Discrimination in Firing

Jan Kabátek Sacha Kapoor Alexander C. Kraaij Dinand Webbink*

November 27, 2025

Abstract

Do firms treat workers from different demographic groups differently when retention costs rise mechanically? We study this question using Dutch administrative data and statutory age-based increases in youth minimum wages, which create sharp, uniform shifts in retention costs at each birthday. Because these shifts are orthogonal to gender, migration background, and productivity, they isolate firing responses from voluntary turnover and allow us to detect differential treatment in the firing margin.

We find systematic demographic differences that vary with tenure. At the first birthday, native and migrant female workers face higher dismissal risks than native males, while migrant males do not exhibit a statistically significant firing gap relative to native males. At later birthdays, after at least a year of observed performance, migrant males become significantly less likely to be dismissed than native males. This pattern arises only among workers hired in loose youth labor markets, consistent with stricter hiring standards for migrant males when labor supply is abundant. In tight markets, migrant males exhibit no corresponding reduction in dismissal risk at later birthdays and face a sizable positive firing gap at the first.

Firm-level estimates reveal substantial heterogeneity in both baseline firing behavior and demographic firing gaps. Empirical Bayes methods indicate dispersion far beyond sampling noise, highlighting distinct clusters of firms with sharply different firing practices.

*Kabatek (j.kabatek@unimelb.edu.au): Melbourne Institute of Applied Economic and Social Research; Kapoor (kapoor@ese.eur.nl): Erasmus University Rotterdam; Kraaij (kraaij@wiso.uni-koeln.de): University of Cologne; Webbink (webbink@ese.eur.nl): Erasmus University Rotterdam and IZA. Kraaij gratefully acknowledges funding from the Reinhard Selten Institute, Universities of Bonn and Cologne. Results based on calculations by researchers from the University of Melbourne, Erasmus University Rotterdam, and the University of Cologne using non-public microdata from Statistics Netherlands. Under certain conditions, these microdata are accessible for statistical and scientific research. For further information: Microdata: Conducting your own research | CBS. We thank Anna Bindler, Robert Dur, Marten Ritterrath, and Dirk Sliwka for valuable comments. We also thank participants at the COMPIE, ESPE, EALE, ODISSEI, and World Labor (SOLE-EALE-AASLE) conferences, the Berlin Network of Labor Market Research (BeNA) Winter Workshop, the European Winter Meeting of the Econometric Society, the Workshop on Labour Economics in Trier, and the Erasmus School of Economics Brown Bag Seminar for helpful feedback. Outstanding research assistance was provided by Ada Sakarya and Lola Salaverri. All omissions and errors are our own.

Keywords: Systemic Discrimination, Gender Discrimination, Ethnic Discrimination, Labor Demand, Firing, Employee Turnover, Minimum Wages, Youth Employment, Immigrant Workers, Retail
JEL Codes: J23, J31, J61, J63, J71, L21, L81, M51

1 Introduction

Audit studies and correspondence experiments consistently show that applicants from historically disadvantaged groups receive fewer callbacks than equally qualified counterparts (Bertrand and Mullainathan, 2004; Kline et al., 2022; Neumark et al., 1996). This differential treatment provides strong evidence of discrimination in the early stages of the hiring process. By contrast, far less is known about how firms treat different demographic groups after hiring—whether discrimination persists among workers whom firms selected for employment, and what post-hiring treatment patterns reveal about its underlying mechanisms.

Once a worker is hired, the firm accumulates information about that worker's actual productivity, reliability, and fit. This learning process means that post-hiring treatment can reveal whether initial hiring disparities reflected inaccurate beliefs that get corrected through experience, or whether they reflect persistent preferences that survive direct observation of worker performance. For instance, if migrant applicants face stricter hiring standards, the marginal migrant hire may be more productive than the marginal non-migrant hire. A profit-maximizing firm forced to choose between these workers later on will fire the non-migrant hire.

This paper investigates discrimination in precisely these firing decisions. Studying firing discrimination at scale poses a fundamental challenge: administrative datasets typically record separations, which do not distinguish voluntary quits from involuntary dismissals (firings). We address this challenge by exploiting exogenous variation in the cost of retaining workers, induced by age-based step increases in the statutory minimum wage schedule. When retention costs increase uniformly across all workers, independent of their characteristics, firms face stronger incentives to dismiss employees. Observable spikes in separation rates around these cost increases allow us to isolate firing behavior from voluntary turnover and examine whether firms respond to identical cost increases differently for workers from different demographic groups.

We implement this approach using administrative data from the Netherlands, where minimum wages increase discontinuously at each birthday for workers aged 15 until the statutory youth rates cease at ages 21, 22, or 23, depending on the period. These legislated increases are substantial, averaging €0.73 per hour, which corresponds to 11.1% of the average gross hourly wage in our sample. They are orthogonal to all worker characteristics including gender, migration background, preferences, and productivity. Previous research shows that these mandated cost increases generate excess separations around birthdays, which are attributed to involuntary dismissals of minimum wage workers who are deemed too expensive to retain (Dayioglu et al., 2022; Kabátek, 2021; Korczak Krzeczowski, 2023; Kreiner et al., 2020). We exploit differences in excess separations between workers of different genders and migration backgrounds to investigate several questions: Do firms fire workers differentially depending on their gender and migration background? Do we observe some firms doing this more than others? And what do these firing patterns reveal about the nature of labor market discrimination?

We document systematic differences in firing patterns across demographic groups that vary between workers' initial and subsequent birthdays at the firm. This split accounts for tenure-related dynamics: passing the first birthday without separating may signal that the firm has a preference to retain the worker, while later birthdays occur after multiple statutory wage increases and may trigger different firing incentives. We find that native male workers approaching their first birthday are 0.68 percentage points more likely to separate than during other months. This is 15.1% of the 4.5% baseline monthly separation rate for native male workers. Native females face an additional 0.09 percentage point (1.9%) increase in the separation probability at the birthday compared to native males. Migrant males show no significant difference at the first birthday, while migrant females show an increase of 0.14 percentage points (2.7%).

The pattern changes notably for subsequent birthdays. Native males no longer show significant separation responses to wage increases, consistent with firms retaining more productive workers past the first year. However, migrant males become significantly *less* likely to be fired around subsequent birthdays than native males. The estimated firing gap is -0.22 percentage points (4.0%) (significant at 1%). This negative firing gap for migrants appears inconsistent with simple taste-based discrimination in firing. Instead, it may suggest that discrimination operates through differential hiring standards: if firms apply stricter criteria when hiring migrants, the marginal migrant worker who stays employed past the first birthday may be more productive than the marginal native worker.

We test this interpretation by examining heterogeneity in labor market conditions at the time of hiring (where tight markets are characterized by low unemployment and high vacancy rates, and loose markets by excess labor supply and elevated unemployment). We use a measure of youth-specific labor market tightness: the municipality-year share of 15–23-year-olds employed in supermarkets. In tight youth labor markets, firms face greater difficulty filling vacancies and thus have less scope to be selective, whereas in loose markets they can apply stricter hiring standards. Consistent with this mechanism, the negative firing gap for migrant males at their second and later birthdays is driven entirely by workers hired in loose labor markets (firing gap of -0.23 percentage points, or 4.2%, significant at 1%), with no statistically significant gap at their first birthday. Among those hired in tight markets, by contrast, migrants exhibit no statistically significant firing gap at later birthdays but a sizable positive gap of 0.45 percentage points (8.2%), significant at 1%, at their first birthday. This pattern supports the interpretation that lower firing rates among incumbent migrant workers reflect stricter hiring thresholds in loose labor markets, rather than employer preferences operating solely at the firing stage.

Pooled estimates mask substantial heterogeneity across firms. We estimate firm-specific firing gaps for the 50 largest firms and find wide dispersion in both baseline birthday firing behavior and demographic firing gaps. Kernel density estimates reveal far more variation in firm-specific parameters than would be expected from sampling variation alone, suggesting genuine heterogeneity in firm-level practices. We apply Empirical Bayes methods (Kline et al., 2022; Walters, 2024) to separate true heterogeneity from estimation noise and identify firms that are extreme in their treatment of different worker groups.

Our study makes several contributions to the literature on labor market discrimination. First, we add to the small body of evidence on discrimination in firing and layoff decisions (Dai et al., 2024; Egan et al., 2022). Beyond documenting firing patterns, we show how these patterns can be informative about discrimination at earlier stages, specifically how firing rates conditional on exogenous cost increases can reveal differential hiring standards across demographic groups.

Second, our approach provides a practical tool for detecting discrimination using observational data and statutory wage variation. Unlike correspondence studies, it involves no deception (Kessler et al., 2019) and can be applied to the universe of firms. Unlike experimental interventions, it leverages naturally occurring variation in labor costs. Firms could apply this method internally to audit practices across branches or managers. Regulators could use it to identify firms that warrant further investigation.

Finally, we contribute to the evaluation of youth minimum wage policies in the Netherlands (Van Bezooijen et al., 2024; Steenks et al., 2025). Our findings suggest these policies not only affect employment levels but also interact with discrimination in complex ways, both enabling age-based churning of workers and revealing heterogeneous treatment of different demographic groups.

The paper proceeds as follows. Section 2 provides institutional background on the Dutch minimum wage system. Section 3 describes our data sources and sample construction. Section 4 presents summary statistics and descriptive evidence on separation patterns. Section 5 introduces our baseline empirical strategy. Section 6 reports pooled results. Section 7 examines distributional results using Empirical Bayes methods. Section 8 concludes.

2 Institutional Context

The Netherlands implemented its first national minimum wage in 1969 under the Minimum Wage and Minimum Holiday Allowance Act ("Wet minimumloon en minimumvakantiebijslag", WML). Key to our analysis is the dependency of the minimum wage on age. 15-year-olds earn the lowest minimum wage and the minimum wage increases until age 23. The minimum wage-age schedule is revised annually on the basis of inflation, cost of living, and the state of the economy. Revisions are typically announced in January and implemented in July of the same year. Figure 1 visualizes the youth minimum wage scheme for the second half of 2025.

The age-dependent structure, by itself, can be viewed as a form of age discrimination. For the government, this approach balances employer and employee interests. There is no consensus on whether retaining youth workers in jobs where the minimum wage is binding aligns with their long-term interests or whether being let go and transitioning to alternative employment or further education is preferable. Moreover, minimum wage legislation does not address other forms of discrimination, such as those based on ethnicity or gender.

For our analysis, we focus on the grocery sector because it is the largest in terms of youth employment. Many employees take jobs in supermarkets, where they usually stock shelves or operate the checkout areas. Typically, they start at age 15 or 16. These jobs are attractive to them because they can work on weekdays after school hours and on weekends. The employers benefit from a workforce that they can pay relatively low hourly wages.

All employers in the Netherlands, including supermarkets, must pay at least the minimum wage specified for each age. On top of this, many sectors are legally bound by collective bargaining agreements which specify age-specific wages that exceed the legal minimum wage. Supermarkets must pay wages at least as high as those specified by the collective agreement of their sector. This agreement not only specifies minimum wages by age, but also by job responsibilities and years of experience within the same job. Further, the agreements are updated every few years following ongoing collective bargaining.

The age-wage profile for youth employment in the Netherlands incentivizes employers to hire young workers and fire them as they age. This is because the cost of retaining a worker

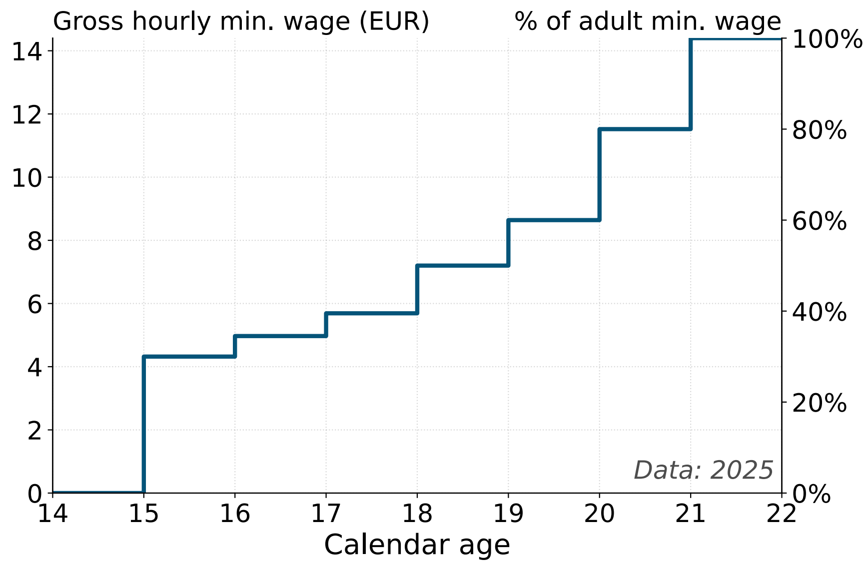

Figure 1: Hourly Minimum Wage in the Netherlands

Notes: Own illustration based on Rijksoverheid (2025).

increases discontinuously on the birthday. If the productivity differences between older and younger youth workers are sufficiently small, firms can benefit from replacing older workers with younger ones.

3 Data

We draw on several administrative datasets that are compiled and maintained by Statistics Netherlands (CBS). Our dataset includes before-tax wages, hours of work, sector, the start and end date of the employment spell if the spell has ended, and demographic characteristics.

The labor market data are taken from the SPOLISBUS datasets of 2006–2020 (Centraal Bureau voor de Statistiek, 2020). This is individual-level data from the tax office and contains a wide range of information on employment start and end dates, wages, working hours, etc. We link the labor market data to basic demographic data on birthdays, gender, and migration background from the dataset GBAPERSOONTAB (Centraal Bureau voor de Statistiek, 2022).

The labor market data consist of employment spells, which we convert into a monthly panel format. We define separations strictly, excluding temporary employment interruptions. Workers are classified as separated only when they cease to be observed in the same firm between 2006 and 2020. Moreover, potential separations occurring in the final month of the observation period are excluded from classification as separations.

The employment spell data, combined with birthdays from the basic demographic data, allow us to construct a daily panel dataset.[1] However, we use a monthly dataset for computational tractability. To preserve the information from sharp wage increases occurring on birthdays, we adjust the panel so that each worker's monthly observations begin on their birthday.

The labor market data indicate for which firm each employee works. However, they do not specify the branch or municipality in which an employee works. We define a branch as the combination of a firm and a municipality. Because a firm may operate multiple branches within a single municipality, this definition is necessarily coarse but remains informative. CBS allocates employees to the different branches of their firms based on their place of residence and the number of jobs each firm reports for its branches. Because the firm reports are collected in December, branch identifiers are typically missing in the other months. We therefore impute missing branch identifiers for worker-month observations in other months, prioritizing branch identifiers recorded in the past and relying on future branch identifiers only when past information is unavailable.

We include scores from nationwide exams taken around age 12, called CITO, as a proxy for cognitive skills. These are the standard scores from the dataset CITOTAB (Centraal Bureau voor de Statistiek, 2005). From the same dataset, we include the teacher's advice for a high school track after learning the student's CITO test score as a proxy for noncognitive skills.

We hand-collect data on minimum wages by age, years of experience, and job responsibilities from the legally binding collective agreement in the grocery sector for each applicable period. These are used to define the age-specific wage increases of each employee, which is one of the key variables in our analysis. Specifically, this is the difference between the employee's next minimum wage and their current minimum wage. We do not use actual wages here, as some firms may uniformly pay a small wage premium to all employees. We assume that firms currently paying this premium will continue to do so after the next age-related wage increase.

We also use the hand-collected minimum wage data to identify employees' hierarchical levels. For each combination of age and years of experience, the binding agreement specifies a minimum wage for one of nine sets of job responsibilities, labeled A to I. By comparing actual wages in the administrative records with the corresponding minimum wages, we infer

[1] The basic demographic data include birthdays only when access has been granted by the data provider.

employees' responsibility sets by assigning them to the highest set whose minimum wage does not exceed their actual wage. Based on these assigned sets, we classify employees into four categories: no manager (A–B), lower manager (C–D), middle manager (E–G), and upper manager (H–I). Lower managers are typically team leaders responsible for a shift or department within a store. Middle and upper managers are usually store managers overseeing one or more local branches of the firm. As additional control variables, we construct dummy variables at the firm-municipality level indicating whether a worker's own ethnicity is represented in lower, middle, or upper management (Benson et al., 2023).

We apply several sample restrictions. We keep regular employees and on-call workers in our sample, thereby excluding directors, interns, temps, and those employed under The Law on Social Employment[2] We also include only those who are employed at the firm for at least three months because the probation time is two months, thereby not counting separations within or at the end of the probation time for new employees.

We exclude observations with clearly implausible values for wages, working hours, or hourly wages. Specifically, we drop cases with monthly wages below €0 or above €6,000, monthly hours below 0 or above 300, and hourly wages exceeding €40. Finally, we keep only workers between the ages of 15 and 25 (inclusive). The control variables measuring representation in management are defined before restricting the sample by age.

We impute missing values for three control variables. First, we impute the CITO standard score using the sample mean. Second, we impute the CITO teacher's advice using the sample mode. Third, we impute the dummy variable indicating whether the worker has a flexible contract using the sample mode, which equals 0. In all regression models that include these variables, we add indicator variables denoting whether each original value was imputed.

4 Descriptive Evidence

This section presents descriptive statistics and visual evidence on variables relevant to our analysis of firing patterns in the grocery sector in the Netherlands for the sample period of 2006–2020.

Table 1 presents descriptive statistics for the main and control variables used in the analysis. We first turn to the firm- and branch-level variables. Average firm size is computed as the mean number of workers per month across all months in which a firm is observed.

[2] In Dutch: Wet sociale werkvoorziening (WSW). This law is for those who cannot work under standard circumstances due to a physical, psychological, or intellectual disability (Overheid, 1997).

Average firm size has a mean of 227 workers with a standard deviation of 1,664. Average branch size is constructed analogously, using the firm-municipality level to proxy for firm branches.

Next, we discuss descriptive statistics for the worker-level variables, which do not change over time. The share of females is 57.1%, making them the majority gender in the Dutch grocery sector. Age is defined as the worker's average age across observed months. Its mean is 18 with a standard deviation of 1.7. Cognitive ability is proxied by standardized CITO test scores, typically taken at age 11 or 12. Although workers take the test in different years, scores are comparable across cohorts (Centraal Bureau voor de Statistiek, 2005). The average score is 535 with a standard deviation of 6.9. 74.2% of the workers do not have a migration background, 5.1% are first-generation migrants, and 20.7% are second-generation migrants. The largest ethnic groups among migrants are Moroccan (4.8%), Turkish (4.1%), and Surinamese (3.2%). The other ethnic groups are pooled and represent 13.6% of the workforce. Contract duration, measured in months, equals the maximum firm tenure observed for each worker. It has a mean of 17.8 months with a standard deviation of 15.9.

Next, we report summary statistics for worker-month-level variables. Monthly working hours average 35.1 with a standard deviation of 32.9. Gross hourly wages average €6.55 with a standard deviation of 3.74. The statutory wage increase on the next birthday averages €0.73 with a standard deviation of 0.42. On average, 19.8% of workers are on permanent contracts in any given month, 2.5% work full time, 29.2% have a flexible contract, and 40.8% are temporary workers. The final three variables indicate whether a worker shares their ethnic background with at least one lower, middle, or upper manager. These shares are 79.1%, 67.2%, and 34.0%, respectively.

Finally, we report the distribution across managerial levels. These statistics are also at the worker-month level and include more observations because the sample is not restricted by age in order to identify all managers. Among worker-month observations, 80.5% are non-managers, 13.1% are in lower management, 5.5% in middle management, and 0.8% in upper management.

Figure A1 displays the separation rate over time. It fluctuates around 5 percent during our sample period. Elevated separation rates near the end of the calendar year may reflect that employment decisions are made at a higher rate as the new year approaches (Kabátek, 2021). Other spikes around August may reflect the end of summer jobs (Kabátek, 2021). The largest spike in the separation rate around the start of 2017 can be attributed to a policy change studied by Van Bezooijen et al. (2024). Importantly, these peaks are not relevant to

Table 1: Descriptive Statistics

	Mean	SD	N
Average firm size over time	227	1,664	1,183
Average branch size over time	17	62	18,193
Female	0.571	–	1,316,829
Age (panel average)	18	1.7	1,316,829
Cognitive ability	535	6.9	1,316,829
Native	0.742	–	1,316,829
1st-generation migrant	0.051	–	1,316,829
2nd-generation migrant	0.207	–	1,316,829
Morocco	0.048	–	1,316,829
Türkiye	0.041	–	1,316,829
Suriname	0.032	–	1,316,829
Other	0.136	–	1,316,829
Max. firm tenure (months)	17.8	15.9	1,316,829
Working hours (monthly)	35.1	32.9	27,601,181
Gross hourly wage (€)	6.55	3.74	27,601,181
Next raise (€)	0.73	0.42	27,601,181
Raise if close to birthday (€)	0.12	0.31	27,601,181
Permanent contract	0.198	–	27,601,181
Full time	0.025	–	27,601,181
Flexible contract	0.292	–	27,601,181
Temp job	0.408	–	27,601,181
Ethnicity repr. in lower mgmt.	0.791	–	27,601,181
Ethnicity repr. in middle mgmt.	0.672	–	27,601,181
Ethnicity repr. in upper mgmt.	0.340	–	27,601,181
Not a manager	0.805	–	34,439,973
Lower manager	0.131	–	34,439,973
Middle manager	0.055	–	34,439,973
Upper manager	0.008	–	34,439,973

Notes: Descriptive statistics for all firms and workers in the Dutch grocery sector from 2006–2020, based on administrative records from Statistics Netherlands. Firm (branch) size is measured at the firm (branch) level. Female, age, cognitive ability, migration and ethnic background, and contract duration are measured at the worker level. Cognitive ability is proxied by standardized CITO tests taken at the end of primary school. Working hours, hourly wages, raises, contract type, full-time status, ethnic representation in management, and hierarchical level are measured at the worker-month level.

our empirical approach because we exploit variation in the cost of retaining workers based on the time relative to birthdays rather than calendar time, which we expect to result in elevated separation rates around workers' birthdays.

This is what we find in Figure 2, which displays the separation rate over the time relative to the birthday. The separation rate starts increasing two months prior to the birthday, reaches its peak one month before the birthday, and remains somewhat elevated one month after the birthday. This pattern is consistent with the findings of Kabátek (2021), based on data from 2006–2011. The difference between the highest and lowest points of the separation rate in this graph is roughly 0.4 percentage points or 8.5%. The pattern is precisely estimated and consistent with the increase in the cost for employers of retaining workers at each birthday.

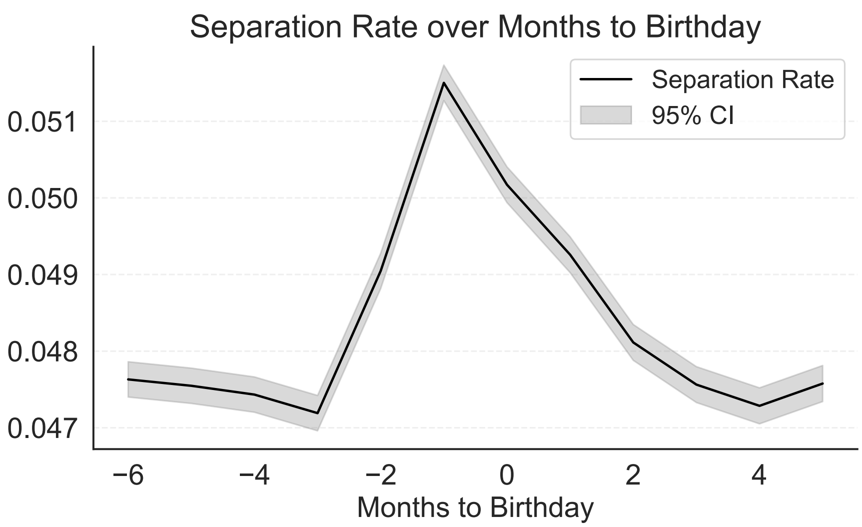

Figure 2: Separation Rate over Birthday Proximity

Notes: Shaded areas indicate 95% confidence intervals. The figure plots the separation rate over birthday proximity. Based on all worker-month observations.

One might wonder whether the spikes in separation rates are driven by other types of separations, specifically temporary contracts that are not extended and quits. However, such separations are unlikely to generate the sharp birthday-related pattern we observe. Insofar as temporary contracts start on different dates, they would tend to raise separation rates more evenly throughout the year. Furthermore, such separations are similar to firings in the sense that they involve a firm's decision to let workers go.

As for quits, the elevated separation rate around, and especially before, the birthday could in principle result from workers voluntarily leaving their jobs shortly before receiving a wage increase. One possible reason, explored by Kabátek (2021), is that workers, anticipating the higher hourly wage, switch to jobs with fewer hours while maintaining the same income. However, Kabátek (2021) found no empirical support for this mechanism, making it an unlikely explanation for the observed pattern.

Panel (a) of Figure 3 presents, again in the same format, the separation rate by birthday proximity, this time disaggregated by demographic group. For the main analysis, we focus on a 2×2 split by gender and migration background, resulting in four groups: native males, native females, migrant males, and migrant females.[3] The graph reveals several notable patterns. First, the elevated separation rates near the birthday are present across all four groups. Second, the separation rates of the four groups do not intersect: native male workers consistently have the lowest separation rate relative to the birthday, followed by native females, migrant females, and migrant males. Third, while native females are, on average, more likely to separate than native males, this pattern is reversed among migrant workers. Finally, the spike in separations at the birthday appears less pronounced for migrant male workers than for the other groups.

The elevated separation rates near the birthday may be partially explained by the chain rule (in Dutch: ketenregeling), which stipulates that employees who have held one or more temporary contracts with the same employer for a certain duration must either be offered a permanent contract or be let go (e.g., Kabátek et al., 2023). To account for this, we residualize the separation rate on the proximity to eligibility for a permanent contract under the chain rule by regressing the separation indicator on this proximity measure and using the residuals, re-centered at the sample mean, in Panel b of Figure 3. The resulting patterns remain largely unchanged. In Panels c–f, we add further controls to the residualization. Panel c shows that the patterns do not notably differ when accounting for tenure at the firm. This could be important if workers start the job near their birthdays and their contracts end exactly one, two, or three years later. In Panel d, we residualize on the CITO variables, our proxies for cognitive and non-cognitive ability. This only slightly changes the figure. Panel e adds dummy variables indicating whether a worker's ethnicity is represented in lower, middle, or upper management, alongside the proxies for ability. This aligns the separation rates of three groups almost exactly, slightly increasing the rate for native males and substantially reducing it for migrant females. The separation rate of native females appears unaffected by the inclusion of representation controls, while that of migrant males shifts downward but

[3] We use the term "migrant" as shorthand for workers with a migration background. We do not distinguish between first- and second-generation migrants in the main analysis.

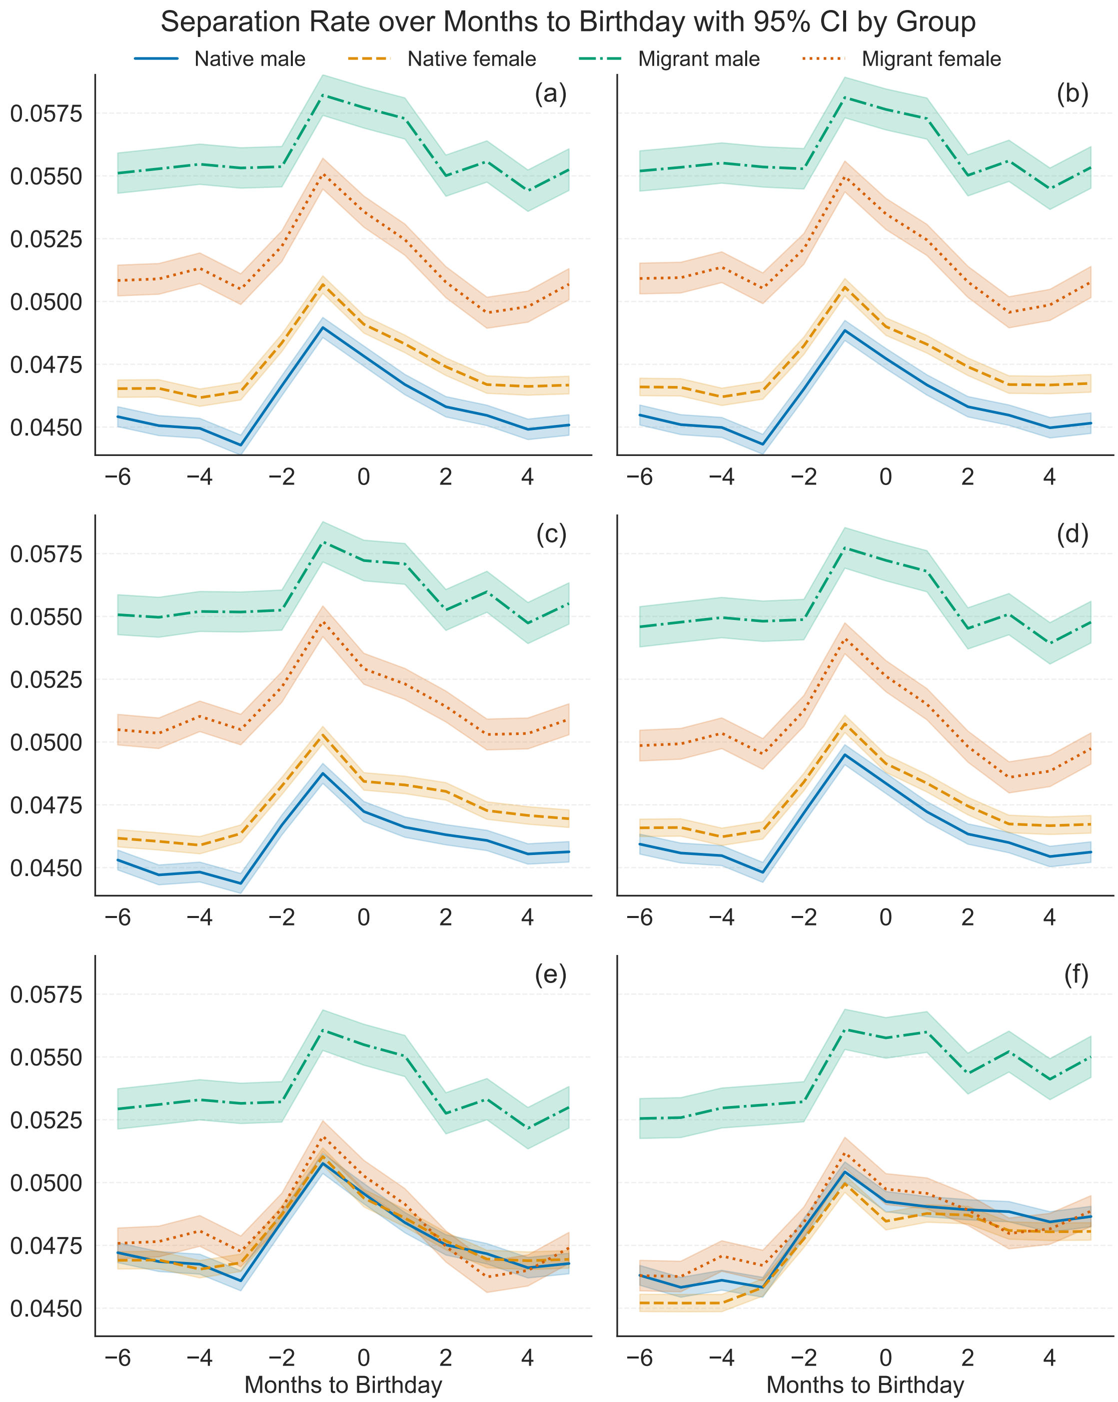

Figure 3: Separation Rate over Birthday Proximity by Group: Raw and Residualized

Notes: Shaded areas indicate 95% confidence intervals. Panel a shows the raw separation rate. Panels b–f show separation rates residualized using different sets of controls: b, chain rule timing; c, tenure at firm; d, CITO score and advice; e, CITO score and advice, and representation in management; f, chain rule timing, tenure at firm, CITO score and advice, representation in management, hierarchical level, and age dummies. The residualized separation rate is constructed by regressing the separation dummy on the specified controls and re-centering the residuals at the sample mean.

remains notably above the rates of other groups. These observations suggest that it matters for the separation rates of native male and migrant female workers whether their ethnic background is represented in the branch's management, but not for the separation rates of native females and migrant males. Finally, when the separation rates are adjusted for chain rule timing, tenure, proxies for ability, representation in management, hierarchical level of the worker, and age dummies, in Panel f the observed pattern changes for all groups. The pre-birthday separation rates are slightly lower than in Panel e, the peak is slightly lower (except for migrant males), and the separation rates remain elevated after the birthday. The strong divide in separation rates between migrant males and the other groups remains.

We divide the firms in our data by firm size, where firm size is defined as the monthly number of workers in each firm, averaged over all months in which the firm is observed. We distinguish between the 50 largest firms and all other firms. Figure A2 shows that the average workforce in the 50 largest firms is 3,856, much higher than the average of 66 in the other firms.

These firms differ not only in size but also in separation patterns. Figure 4 shows that the largest firms have lower separation rates across the four groups, with migrant males separating from those firms at a lower rate than native males from the other firms. Additionally, the birthday spikes appear more consistent across groups for the largest firms than for the others. In the other firms, the separation rate of migrant males peaks at the birthday, whereas the spike for migrant females occurs one month before the birthday. The birthday spikes in the smaller firms also appear less pronounced than those in the 50 largest firms, suggesting that separation timing in the largest firms may reflect more strategic firing behavior. These differences between the 50 largest firms and the other firms motivate our choice to focus on these firms in the firm-level analysis in Section 7.

5 Empirical Strategy

We rely on discontinuities in the cost of retaining workers, generated by minimum wage increases at birthdays, to identify the firing gaps of interest. Specifically, we estimate the extent to which the four groups (migrant–native and female–male) are differentially treated in firms' firing decisions around birthdays. The functional form of our baseline model spec-

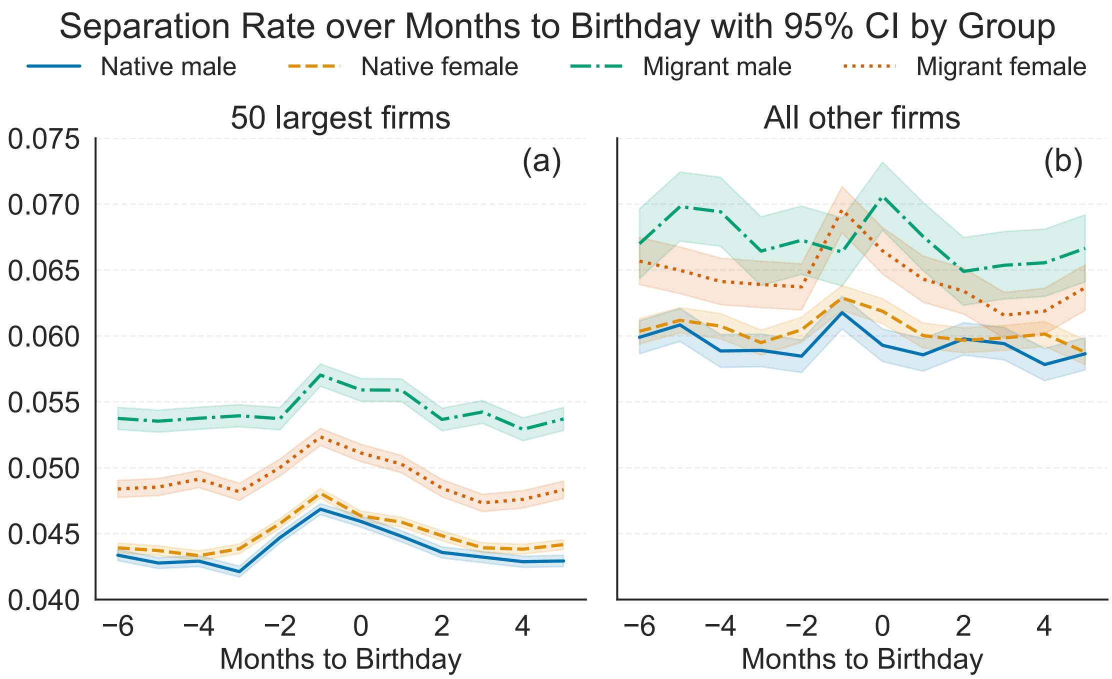

Figure 4: Separation Rate over Birthday Proximity by Group: 50 Largest Firms vs. Smaller Firms

Notes: Shaded areas indicate 95% confidence intervals. The separation rates are not residualized.

ification is as follows:

$$Y_{ift} = \beta_0 + \sum_{k=1}^{3} \beta_{1k}\text{Group}_{ik} + \beta_2 \mathbb{1}\{\text{Close to BD}\}_{it} \\ + \sum_{k=1}^{3} \beta_{3k}\mathbb{1}\{\text{Close to BD}\}_{it} \times \text{Group}_{ik} + \boldsymbol{X}_{ift}\gamma + \xi_{ft} + \varepsilon_{ift}. \tag{1}$$

Here Y_{ift} is a separation dummy for individual (worker) i in firm f at time t (monthly), Group has native male as the reference category and further distinguishes between native females, migrant males, and migrant females, and $\mathbb{1}\{\text{Close to BD}\}$ is defined as the time period t being one or two months prior to the worker's birthday. This is the period when separation rates peak (see Figure 3). We exclude the birthday month and the month that follows to ensure that the estimates are conservative and primarily capture strategic, preemptive firings, rather than separations driven by expiring one-year contracts initiated after birthdays or by birthday-related behavioral incidents.[4] $\boldsymbol{X}_{ift}$ denotes the set of control variables which

[4]Studying a different question, Bindler et al. (2023) control for potential behavioral shocks that may occur on birthdays.

consists of age dummies, a manager dummy (pooling lower, middle, and upper managers), dummies for worker's ethnicity represented in lower, middle, and upper management, branch size quintiles (where branch refers to the firm-municipality combination), the CITO standard score (cognitive ability), the CITO teacher recommendation (non-cognitive ability), and interactions of the ability and representation variables with the close-to-birthday indicator variable. ξ_{ft} denotes firm-month fixed effects. Standard errors are clustered at the firm level.

The firing gaps of interest are identified by $\boldsymbol{\beta}_3$. These could reflect (i) direct discrimination in firing, (ii) discrimination in hiring, and (iii) productivity differences observable to firms but unobserved to us. We address (ii) by studying heterogeneity with respect to labor market tightness at the time of hire. Regarding (iii), we make three mitigating assumptions. First, baseline separation rates account for heterogeneity in productivity and outside options. Second, the difference between quits close to the birthday and other quits is the same across groups. Third, the firing probability as a function of productivity is the same close to the birthday and far from the birthday. We relax this third assumption by interacting our ability proxies with the close-to-birthday dummy. The intuition of our empirical approach is that firms tend to fire workers around their birthdays due to statutory wage increases, and potentially treat groups differently in doing so.

Employing a simple OLS model in this setting comes with the drawback that workers move back and forth between being close to their birthdays and not being close to their birthdays as their time at the firm progresses. There are potential dynamic effects because workers passing the first birthday without exiting the firm reveal a preference of the firm to retain these workers, for productivity reasons or other considerations. On the other hand, workers passing multiple birthdays may be more likely to be let go given that they have received several statutory wage increases. To address these concerns, we estimate the regression model of Equation (1) separately for those who have not passed any birthday at the firm yet, those who have passed their first birthday at the firm, those who have passed their second birthday at the firm, and so on. Having inspected separate results up to the fourth birthday at the firm, we show a simple two-way split in the remainder of the paper: workers who are approaching their first birthday at the firm, and those who are approaching any subsequent birthdays.

After presenting results for these models, we also turn to a richer specification incorporating not only birthday proximity but also the magnitudes of statutory wage increases attached to those birthdays. We estimate

$$\begin{aligned} Y_{ift} = \beta_0 + \sum_{k=1}^{3} \beta_{1k}\text{Group}_{ik} + \beta_2 \Delta\text{Wage if close to BD}_{it} \\ + \sum_{k=1}^{3} \beta_{3k} \Delta\text{Wage if close to BD}_{it} \times \text{Group}_{ik} + \boldsymbol{X}_{ift}\boldsymbol{\gamma} + \xi_{ft} + \varepsilon_{ift}, \end{aligned} \tag{2}$$

where ΔWage if close to BD_{it} is the product of the indicator for birthday proximity and the statutory wage increase the worker is due to receive on that birthday. All other terms are defined as in Equation (1). In this specification, β_2 and $\boldsymbol{\beta}_3$ are interpreted as changes in the separation rate per €1 increase in statutory wages near birthdays.

6 Pooled Results

Table 2 presents the regression results using the full population of firms in the Dutch grocery sector. The dependent variable is the separation dummy multiplied by 100, so that the estimates should be interpreted in percentage points. Column 1 is estimated on the full population of workers approaching their first birthday at their firm. According to the coefficient of Close to birthday in Column 1, native male workers are on average 0.68 percentage points more likely to be fired close to their first birthdays, which corresponds to 15.1% of their baseline separation rate of roughly 4.5% (see Figure 3). This is significant at the 1% level. Column 2 presents estimates of workers who are still employed by the firm after their first birthday and are approaching any subsequent birthday at the firm. The estimate shrinks toward zero (0.04 percentage points) and is no longer significant when analyzing the subsequent birthdays, which may suggest that native male workers whom the firm kept through their first birthdays are relatively productive.

The interactions with the gender and migration background categories tell us how much more or less likely each of the groups is to be fired around their birthdays than native males close to their birthdays. We refer to these as firing gaps. The firing gap for native females is 0.09 in Column 1, meaning that native females are on average 0.09 percentage points more likely to be fired around their first birthdays than natives. This is significant at the 5% level and corresponds to 1.9% of their baseline separation rate of roughly 4.7%. The corresponding estimates in Column 2 no longer reveal a statistically significant firing gap for native females.

The firing gap for migrant males is -0.05 percentage points and not significant at 10

Table 2: Pooled Regression Results: Main Specification

	Dep. var.: $\mathbb{1}\{\text{Separation}\} \times 100$	
	(1)	(2)
Close to birthday	0.681***	0.0394
	(0.158)	(0.146)
Native female	-0.481***	-0.0391
	(0.0462)	(0.0485)
Migrant male	-0.646***	0.362***
	(0.117)	(0.0585)
Migrant female	-1.643***	-0.159*
	(0.170)	(0.0877)
Close to birthday	0.0894**	0.0571
×Native female	(0.0381)	(0.0411)
Close to birthday	0.0492	-0.216***
×Migrant male	(0.0841)	(0.0650)
Close to birthday	0.144*	-0.0221
×Migrant female	(0.0754)	(0.0771)
Close to birthday	0.00124	0.0128***
×CITO score, demeaned	(0.00375)	(0.00324)
R^2	0.113	0.090
Description	1st BD	2nd+ BD
Controls	✓	✓
Firms	1,070	953
Workers	1,228,168	837,753
Observations	7,779,389	10,774,618

Notes: Regression results of all firms and workers in the Dutch grocery sector from 2006–2020, based on administrative records from Statistics Netherlands. The dependent variable is a dummy for worker separation, scaled by 100; coefficients are thus interpreted as percentage point changes. Regressors include gender, migration background, and their interactions with the wage change close to the birthday. The reference group is native Dutch men. Columns 1 and 2 split the sample by whether workers approach their first or a subsequent birthday at the firm. Control variables consist of age dummies, a manager dummy (pooling lower, middle, and upper managers), dummies for own ethnicity represented in lower, middle, and upper management, branch size quintiles (where branch refers to the firm-municipality combination), the CITO standard score (cognitive ability), the CITO teacher recommendation (non-cognitive ability), interactions of the ability and representation variables with the close-to-birthday indicator, and firm-month fixed effects. Standard errors clustered at the firm level in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

percent in Column 1. In Column 2, however, the firing gap becomes large and negative at 0.22 percentage points, corresponding to 4.0% of their baseline separation rate of 5.5%. This is significant at the 1% level. This suggests that migrants who are able to stay at the firm through their first birthday are on average less likely to be dismissed close to their subsequent birthdays than native males. Migrant females are on average 0.14 percentage points more likely to be fired close to their first birthdays, or 2.7% from their baseline of roughly 5.1%, though this is only significant at the 10% level. Turning to Column 2, we find no statistically significant firing gap between migrant females and native males with respect to the subsequent birthdays.

The negative firing gap between migrant and native male workers raises the question of why migrants are less likely to be dismissed around their subsequent birthdays than native workers. One possible explanation is differential hiring standards: if the productivity bar is higher for migrants to be hired, incumbent migrant workers may, on average, be more productive than their native counterparts. To assess the relevance of this explanation, we examine whether the migrant firing gap varies with labor market conditions at the time of hiring. In loose labor markets, defined by excess labor supply and elevated unemployment, firms have more room to be selective and potentially discriminate than in tight markets, which are characterized by low unemployment and high job vacancy rates. If migrant workers face stricter hiring thresholds, we would expect the firing gap to become more negative when hiring occurs in loose labor markets and to narrow toward zero or positive values in tighter markets.

Specifically, we construct a municipality-year-level measure of labor market tightness: the share of 15–23-year-olds employed by supermarkets. This employment-based measure is preferable to an unemployment-based one, particularly if young workers are unlikely to register as unemployed. We estimate the same model using a median split on this measure and find results consistent with the interpretation that firms have more room to discriminate when hiring in loose labor markets than in tight ones. Figure 5 presents these estimates in a coefficient plot, and Table A1 reports them in table format.

The top panel of Figure 5 reproduces the estimates for all workers from Table 2 and is included here for comparison. The middle and bottom panels present estimates for workers hired in loose and tight labor markets, respectively.

Among native males approaching their first birthday at the firm, those hired in loose labor markets are 0.82 percentage points (18.2%) more likely to be dismissed when near their birthday compared to when they are farther from it (significant at the 1% level). For

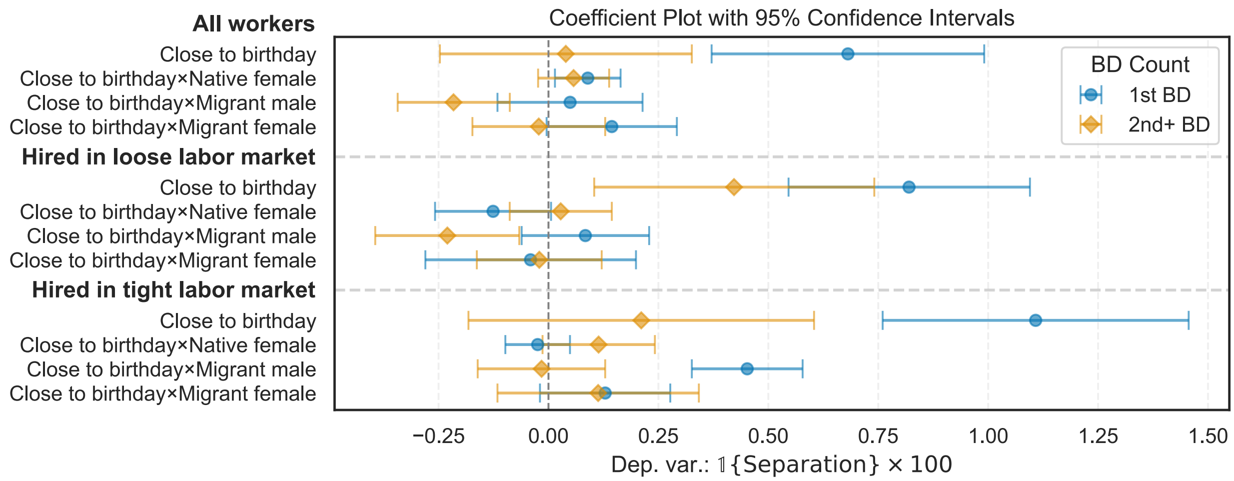

Figure 5: Regression Results: Sample Split by Labor Market Tightness at Time of Hire

Notes: Bars display 95% confidence intervals. The dependent variable is a dummy for worker separation, scaled by 100; coefficients are thus interpreted as percentage point changes. Each of the three sets starts with the Close to birthday coefficient, followed by its interactions with the three groups. Labor market tightness is measured at the municipality-year level as the share of 15–23-year-olds employed in supermarkets. A median split on this measure is used to define tight and loose labor markets: the middle panel shows estimates for hires in loose markets, and the bottom panel for tight markets. The full sample includes 11,992,125 observations of 585,175 male workers at 968 firms and 17,243,108 observations of 786,983 female workers at 1,043 firms.

those hired in tight labor markets, the increase is larger at 1.11 percentage points (24.7%) and also significant at the 1% level. However, these differences should be interpreted with caution due to wide confidence intervals.

For native males approaching any subsequent birthday, the probability of being fired near the birthday is higher by 0.42 percentage points (9.3%) in the loose-market group (significant at the 1% level) and by 0.21 percentage points (4.7%) in the tight-market group (not significant at the 10% level), relative to non-birthday periods. Again, while the point estimates differ, they lack statistical precision. The difference between first- and subsequent-birthday effects, however, is relatively precisely estimated.

Native females hired in loose labor markets exhibit a firing gap of −0.13 percentage points (2.8%) relative to native males at their first birthday (significant at the 10% level). This gap nearly disappears among those hired in tight markets (0.02 percentage points, 0.4%, not significant at the 10% level). For subsequent birthdays, the firing gaps differ across labor-market conditions: 0.03 percentage points (0.6%, not significant at the 10% level) in loose markets and 0.11 percentage points (2.3%, significant at the 10% level) in tight markets.

Migrant males display the clearest heterogeneity. As shown in Table 2, they are 0.22

percentage points less likely to be dismissed than native males near their second or later birthdays at the firm, while showing no significant difference at their first birthday. When restricting the sample to those hired in loose labor markets, both estimates closely mirror the full-sample results: the subsequent-birthday firing gap is −0.23 percentage points (4.2%, significant at 1%), and the first-birthday gap likewise remains close to zero and insignificant. In contrast, among those hired in tight markets, the pattern reverses. Migrant males show a positive firing gap of 0.45 percentage points (8.2%) at their first birthday, while the gap at subsequent birthdays is again close to zero and insignificant.

These results point toward an explanation based on differential hiring standards across labor-market conditions. In loose labor markets, firms may apply stricter standards to migrant males than to native males, so those migrant males who are hired are relatively productive and therefore less likely to be dismissed later. In tight labor markets, however, firms have much less scope to apply such differential standards in hiring, and migrant males are more likely to be dismissed when they approach their first birthday at the firm. Among those not dismissed at their first birthday, there is no statistically significant difference relative to native males at subsequent birthdays, consistent with the idea that retention itself indicates that a worker is sufficiently productive.

Migrant females show little heterogeneity in firing patterns by labor market tightness. Only their first-birthday firing gap of 0.13 percentage points (2.5%) when hired in tight labor markets is significant at the 10% level.

Next, we consider not only proximity to birthdays but also the magnitude of the statutory wage increases workers receive on their birthdays. Table 3 presents results based on Equation (2), where the birthday-proximity indicator is interacted with the corresponding statutory wage increase.

For native males, the probability of dismissal increases by 0.39 percentage points (8.7%) for each €1 wage increase near their first birthday at the firm (significant at the 10% level). For subsequent birthdays, the estimate is −0.20 percentage points (4.4%, not significant at the 10% level).

Native females are 0.15 percentage points (3.2%) more likely to be dismissed near their first birthday for each €1 of mandated wage growth (significant at 1%). The corresponding estimate for subsequent birthdays is 0.08 percentage points (1.7%, significant at the 10% level).

For migrant males, the estimated firing gaps are 0.24 percentage points (4.4%) near their first birthday (significant at 1%) and −0.15 percentage points (2.7%, also significant at 1%)

Table 3: Pooled Regression Results: Accounting for Wage Increases

	Dep. var.: $\mathbb{1}\{\text{Separation}\} \times 100$	
	(1)	(2)
ΔWage if close to birthday	0.387*	-0.200
	(0.229)	(0.154)
Native female	-0.484***	-0.0439
	(0.0426)	(0.0560)
Migrant male	-0.708***	0.322***
	(0.126)	(0.0532)
Migrant female	-1.729***	-0.191**
	(0.200)	(0.0884)
ΔWage if close to birthday	0.150***	0.0840*
×Native female	(0.0521)	(0.0493)
ΔWage if close to birthday	0.235***	-0.151***
×Migrant male	(0.0768)	(0.0582)
ΔWage if close to birthday	0.442***	0.0715
×Migrant female	(0.0657)	(0.0816)
R^2	0.113	0.090
Description	1st BD	2nd+ BD
Controls	✓	✓
Firms	1,070	953
Workers	1,228,168	837,753
Observations	7,779,389	10,774,618

Notes: Regression results of all firms and workers in the Dutch grocery sector from 2006–2020, based on administrative records from Statistics Netherlands. The dependent variable is a dummy for worker separation, scaled by 100; coefficients are thus interpreted as percentage point changes. The main independent variable is Wage increase if close to BD, which equals zero for workers not close to their birthday and the amount of the statutory wage increase (in €) for those who are. Regressors include gender, migration background, and their interactions with Wage increase if close to BD. The reference group is native Dutch men. Columns 1 and 2 split the sample by whether workers approach their first or a subsequent birthday at the firm. Control variables consist of age dummies, a manager dummy (pooling lower, middle, and upper managers), dummies for own ethnicity represented in lower, middle, and upper management, branch size quintiles (where branch refers to the firm-municipality combination), the CITO standard score (cognitive ability), the CITO teacher recommendation (non-cognitive ability), interactions of the ability and representation variables with the close-to-birthday indicator, and firm-month fixed effects. Standard errors clustered at the firm level in parentheses. * p $<$ 0.1, ** p $<$ 0.05, *** p $<$ 0.01.

near their subsequent birthdays. This statistically significant first-birthday result contrasts with the weaker and insignificant estimate in the pooled specification reported in Table 2.

Migrant females have a first-birthday estimate of 0.44 percentage points (8.6%, significant at the 1% level) and 0.07 percentage points (1.4%, not significant at the 10% level) for subsequent birthdays.

We extend this model specification in Section 7 by estimating it separately for each of the largest firms to explore whether, and to what extent, the firing gaps vary across firms.

7 Distributional Results

We estimate firm-specific regressions (building on Equation (2)) for each of the 50 largest firms in the Dutch grocery sector to analyze heterogeneity in firing behavior across firms. Our aim is to present the distributions of four key coefficients for each firm: the baseline coefficient for native males and the three interaction terms capturing firing gaps for the other groups.

To achieve this, we follow the three-step empirical Bayes procedure outlined in Walters (2024). The first step involves estimating the firing gap coefficients and standard errors for each firm. Second, we estimate the prior distribution using our data by pooling these estimates. This step is commonly referred to as deconvolution (Walters, 2024). Finally, we obtain posterior predictions for each firm using the estimated prior. We follow Kline et al. (2022) and Walters (2024) by using a nonparametric deconvolution algorithm (Efron, 2016), as well as considering the linear shrinkage estimators that result from assuming a normal prior distribution (Walters, 2024).

The empirical Bayes approach is motivated by the fact that heterogeneity in the raw estimates can reflect true heterogeneity in the firm parameters on the one hand and noise on the other. The deconvolution procedure aims to separate the two by estimating the distribution of true firm-level parameters, effectively filtering out sampling noise. The posterior predictions then combine this estimated prior with each firm's data to produce improved estimates that balance firm-specific information against what we learn from the broader distribution (Kline et al., 2022; Walters, 2024).

We begin with step 1 by estimating Equation (2) separately for each firm:

$$
\begin{aligned}
Y_{it}^f = \beta_0^f + \sum_{k=1}^{3} \beta_{1k}^f \text{Group}_{ik} + \beta_2^f \Delta\text{Wage if close to BD}_{it} \\
+ \sum_{k=1}^{3} \beta_{3k}^f \Delta\text{Wage if close to BD}_{it} \times \text{Group}_{ik} + \boldsymbol{X}_{it}^f \boldsymbol{\gamma}^f + \xi_t^f + \varepsilon_{it}^f.
\end{aligned}
\tag{3}
$$

We then collect estimates of β_2^f and $\boldsymbol{\beta}_3^f$ for the 50 largest firms, following Kline et al. (2022). To account for the remainder of the sample, we aggregate all other firms into a composite 51st firm. Unlike the earlier presentation of the results, we report coefficients in their natural scale rather than multiplying them by 100. Finally, we apply an empirical Bayes procedure to improve inference about firm-specific parameters by borrowing strength from the distribution of firm-level estimates, as in Walters (2024).

Beyond the linear shrinkage approach that assumes a normal mixing distribution, we use the nonparametric deconvolution algorithm of Efron (2016), implemented via the statistical package of Narasimhan and Efron (2020), with a natural cubic spline of degree 5, 200 support points, and a penalization parameter of 0.1.

To ensure compliance with data protection requirements, we use smoothed kernel density estimates rather than histograms when reporting the distributions of raw estimates, priors, and posteriors.[5] We implement standard adjustments including winsorizing extreme values and tuning kernel density bandwidth when visualizing the distributions. Specifically, we winsorize observations that fall more than four median absolute deviations from the median, which provides robust protection against extreme outliers while preserving the central distribution. For kernel density estimation, we use an Epanechnikov kernel with 200 evaluation points and scale the reference bandwidth by a factor ranging from 1.0 to 1.3 depending on the variable, ensuring sufficient smoothing to prevent disclosure of small cells. These adjustments ensure that individual firms cannot be identified in the published figures, while all underlying estimation (including deconvolution and posterior computation) is performed on the original firm-level data.

Figure 6 presents the distributional results of all parameters of interest of Equation (3). The different colors and patterns distinguish between the raw estimates, priors, and posteriors. The vertical line and shaded area indicate the estimates from Equation (2) and Table 3, pooling all firms. A quick glance at Figure 6 reveals that those pooled estimates mask sub-

[5] We report kernel density estimates because we are not allowed to report estimates for individual or small cells of firms.

stantial heterogeneity in the baseline birthday firing behavior (for native male workers) as well as the firing gaps of female and migrant workers.

The raw estimates exhibit substantial dispersion across all eight parameters of interest. The distribution for native males near their 2nd+ birthdays (panel b) shows a pronounced positive tail, indicating a mass of firms that are systematically more likely to dismiss this group near birthdays. In contrast, the distributions for native females and migrant females (panels d and h) show negative tails, indicating firms that are systematically less likely to dismiss these groups near birthdays than native males. However, these tails may be partly or entirely driven by sampling noise, which is precisely what the empirical Bayes methods are designed to filter out.

The priors obtained by nonparametric deconvolution are more concentrated than the raw estimates for all parameters, as expected when separating signal from sampling noise, although the difference is less pronounced in panel f. Most distributions have modes close to zero, while some display modes substantially below zero (panels f–h). The departures from normality suggest that the flexible nonparametric approach may better capture the true distribution of firm-level firing gaps than linear shrinkage, which assumes a normal mixing distribution.

The linear shrinkage posteriors, which assume a normal mixing distribution, shrink some of the raw estimates toward the mean, most notably in panel c, but also in panel e. The distributions of the posterior estimates are similar to the raw estimates in panels b and f–h, and are almost identical for the remaining parameters (panels a and d).

The nonparametric posteriors shrink the raw estimates more aggressively. The modes of the distributions in panels a, e, and g lie close to zero, while other posterior distributions exhibit pronounced positive modes (panels b–d) or negative modes (panels f and h). Several tails stand out, including negative tails in panels a, b, e, f, and h, and positive tails in panels b, c, and f–h, indicating substantial heterogeneity in firm-level firing gaps even after accounting for sampling noise. The posterior distribution for migrant males near their 2nd+ birthdays (panel g) is particularly noteworthy, as it exhibits trimodality, with two strongly negative modes and one slightly positive mode.

8 Conclusion

This paper studies whether firms fire workers from different demographic groups at different rates when retention costs rise mechanically at age-based minimum wage thresholds. Using

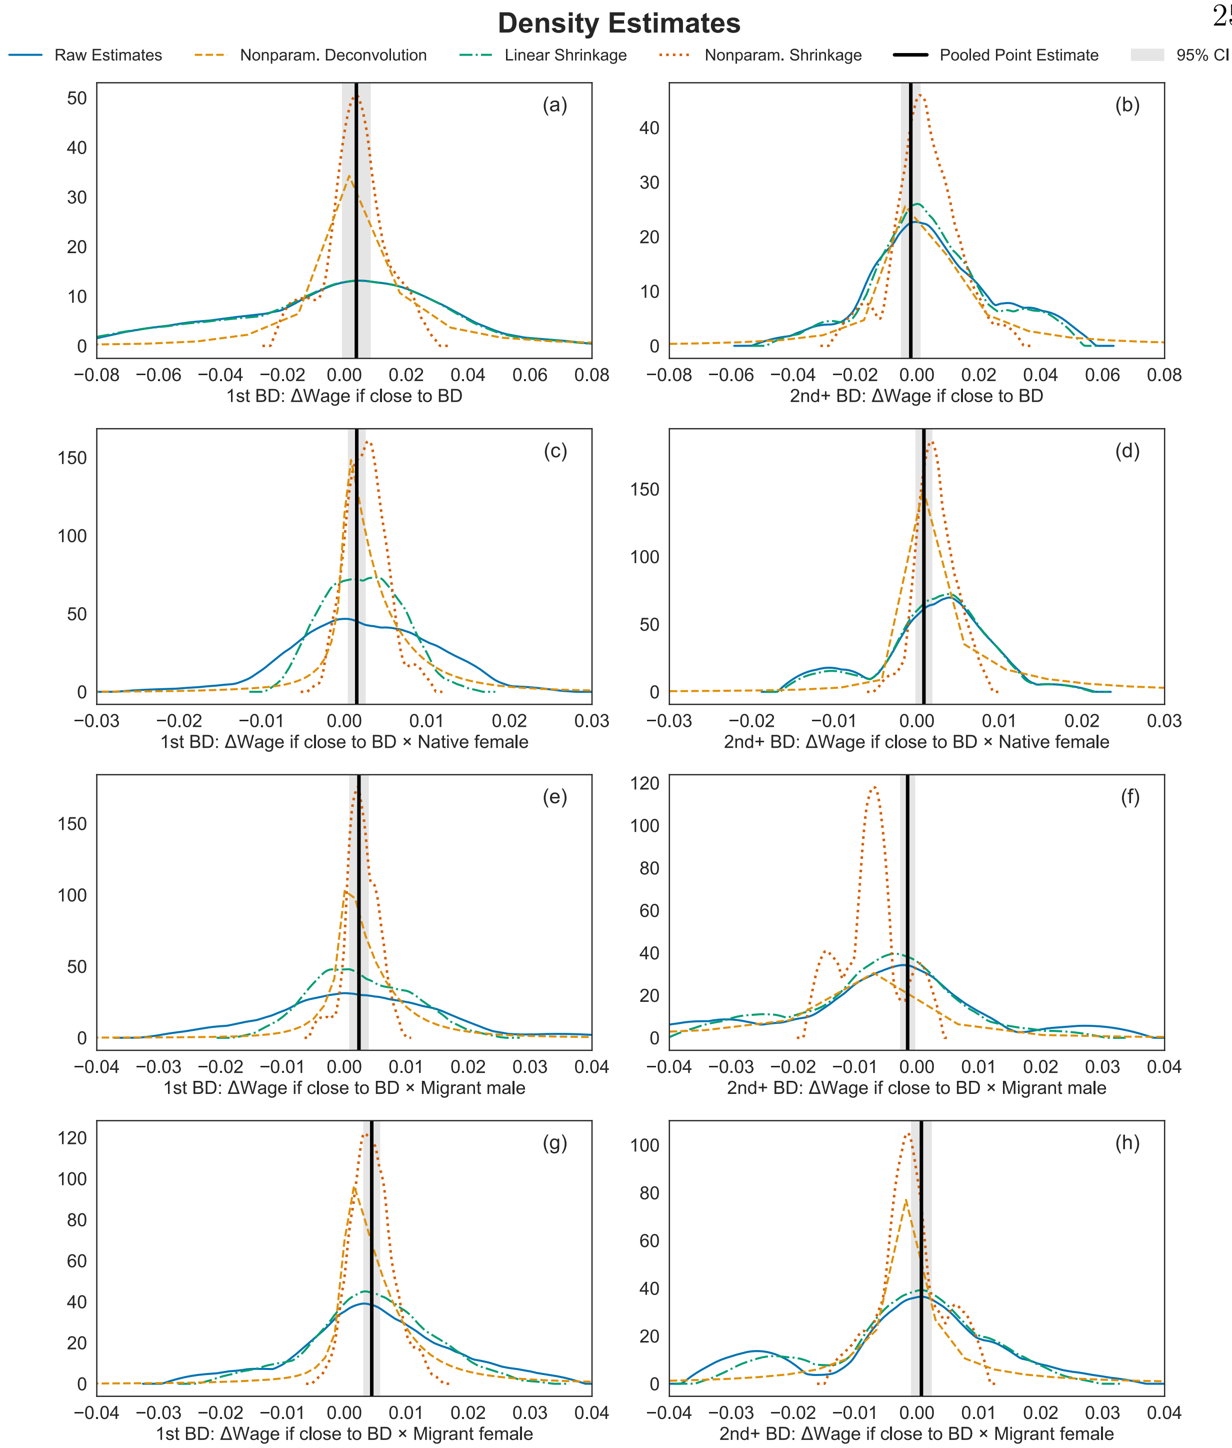

Figure 6: Density plots of firm-specific parameters by demographic group

Notes: Distributional estimates based on the 50 largest firms, with all remaining firms aggregated into a composite 51st firm. The raw estimates are obtained from Equation (3) and processed using an empirical Bayes procedure following Kline et al. (2022) and Walters (2024). We estimate the prior distribution via nonparametric deconvolution using the algorithm of Efron (2016), implemented through Narasimhan and Efron (2020). Kernel density estimates (with winsorization and bandwidth adjustments) are reported for disclosure control. The pooled estimates and their 95% confidence intervals are taken from Table 3.

Dutch administrative data and sharp statutory increases in youth minimum wages, we isolate firing responses that are exogenous to worker characteristics such as gender, migration background, and productivity. This allows us to examine differential treatment in the firing margin at scale, even though administrative data typically do not distinguish quits from involuntary separations.

Our results reveal systematic differences in firing patterns across demographic groups that vary markedly with tenure. At the first birthday, native females and migrant females face higher dismissal risks than native males, while migrant males show no differential treatment. At subsequent birthdays, after firms have observed performance for at least a year, migrant males become significantly less likely to be dismissed than comparable native males. Exploring heterogeneity by labor market conditions at the time of hiring shows that these patterns are concentrated among workers hired in loose youth labor markets, consistent with the idea that firms apply stricter hiring criteria when selecting migrant applicants. Among workers hired in tight markets, migrant males display no such negative firing gap at later birthdays and instead show a sizable positive firing gap at the first birthday. These findings support an interpretation in which differential hiring standards, rather than taste-based firing discrimination, drive much of the observed heterogeneity for migrant males.

Pooled estimates mask substantial variation in firm behavior. Using empirical Bayes methods, we recover firm-specific firing parameters and document wide dispersion in baseline birthday firing behavior and in firing gaps for female and migrant workers. The raw firm-level estimates exhibit pronounced positive and negative tails across multiple parameters, indicating that some firms are systematically more likely and others systematically less likely to dismiss certain demographic groups near birthday thresholds. The nonparametric priors obtained via deconvolution are more concentrated but remain distinctly non-normal, consistent with genuine underlying heterogeneity. The empirical Bayes posteriors continue to display substantial tails even after accounting for sampling noise. Particularly notable is the posterior distribution for migrant males at their second and later birthdays, which is trimodal, indicating distinct clusters of firms with sharply different firing patterns. These findings reveal meaningful firm-level heterogeneity in demographic firing gaps, beyond what sampling variation alone would generate.

This paper contributes to the literature on labor market discrimination by showing how firing responses to exogenous cost shocks can be used to probe mechanisms that are consistent with differential treatment operating through earlier stages of employment, in particular through differential hiring standards that shape the composition of the retained workforce.

While our observational data do not allow us to establish discrimination in the legal or causal sense, the patterns we document are informative about disparities in firm responses to identical cost increases across demographic groups. In addition, our approach provides a scalable, low-cost method for detecting firms with extreme firing behavior using naturally occurring statutory wage variation and empirical Bayes methods. This approach also has implications for regulators in countries with age-based minimum wage schedules, who may use such variation to detect systematic discriminators both across and within firms. Finally, we contribute to the evaluation of youth minimum wage policies in the Netherlands by showing that these policies influence not only employment levels but also the distribution of separations across demographic groups.

References

Alan Benson, Simon Board, and Moritz Meyer-ter Vehn. Discrimination in hiring: Evidence from retail sales. *The Review of Economic Studies*, 91(4):1956–1987, 2023. doi: 10.1093/restud/rdad087.

Marianne Bertrand and Sendhil Mullainathan. Are Emily and Greg More Employable Than Lakisha and Jamal? a Field Experiment on Labor Market Discrimination. *American Economic Review*, 94(4):991–1013, 2004. doi: 10.1257/0002828042002561.

Anna Bindler, Randi Hjalmarsson, Nadine Ketel, and Andreea Mitrut. Discontinuities in the age-victimisation profile and the determinants of victimisation. *The Economic Journal*, 134(657):95–134, 2023.

Centraal Bureau voor de Statistiek. Kenmerken van deelnemers aan de Eindtoets Basisonderwijs van Cito, 2005. URL https://doi.org/10.57934/0B01E410801D09AC.

Centraal Bureau voor de Statistiek. Banen en lonen op basis van de Polisadministratie, 2020. URL https://doi.org/10.57934/0B01E410804CB681.

Centraal Bureau voor de Statistiek. Persoonskenmerken van alle in de Gemeentelijke Basisadministratie Persoonsgegevens (GBA) ingeschreven personen, niet-gecoördineerd, 2022. URL https://doi.org/10.57934/0B01E4108071DAF2.

Rui Dai, Tingting Liu, and Cong Wang. Employee Turnover in the Spotlight: Evidence from Mergers and Acquisitions. *The Wharton School Research Paper*, 2024. Available at SSRN: https://ssrn.com/abstract=5280455.

Meltem Dayioglu, Müşerref Küçükbayrak, and Semih Tumen. The impact of age-specific minimum wages on youth employment and education: A regression discontinuity analysis. *International Journal of Manpower*, 43(6):1352–1377, 2022.

Bradley Efron. Empirical Bayes deconvolution estimates. *Biometrika*, 103(1):1–20, 02 2016. ISSN 0006-3444. doi: 10.1093/biomet/asv068. URL https://doi.org/10.1093/biomet/asv068.

Mark Egan, Gregor Matvos, and Amit Seru. When Harry Fired Sally: The Double Standard in Punishing Misconduct. *Journal of Political Economy*, 130(5):1184–1248, 2022.

Jan Kabátek, Ying Liang, and Kun Zheng. Are Shorter Cumulative Temporary Contracts Worse Stepping Stones? Evidence from a Quasi-Natural Experiment. *Labour Economics*, 84:102427, 2023. doi: 10.1016/j.labeco.2023.102427.

Jan Kabátek. Happy birthday, you're fired! Effects of an age-dependent minimum wage on youth employment flows in the Netherlands. *ILR Review*, 74(4):1008–1035, 2021.

Judd B. Kessler, Corinne Low, and Colin D. Sullivan. Incentivized resume rating: Eliciting employer preferences without deception. *American Economic Review*, 109(11):3713–44, November 2019. doi: 10.1257/aer.20181714. URL https://www.aeaweb.org/articles?id=10.1257/aer.20181714.

Patrick Kline, Evan K Rose, and Christopher R Walters. Systemic Discrimination Among Large U.S. Employers. *The Quarterly Journal of Economics*, 137(4):1963–2036, 06 2022.

Aaron Korczak Krzeczowski. Junior Minimum Wages and Employment Dynamics: An Analysis Using Australian Administrative Data. *mimeo*, 2023.

Claus Thustrup Kreiner, Daniel Reck, and Peer Ebbesen Skov. Do Lower Minimum Wages for Young Workers Raise Their Employment? Evidence from a Danish Discontinuity. *The Review of Economics and Statistics*, 102(2):339–354, 05 2020. ISSN 0034-6535. doi: 10.1162/rest_a_00825. URL https://doi.org/10.1162/rest_a_00825.

Balasubramanian Narasimhan and Bradley Efron. deconvolveR: A g-modeling program for deconvolution and empirical Bayes estimation. *Journal of Statistical Software*, 94(11): 1–20, 2020. doi: 10.18637/jss.v094.i11.

David Neumark, Roy J. Bank, and Kyle D. Van Nort. Sex Discrimination in Restaurant Hiring: An Audit Study. *The Quarterly Journal of Economics*, 111(3):915–941, 08 1996. ISSN 0033-5533. doi: 10.2307/2946676. URL https://doi.org/10.2307/2946676.

Overheid. Wet sociale werkvoorziening, 1997. URL https://wetten.overheid.nl/BWBR0008903/2015-07-01. Version as of July 1, 2015. Accessed: November 19, 2025.

Rijksoverheid. Bedragen minimumloon 2025, 2025. URL https://www.rijksoverheid.nl/onderwerpen/minimumloon/bedragen-minimumloon/bedragen-minimumloon-2025. Accessed: November 19, 2025.

K. Steenks, A. Heyma, and T. Vervliet. The Effects of Dutch Youth Minimum Wage Increases on Income Inequality. *De Economist*, 173:299–330, 2025. doi: 10.1007/s10645-025-09451-z.

E. Van Bezooijen, W. Van den Berge, and A. Salomons. The Young Bunch: Youth Minimum Wages and Labor Market Outcomes. *ILR Review*, 77(3):428–460, 2024. doi: 10.1177/00197939241239317.

Christopher Walters. Chapter 3 – Empirical Bayes methods in labor economics. In Christian Dustmann and Thomas Lemieux, editors, *Handbook of Labor Economics*, volume 5, pages 183–260. Elsevier, 2024. doi: 10.1016/bs.heslab.2024.11.001.

A Appendix

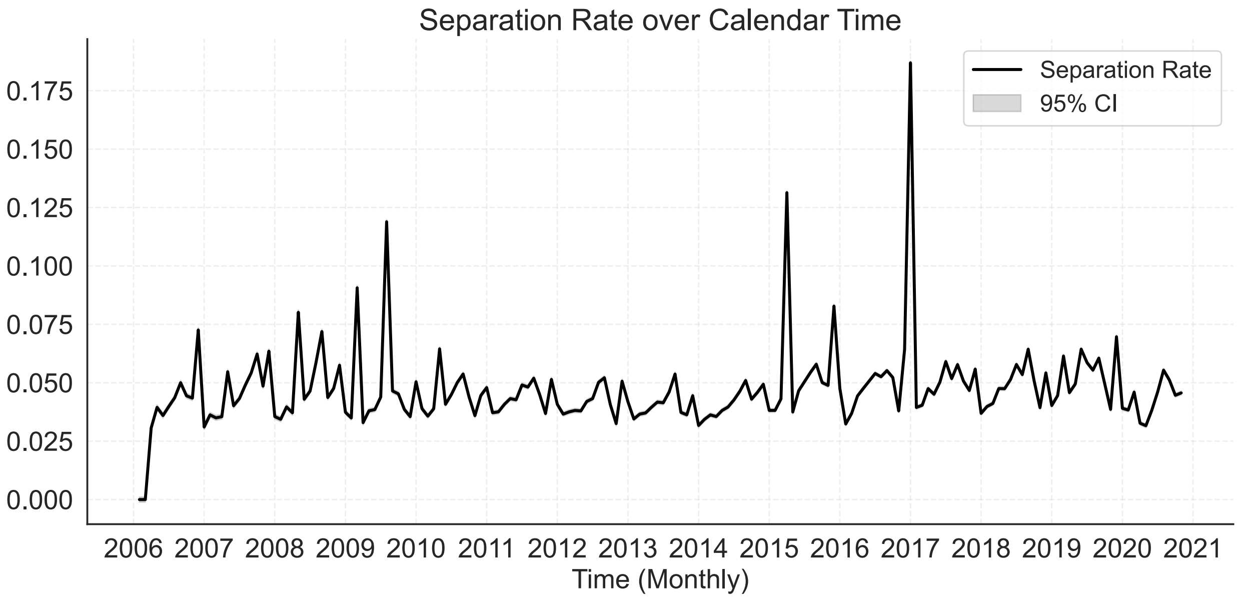

Figure A1: Separation Rate over Calendar Time

Notes: This figure plots the separation rate in the grocery sector in the Netherlands from 2006 through 2020. The rate fluctuates around 5% throughout the sample. Seasonal spikes, such as those in August and December, and the 2017 policy-induced spike are unrelated to our identification strategy, which exploits variation around workers' birthdays rather than calendar time.

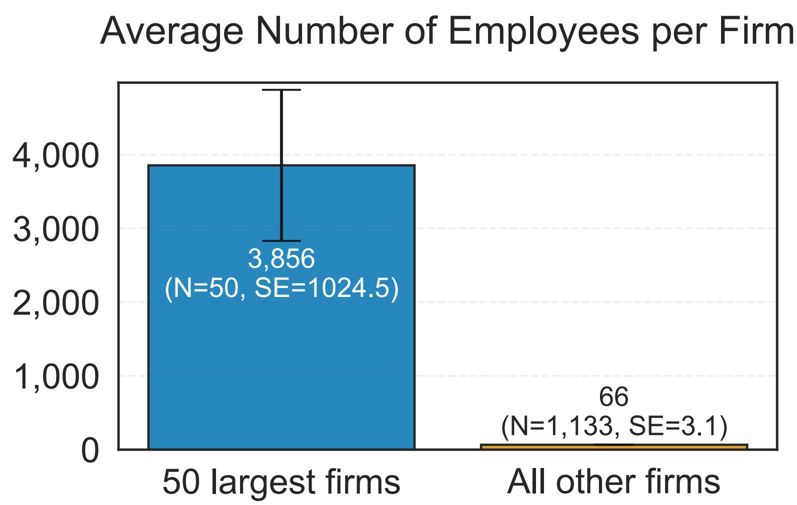

Figure A2: Average Firm Size: 50 Largest Firms vs. Smaller Firms

Notes: Error bars represent standard errors.

Table A1: Pooled Regression Results: Sample Split by Labor Market Tightness at Time of Hire

	Dep. var.: $\mathbb{1}\{\text{Separation}\} \times 100$			
	(1)	(2)	(3)	(4)
Close to birthday	0.820***	1.108***	0.422***	0.211
	(0.140)	(0.177)	(0.162)	(0.200)
Native female	-0.114**	-0.180***	-0.00642	-0.108***
	(0.0500)	(0.0351)	(0.0705)	(0.0402)
Migrant male	0.664***	0.497***	0.555***	0.415***
	(0.0484)	(0.0461)	(0.0631)	(0.0732)
Migrant female	0.149***	0.0721	0.0673	-0.0541
	(0.0484)	(0.0752)	(0.0630)	(0.104)
Close to birthday	-0.126*	-0.0246	0.0278	0.114*
×Native female	(0.0671)	(0.0374)	(0.0591)	(0.0652)
Close to birthday	0.0839	0.452***	-0.230***	-0.0157
×Migrant male	(0.0738)	(0.0641)	(0.0834)	(0.0739)
Close to birthday	-0.0408	0.129*	-0.0209	0.113
×Migrant female	(0.122)	(0.0756)	(0.0721)	(0.117)
R^2	0.178	0.184	0.100	0.105
Description	1st BD, loose	1st BD, tight	2nd+ BD, loose	2nd+ BD, tight
Controls	✓	✓	✓	✓
Firms	612	746	551	686
Workers	446,908	463,702	365,443	372,473
Observations	2,488,937	2,523,667	4,153,739	4,126,610

Notes: The dataset covers all firms and workers in the Dutch grocery sector from 2006–2020, based on administrative records from Statistics Netherlands. The dependent variable is a dummy equal to one if employment is terminated, multiplied by 100; coefficients are thus interpreted as percentage point changes. Regressors include gender, migration background, and their interactions with the wage change close to the birthday. The reference group is native Dutch men. The columns split the sample by whether workers were hired in a loose or tight labor market, based on a median split of labor participation at the year-municipality level. Control variables consist of age dummies, a manager dummy (pooling lower, middle, and upper managers), dummies for own ethnicity represented in lower, middle, and upper management, branch size quintiles (where branch refers to the firm-municipality combination), the CITO standard score (cognitive ability), the CITO teacher recommendation (non-cognitive ability), interactions of the ability and representation variables with the close-to-birthday indicator, and firm-month fixed effects. Standard errors clustered at the firm level in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.