2), 682–713.
- Fortin, N. M., T. Lemieux, and N. Lloyd (2021). Labor market institutions and the distribution of wages: The role of spillover effects. *Journal of Labor Economics* 39(S2), S369–S412.
- Gopalan, R., B. Hamilton, A. Kalda, and D. Sovich (2021). State minimum wages, employment, and wage spillovers: Evidence from administrative payroll data. *Journal of Labor Economics* 39(3), 673–707.
- Green, D. A., R. Morissette, and B. M. Sand (2019). Economy wide spillovers from booms: Long distance commuting and the spread of wage effects. *Journal of Labor Economics* 37, S642–S687.
- Grossman, J. B. (1983). The impact of the minimum wage on other wages. *Journal of Human Resources* 18(3), 359–378.
- Haanwinckel, D. (2020). Supply, demand, institutions, and firms: A theory of labor market sorting and the wage distribution. Technical report, University of California Los Angeles.
- Jales, H. (2017). Estimating the effects of the minimum wage in a developing country: A density discontinuity approach. *Journal of Applied Econometrics* 33(1), 29–51.

- Lamadon, T., M. Mogstad, and B. Setzler (2022). Imperfect competition, compensating differentials, and rent sharing in the US labor market. *American Economic Review* 112(1), 169–212.
- Lazear, E. P. and S. Rosen (1981). Rank-order tournaments as optimum labor contracts. *Journal of Political Economy* 89(5), 841–864.
- Lee, D. S. (1999). Wage inequality in the United States during the 1980s: Rising dispersion or falling minimum wage. *Quarterly Journal of Economics* 114(3), 977–1023.
- Malcomson, J. (1984). Work incentives, hierarchy, and internal labor markets. *Journal of Political Economy* 92(3), 486–507.
- Manning, A. (2003). *Monopsony in Motion: Imperfect Competition in Labor Markets*. Princeton University Press.
- Manning, A. (2021). Monopsony in labor markets: A review. *ILR Review* 74(1), 3–26.
- Meyer, B. D. (1990). Unemployment insurance and unemployment spells. *Econometrica* 58(4), 757–782.
- Meyer, R. H. and D. A. Wise (1983). Discontinuous distributions and missing persons: The minimum wage and unemployed youth. *Econometrica* 51(6), 1677–1698.
- Neumark, D., M. Schweitzer, and W. Wascher (2004). Minimum wage effects throughout the wage distribution. *Journal of Human Resources* 39(2), 425–450.
- Ryu, K. (1992). Group duration analysis of the proportional hazard estimator: Minimum chi-squared estimators and specification tests. Technical report, University of California Los Angeles.
- Stewart, M. B. (2012). Quantile estimates of counterfactual distribution shifts and the effect of minimum wage increases on the wage distribution. *Journal of the Royal Statistical Society, A* 175(1), 263–287.
- Teulings, C. N. (2000). Aggregation bias in elasticities of substitution and the minimum wage paradox. *International Economic Review* 41(2), 359–398.

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

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

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

$$N_{S1}^M[F_{S1}^M(w_2) - F_{S1}^M(w_1)] = N_{S1}^{NM}[F_{S1}^{NM}(w_2) - F_{S1}^{NM}(w_1)] \quad (15)$$

$$N_{J1}^M[F_{J1}^M(w_2) - F_{J1}^M(w_1)] = N_{J1}^{NM}[F_{J1}^{NM}(w_2) - F_{J1}^{NM}(w_1)] \quad (16)$$

$$N_{L0}^M[F_{L0}^M(w_2) - F_{L0}^M(w_1)] = N_{L0}^{NM}[F_{L0}^{NM}(w_2) - F_{L0}^{NM}(w_1)] \quad (17)$$

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

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

$$\begin{aligned} F_1^M(w_2) - F_1^M(w_1) &= \frac{1}{E_1^M}[N_S^M(F_{S1}^M(w_2) - F_{S1}^M(w_1)) \\ &\quad + N_{J1}^M(F_{J1}^M(w_2) - F_{J1}^M(w_1))] \end{aligned} \quad (18)$$

and,

$$F_1^{NM}(w_2) - F_1^{NM}(w_1) = \frac{1}{E_1^{NM}} [N_S^{NM}(F_{S1}^{NM}(w_2) - F_{S1}^{NM}(w_1)) + N_{J1}^{NM}(F_{J1}^{NM}(w_2) - F_{J1}^{NM}(w_1))] \quad (19)$$

Under the no-minimum-wage effects on shape and selection assumption above, the right hand sides of equations 18 and 19 are equal except for the scaling factors, $\frac{1}{E_1^M}$ and $\frac{1}{E_1^{NM}}$. These scaling factors mean that the observed probability distribution function in this interval will change with the change in the minimum wage even though there was no direct economic effect of the minimum wage change in this part of the wage range. This is equivalent to what we described in the simple truncation case in the text.

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

$$1 - F_1^M(w_1) = \frac{1}{E_1^M} [N_S^M(1 - F_{S1}^M(w_1)) + N_{J1}^M(1 - F_{J1}^M(w_1))] \quad (20)$$

and,

$$1 - F_1^{NM}(w_1) = \frac{1}{E_1^{NM}} [N_S^{NM}(1 - F_{S1}^{NM}(w_1)) + N_{J1}^{NM}(1 - F_{J1}^{NM}(w_1))] \quad (21)$$

The hazard rate for M is given by the ratio of (18) to (20) and the hazard rate for NM is given by the ratio of (19) to (21). Note that the scaling factors $\frac{1}{E_1^M}$ and $\frac{1}{E_1^{NM}}$ cancel out, leaving terms that are equivalent according to our assumption that the minimum wage has no scale or selection effects in this range. Thus, while the pdf would show scaling effects that might be mistaken for spillovers from the minimum wage into the $[w_1, w_2]$, the hazard rate would show that there are no selection or shape channel effects.

B On-line Appendix: Data Linkage Details

In this appendix, we describe the process by which we link individuals across Labour Force Survey (LFS) months. The LFS has a rotating panel design in which households remain in the sample for six consecutive months, with 1/6 of the sample replaced each month. Attempts to create a panel from waves of the LFS are complicated by the fact that it follows dwellings rather than individuals. Individuals who change dwellings exit

the survey. Fortunately, we can uniquely identify individuals across monthly files with a combination of LFS variables. Changes over time in geographical identifiers (e.g., EI regions) mean that different identifying variables must be used for different periods but since our sample period begins in 1996, we only need to use one set of consistent variables (Brochu (2021)). In our description of the procedure for identifying individuals and linking them across months that follows we name variables (in capital letters) using the same names provided in the LFS codebook.

We linked observations across months using the month, the one-digit province code (PROV1), the pseudo UIC regions (PSEUDOUIC), the regional strata (FRAME), the superstratum (STRAFRAM), the sample design type (TYPE), the first-stage sampling unit (CLUST), the rotation number (ROTATION), the number assigned to dwellings within a cluster (LISTLINE), the multiple dwelling code for structures that have more than one dwelling (MULT) and the LINE variables. After we created our data, the LFS has added a new variable, HHLDID, which makes the linkage much more straightforward.³³

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

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

In this appendix, we present extra implementation details for the events study exercise.

One complication faced in running the event study is that there is overlap in the event-study windows when two large changes in the minimum wage take place less than 4 years apart. One possible approach for dealing with this issue is to restrict the analysis to minimum wage changes that take place more than four years away from each other.

³³See Brochu (2021) for a complete discussion of linkage issues in different periods and other issues with LFS data.

This would, unfortunately, result in a sharp drop in the number of large minimum wage changes to be used in the event-study analysis. A more efficient approach (and the one we use here) is to add a set of “treatment window dummies” to control for the possible overlap in the event-study windows.

In the case where the minimum wage only changes once, say in January 2007, we include dummies for three province-specific treatment windows: a “pre-treatment” window for the 1997-2004 period; a “treatment” window going from January 2005 to December 2008; a “post-treatment” window going from 2009 to 2016. In this simple example, the “trigger” for creating a new window is hitting 8 quarters before or after the minimum wage change.

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

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

D On-line Appendix: Minimum Wage Effect Estimates for Forward and Backward Stayers

Table 3: Minimum Wage Effects on the Hazard
with No Current Change in the Minimum Wage
Stayers

	Male Stayers		Female Stayers	
	Forward	Backward	Forward	Backward
Min wage effects:				
More than \$0.5 below	-1.550 (0.331)	-1.524 (0.320)	-2.026 (0.321)	-2.043 (0.299)
\$0.30 to 0.50 below	-0.424 (0.444)	-0.476 (0.453)	-0.991 (0.396)	-1.072 (0.410)
\$0.10 to 0.30 below	-0.347 (0.115)	-0.433 (0.098)	-0.763 (0.109)	-0.849 (0.108)
At minimum wage	1.885 (0.246)	1.836 (0.259)	1.633 (0.220)	1.602 (0.227)
\$0.10 to 0.30 above	0.658 (0.104)	0.690 (0.111)	0.559 (0.144)	0.558 (0.138)
\$0.30 to 0.50 above	0.596 (0.109)	0.601 (0.106)	0.437 (0.114)	0.451 (0.109)
\$0.50 to 1 above	0.329 (0.037)	0.329 (0.054)	0.222 (0.064)	0.238 (0.055)
\$1 to 1.50 above	0.141 (0.052)	0.127 (0.057)	0.084 (0.059)	0.089 (0.062)
\$1.50 to 2.00 above	0.168 (0.064)	0.155 (0.072)	0.121 (0.072)	0.132 (0.074)

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

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

Table 4: Minimum Wage Effects on the Hazard
with Recent or Upcoming Minimum Wage Change
Stayers

	Male Stayers		Female Stayers	
	Forward	Backward	Forward	Backward
Minimum wage effects:				
More than 50 below old MW	0.133 (0.122)	-0.154 (0.078)	0.032 (0.092)	0.092 (0.063)
30 to 50 below old MW	-0.114 (0.087)	-0.573 (0.073)	-0.298 (0.087)	-0.674 (0.082)
10 to 30 below old MW	-0.037 (0.182)	-0.180 (0.229)	-0.154 (0.125)	-0.254 (0.296)
At old minimum wage	-0.063 (0.068)	-1.305 (0.080)	-0.119 (0.061)	-1.330 (0.139)
Between old and new MW	0.063 (0.063)	-0.620 (0.052)	-0.065 (0.054)	-0.741 (0.104)
At new MW	0.287 0(.081)	1.054 (0.123)	0.250 (0.069)	0.819 (0.075)
10 to 30 above new MW	0.138 (0.092)	0.518 (0.063)	0.158 (0.082)	0.419 (0.066)
30 to 50 above new MW	0.119 (0.06)	0.144 (0.075)	0.055 (0.036)	0.151 (0.045)
50 to \$1 above new MW	0.002 (0.037)	0.055 (0.041)	0.067 (0.026)	0.076 (0.036)
\$1 to \$1.50 above new MW	-0.032 (0.076)	-0.093 (0.048)	-0.107 (0.044)	-0.022 (0.035)
\$1.50 to \$2.00 above new MW	-0.11 (0.03)	-0.031 (0.028)	-0.107 (0.044)	-0.051 (0.033)

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

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