NBER WORKING PAPER SERIES

EMPLOYMENT RELATIONSHIPS, WAGE SETTING, AND LABOR MARKET POWER

Francesco Agostinelli
Domenico Ferraro
Giuseppe Sorrenti
Leonard Treuren

Working Paper 34439
http://www.nber.org/papers/w34439

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

We thank Juan Pablo Atal, Isaac Baley, Hanming Fang, Chao Fu, Rasmus Lentz, Robert Moffitt, Luigi Pistaferri, B. Ravikumar, Edouard Schaal, Jeffrey Smith, Chris Taber, Nicolas Werquin, Matt Wiswall, and attendees at several seminars and conferences for helpful comments and suggestions. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2025 by Francesco Agostinelli, Domenico Ferraro, Giuseppe Sorrenti, and Leonard Treuren. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Employment Relationships, Wage Setting, and Labor Market Power
Francesco Agostinelli, Domenico Ferraro, Giuseppe Sorrenti, and Leonard Treuren
NBER Working Paper No. 34439
October 2025
JEL No. J0, J42, L10

ABSTRACT

We ask to what extent the quantification of labor market power depends on the modeling of the long-term worker-firm employment relationship. We develop an oligopsony model with dynamic wage contracts. Workers decide whether and where to work, choosing among firms providing different amenities and solving a dynamic discrete choice labor supply problem with firm-specific human capital. As a result, firms optimally choose wage-tenure contracts to attract and retain workers. We find that such contracts mitigate firms' incentives to impose large instantaneous wage markdowns—compared to standard static wage-setting models—thereby reducing the share of socially inefficient worker-firm separations. As a consequence, we show that the empirical approaches based on "sufficient statistics" tend to overestimate the extent of labor market power: low levels of firm-specific labor supply elasticities do not necessarily indicate rent extraction, but instead reflect firms' ability to retain workers by offering long-term value through human capital accumulation.

Francesco Agostinelli
University of Pennsylvania
and NBER
fagostin@upenn.edu

Domenico Ferraro
Arizona State University
Department of Economics
domenico.ferraro@asu.edu

Giuseppe Sorrenti
University of Lausanne
giuseppe.sorrenti@unil.ch

Leonard Treuren
KU Leuven
leonard.treuren@kuleuven.be

1. Introduction

A vast and growing body of empirical evidence indicates that employers possess labor market power—that is, they exercise wage-setting power that results in wages below competitive levels (e.g., Lamadon, Mogstad, and Setzler, 2022). Studying the mechanisms that allow employers to suppress wages and quantifying the effects of such practices on firms' economic rents and total welfare are key to designing active labor market policies such as minimum wages, tax benefits and exemptions, and employment protection legislation (EPL).

As labor market power cannot be directly observed in data, researchers need to take a stand on how to model labor markets to infer it. The norm in the literature is to use a *static* model of imperfectly competitive labor markets—static in the sense that firms set wages period-by-period to maximize current-period profit. For instance, in the textbook monopsony model, the firm picks a point on its upward-sloping labor supply curve, trading off employment against profit per worker (Manning, 2003). Higher wages coincide with higher employment and lower markdowns of the wage with respect to the marginal revenue product of labor—the "wage markdown"—and vice versa.

Static models abstract from dynamic considerations that might influence firms' wage-setting policies. In particular, firms might want to retain workers who accumulate firm-specific human capital over time.[1] Indeed, there is a long tradition of models with implicit or wage-tenure contracts, in which the current wage alone does not fully reflect the compensation for the exchange of labor services between an employer and the worker (Azariadis, 1975; Holmstrom, 1981, 1983; Thomas and Worrall, 1988; Burdett and Coles, 2003; Menzio and Shi, 2010; Lentz, 2014; Balke and Lamadon, 2022). In these settings, firms choose both the current wage and continuation utilities (promised values) over the workers' tenure as a device to attract and retain workers, ultimately aiming to maximize profits and to reduce inefficient employment terminations. These papers, however, abstract from the quantification of labor market power.

In this paper, we develop and estimate a novel oligopsony model that allows workers to make labor supply decisions with dynamic considerations related to human capital accumulation, and firms to offer dynamic wage contracts. Workers decide whether and where to work, choosing among firms that are heterogeneous in non-wage amenities and revenue productivity, and solving a dynamic discrete-choice labor-supply problem with both firm-specific and transferable human capital. Firms choose wage-tenure contracts to attract and retain workers. We estimate this model using a matched employer-employee panel dataset from the Netherlands covering the years 2006 to 2021. We estimate wage markdowns and compare them to estimates obtained

[1]Human capital accumulation has been shown to be quantitatively important for observed labor market outcomes and labor supply elasticities (Imai and Keane, 2004; Huggett, Ventura, and Yaron, 2011).

under the assumption that wage setting is static. On average, the wage markdown with respect to the marginal revenue product of labor implied by wage-tenure contracts is modest—about 10% over the tenure period on average—a figure that is orders of magnitude smaller than the average wage markdown in an estimated version of the model that replaces wage-tenure contracts with static (period-by-period) wage setting. This is because static wage setting tilts the firms' trade-off toward static profits per worker (a higher markdown) and away from firm size (attracting and retaining workers).

In addition, we show how heterogeneity across firms in amenities and revenue productivity is related to the dynamic wage contracts they offer and the dynamics of wage markdowns over the employment relationship. Empirically, we rely on a temporary tax cut for migrants that generates a predictable change in the relative price of compensating employees with wages rather than with amenities. We show that at high-amenity firms, wages and wage markdowns vary substantially with this relative price. Indeed, wages are marked up relative to the marginal revenue product of labor in the years leading up to the expiration of the tax cut and marked down thereafter. In contrast, wages trend with workers' productivity at low-amenity firms, as these firms have no other way to compensate employees than with wages.

We start by developing a model of oligopsonistic competition in the labor market in which firms use wage-tenure contracts to attract and retain workers. Non-wage amenities are a source of vertical job differentiation in the spirit of Rosen (1986). Workers accumulate both non-transferable and transferable skills, both of which improve revenue productivity. The firms' strategic variable is the promised value they commit to deliver to their employees via a sequence of wages conditional on an employee's tenure at the firm. Workers can switch employers or leave employment at any time if the value of doing so exceeds the value of staying in the current relationship, but this comes at a fixed switching cost. Firms optimally allocate wage expenses over today's wage and future workers' values, taking into account the one-sided commitment, the tax environment, and workers' preferences.

The model's first key insight is that the labor supply elasticities faced by firms vary considerably over a worker's tenure at a firm. This variation comes from the accumulation of firm-specific human capital, which naturally creates incentives for the firm to intertemporally substitute compensation so as to optimally provide incentives for the worker to remain in the employment relationship. The provision of incentives through wage-tenure contracts consists of a static component via the current wage and a dynamic component via future workers' continuation utility or promised value. Moreover, the more non-wage amenities a firm has, the more freely it can intertemporally shift worker compensation from wages to amenities or vice versa in response to changes in their relative price—e.g., due to anticipated tax changes. Such substitution forms a second driver of labor supply elasticity variation over workers' tenure.

A second key insight is that the relevant notion of a firm-specific labor supply elasticity is not the static one, typically used to identify wage markdowns in static models of imperfectly competitive labor markets. Instead, it is the labor supply elasticity that a firm anticipates facing over the tenure of the worker at the firm given the accumulation of firm-specific human capital, as well as the extent of competition from other firms. Neglecting these dynamic considerations leads to biased estimates of wage markdowns. As stated, firms actively influence this long-term labor supply elasticity by using wage-tenure contracts to retain workers, thereby reducing this elasticity, and by optimally shifting wages across tenure to accommodate changing relative prices of wage-based and amenity-based compensation. As a result, the static firm-specific labor supply elasticity is no longer a sufficient statistic for the wage markdown. In fact, in the presence of dynamic wage setting, a low elasticity faced by firms on a period-by-period basis is generally consistent with a low, not a high, wage markdown, as it indicates successful retention of workers.[2]

To estimate the model, we assemble a matched employer–employee panel covering employees of Dutch firms from 2006 to 2021, obtained from Statistics Netherlands. This dataset allows us to observe workers' individual wages and paid taxes, and track them as they move between different firms and accumulate tenure and experience. While this data would suffice to estimate the model, we additionally obtain data on workers who use the Migrants' Tax Exemption Policy (MTEP) from the Ministry of Finance. The MTEP temporarily grants eligible high-skilled migrants to the Netherlands a partial income-tax exemption—up to 30% of salary tax-free. The MTEP amounts to a fully anticipated and dramatic change in the cost of compensating employees with wages over workers' tenure. Therefore, it is the ideal setting to illustrate how firm heterogeneity determines the optimal wage contract, and how this affects wage markdowns in a particular year and over time.

Using the matched employer-employee dataset, we estimate our model's parameters and quantify markdowns. We allow firms to be heterogeneous in terms of the revenue-generating technologies they use and the non-wage amenities they provide to their employees. More revenue-productive and higher-amenity firms have a competitive advantage in attracting and retaining workers. Such an advantage is further amplified by the wage-tenure contracts as such firms are better equipped to retain workers, implying a longer average duration of employment spells, and, thereby, higher average productivity due to tenure effects. Both the exogenous and the endogenous components of firm heterogeneity allow the model to reproduce the observed variation in wages over workers' tenure and across firms.

[2]Similarly, Erdem, Imai, and Keane (2003) and Hendel and Nevo (2006) show that relying solely on short-term price variation can lead to misleading estimates of price elasticities and markups in product markets where dynamics are important—highlighting the importance of accounting for dynamic demand considerations, which is central to our analysis.

The model also replicates the observed decline in wages at the expiration date of tax benefits for foreign workers in the Netherlands. Indeed, the heterogeneity across firms in this observed decline is fully rationalized by the underlying heterogeneity in amenity values that the firms can provide to their employees. Wages are taxed, amenities are not. So, even if anticipated, the change in the tax rate faced by workers alters the way firms optimally structure wage compensation over the workers' tenure to induce as much retention as possible given the extent of prevailing competition from other firms. This mechanism applies to all firms, but the extent to which a given firm can leverage it depends on the level of amenities it provides.

Specifically, high-amenity firms set their wages to steeply increase for a few years before the expiration of tax benefits, decline at the time of expiration, and rise again thereafter. This pattern of wages around the tax benefit expiration date naturally implies intertemporal variation in wage markdowns in continuing employment relationships. In contrast, lower-amenity firms have much less room to shift compensation over the tenure of the worker at the firm, as declines in the monetary component of worker compensation would lead to reduced retention with the associated loss in valuable firm-specific human capital. The observed sawtooth patterns of wages around the expiration date of the tax benefits and the heterogeneity across firms indicate that firms can indeed structure compensation in a way that is consistent with the existence of long-term employment relationships.

The use of wage-tenure contracts shifts the trade-off faced by firms away from static profits per worker, and thus high wage markdowns, toward firm size through attraction and retention of workers. Although the specific numbers we estimate for the average wage markdown may and will likely change depending on the country and context studied, the far-reaching implications of our analysis remain. Dynamic wage setting fundamentally changes the way we think about how firms exercise wage-setting power, bringing the intertemporal aspects of employment relationships to the forefront of markdown estimation.

This paper is organized as follows. Section 2 discusses how this paper relates to the literature. Section 3 introduces the key insight of our paper using the canonical textbook monopsony framework. Section 4 presents the full oligopsony model with wage-tenure contracts. Section 5 describes the matched employer-employee data and the tax exemption policy that we use to estimate the model. In Section 6, we discuss how wage-tenure contracts help reconcile the model with the tenure profile of wages in the data and we develop insight into the identification of key parameters that govern heterogeneity in technology and amenities. We leverage these insights in Section 7, where we estimate the model using cross-sectional and longitudinal moments from our dataset. Section 8 quantifies the main implications of the model for wage markdowns, underscoring how static wage setting can generate misleading conclusions about the extent of labor market power exercised by firms. Section 9 concludes.

2. Related Literature

Our paper is related to several strands of literature. First, we relate to articles that examine labor market power using oligopsony models.[3] Historically, such work has built on the insights of Robinson (1933) and emphasized strategic interaction among large firms offering undifferentiated jobs. Recently, scholars have followed the approach of Card et al. (2018) to integrate insights from modern industrial organization into the study of labor market power by allowing idiosyncratic preferences over jobs to generate (additional) wage-setting power (e.g., Azar, Berry, and Marinescu (2022)). Our work extends this approach by incorporating dynamic considerations, which are arguably more prominent in labor markets than in product markets.[4] Perhaps closest to our paper in this line of research is Lamadon, Mogstad, and Setzler (2022), who quantify compensating differentials and rents using a *static* oligopsony model. We build on their work in two main ways: firms post wage-tenure contracts to maximize the discounted stream of future profits, and workers choose whether and where to work, taking into account the accumulation of transferable and non-transferable human capital within employment relationships.[5]

Recent developments in the literature on dynamic models of labor market power without search frictions include Han (2024), who analyzes a special case of wage posting in which the firm's policy is log-linear in worker experience, and Jungerman (2023), who considers dynamic profit maximization with wage posting in the spirit of the textbook static monopsony model, but without introducing dynamic contracts. In contrast, our framework highlights that once firms can offer dynamic contracts, they gain the ability to nonlinearly shift rents across the employment relationship to retain workers, which opens up new insights into wage setting and implications for markdown estimation relative to linear wage-posting approaches.

Second, there is a strand of literature on labor market power that considers search frictions. Jarosch, Nimczik, and Sorkin (2024) develop a search-theoretic model of the labor market that incorporates bargaining, in which a finite number of firms compete for workers by influencing workers' outside options. Bagga (2023) shows that such a framework goes a long way toward explaining the observed decline in employer-to-employer transitions and wages relative to productivity in the United States. Gottfries and Jarosch (2023) study the impact of noncompete

[3] Azar and Marinescu (2024) and Kline (2024) review the literature on wage setting power.

[4] We replace the typical assumption of unilateral static profit maximization with unilateral dynamic profit maximization. Delabastita and Rubens (2025) study the quantification of wage markdowns when employers coordinate their wages in a static setting using production and cost data. Such data are also employed in other studies to quantify markdowns under static cost-minimization assumptions (e.g., Yeh, Macaluso, and Hershbein (2022)).

[5] Volpe (2024) builds on Lamadon, Mogstad, and Setzler (2022) by allowing workers to have heterogeneous marginal rates of substitution between wages and amenities. This extension introduces a direct link between the distribution of workers' preferences in the economy and the shape of firm-specific labor supply curves.

agreements on wages and employment with a dynamic model that features on-the-job search, wage posting, and a finite number of firms, and Gottfries and Jarosch (2025) extend this model to employer wage coordination. As Card (2022) stresses, search frictions directly model job turnover at the cost of assuming that workers are imperfectly informed about job opportunities, while the static idiosyncratic preference approach does not rely on imperfect information but assumes that workers have their most-preferred job. Our dynamic model allows for both informed workers and job turnover, and emphasizes the role of accumulated human capital in restricting job mobility as workers' idiosyncratic preferences over jobs change over time.[6]

In macroeconomics, scholars have evaluated the welfare effects of labor market power in general-equilibrium models with oligopsonistic labor markets. Berger, Herkenhoff, and Mongey (2022) demonstrate how firm labor market power distorts wages and employment, leading to substantial welfare losses and output reductions linked to increasing labor market concentration. Like us, they caution against the use of reduced-form estimates of labor supply elasticities to infer wage markdowns and wage-setting power. Building on this foundation, Berger, Herkenhoff, and Mongey (2025) show that firm productivity dispersion limits the ability of minimum wages to overcome the welfare distortions of labor market power.[7] Deb et al. (2024) integrate product market power into this approach to study how labor and product market power, as well as firm heterogeneity in technology, influence wage inequality and premia. Azar and Vives (2021) consider labor and product market power in general equilibrium in the context of rising common ownership. In contrast to these studies, our analysis focuses on dynamic worker–firm relationships and abstracts from general-equilibrium effects. We demonstrate that explicitly incorporating long-term contracts fundamentally changes the view of labor market power, with important implications for the quantification and empirical measurement of wage markdowns. Long-term contracts enable firms to flexibly schedule compensation in a manner that tilts the trade-off away from per-worker profitability—and thus higher markdowns—toward worker retention and larger firm size achieved through lower markdowns.

One implication of our work is that the short-run labor supply elasticity is not a sufficient statistic for labor market power or the wage markdown. Reduced-form estimates of such elasticities tend to produce high markdowns—see, for example, several papers discussed in Ashenfelter, Farber, and Ransom (2010)—but they abstract from dynamic considerations that induce firms to increase wages. Several authors have pointed out that labor supply elasticities need not be a sufficient statistic for labor market power, even in a static setting. In addition to labor market power, wedges between the marginal revenue product of labor and wages might reflect,

[6]Several authors consider dynamic models in the search tradition without a focus on labor market power. For instance, dynamic bilateral competition featuring wage bargaining (e.g., Cahuc, Postel-Vinay, and Robin (2006); Bagger et al. (2014); Bagger and Lentz (2019)), or dynamic contracting (Balke and Lamadon, 2022).

[7]Berger et al. (2024) expand on the mechanisms underlying wage markdowns by incorporating worker–firm-specific preferences, firm granularity, and job search frictions.

for instance, output market power (Kroft et al., 2025), market power in other input markets (Treuren, 2022), or factor-biased technological change (Rubens, Wu, and Xu, 2025). These papers do not study dynamic considerations, which is our key contribution.

3. The Main Idea in a Nutshell

This section uses a toy model to illustrate how dynamic considerations in wage-setting affect the standard monopsony framework and imply that the short-run labor-supply elasticity is no longer a sufficient statistic for the wage markdown. Indeed, if firms use their wage policy to attract *and* retain workers, the relevant notion of labor-supply elasticity is a long-run elasticity over the worker's tenure at the firm. For expositional simplicity, we stay as close as possible to the seminal textbook monopsony model of Manning (2003). In Section 4, we develop the full oligopsony model that introduces specific dynamic labor-supply considerations—for example, workers accumulate firm-specific human capital and face switching costs when changing employers—and incorporates idiosyncratic preferences across jobs. The intuition developed in this section carries over to this richer setting.

The Textbook Monopsony Model As in Manning (2003), the monopsonist chooses the wage to maximize static profits, facing an imperfectly elastic firm-specific labor supply curve $D(w)$, with $\epsilon_w^D(w) \equiv (D'(w)/D(w))\, w$ and $0 < \epsilon_w^D(w) < \infty$. In choosing the wage w, the firm faces a trade-off between the profit per worker and the size of the firm

$$\max_{w \geq 0} \Pi \equiv \overbrace{(A - w)}^{\text{profit per worker}} \cdot \underbrace{D(w)}_{\text{mass of workers}}, \tag{1}$$

where A is the marginal revenue product of labor. Such a specification captures the extensive margin of labor supply where the wage determines the probability (or fraction) of workers the firm attracts.

The first-order condition (FOC) for w implies the following inverse relationship between the wage markdown and the firm-specific labor supply elasticity

$$\underbrace{\left(\frac{A}{w} - 1\right)}_{\text{wage markdown}} = \frac{1}{\epsilon_w^D(w)}. \tag{2}$$

Given an estimate of $\epsilon_w^D(w)$, the implied wage markdown follows directly from (2); in this

sense, $\epsilon_w^D(w)$ is a sufficient statistic for the markdown. The firm's problem in (1) is static: the monopsonist sets the wage period-by-period because the firm-specific labor supply depends only on the current wage. Consequently, the FOC in (2) holds every period, and the current wage choice has no bearing on the firm's future size.

A Dynamic Monopsony Model We now incorporate dynamic considerations by allowing the firm-specific labor supply curve to evolve with tenure, i.e., with the years an employee spends at the firm. As a result, the size of the firm's workforce depends on the entire sequence of wages offered over the tenure profile. While this feature could be given several different microfoundations, we deliberately keep the framework general.

For expositional simplicity, we consider two periods, but the insights straightforwardly apply to multi-period settings more generally. Consider a single cohort of workers, each deciding whether to work for the monopsonist this period and, conditional on working this period, whether to remain in the next period. Staying one more year makes a worker more productive. The monopsonist chooses the wage in the two periods, w_1 and w_2, to maximize the sum of period profits

$$\max_{w_1,w_2\geq 0} \Pi \equiv (A_1 - w_1) D_1(w_1, w_2) + (A_2 - w_2) D_2(w_1, w_2), \tag{3}$$

where both $D_1(w_1, w_2)$ and $D_2(w_1, w_2)$ depend on the wages in both periods, and $A_2 > A_1$ captures the increase in productivity with tenure. The dependence of the first period's labor supply $D_1(w_1, w_2)$ on the future wage w_2 captures the idea that for any given w_1 in the first period, if the monopsonist offers a higher future wage w_2, it attracts more workers in the first period. The dependence of labor supply in the second period $D_2(w_1, w_2)$ on w_1 captures the idea that the number of workers the monopsonist will employ in the second period equals the number of workers it attracts in the first period, which depends on w_1, times the fraction of them who choose to work in the second period. Unlike the textbook monopsony model, firm size in a given period depends on wages in all periods. The cost of increasing the wage in a given period is that the margin per worker decreases—note that the productivity increase captured by $A_2 > A_1$ puts upward pressure on this margin over time. The FOCs for w_1 and w_2 capture this intertemporal trade-off due to dynamic wage setting

$$-D_1(w_1, w_2) + (A_1 - w_1)\frac{\partial D_1(w_1, w_2)}{\partial w_1} + (A_2 - w_2)\frac{\partial D_2(w_1, w_2)}{\partial w_1} = 0, \tag{4}$$

$$-D_2(w_1, w_2) + (A_1 - w_1)\frac{\partial D_1(w_1, w_2)}{\partial w_2} + (A_2 - w_2)\frac{\partial D_2(w_1, w_2)}{\partial w_2} = 0. \tag{5}$$

These FOCs show that if one ignores the dependence of $D_2(w_1, w_2)$ on w_1 and $D_1(w_1, w_2)$ on

w_2—that is, assumes $\frac{\partial D_2(w_1,w_2)}{\partial w_1} = 0$ and $\frac{\partial D_1(w_1,w_2)}{\partial w_2} = 0$—the static monopsony condition in equation (2) holds in every period. In other words, once dynamic considerations are introduced on the labor supply side, they alter the firm's wage-setting incentives.

Rearranging terms yields an expression for the wage markdown that depends on the inverse of the contemporaneous labor-supply elasticity—as in the static model—and additionally on terms that capture dynamic wage setting, as the following proposition summarizes.

Proposition 1 (**Static vs. Dynamic Wage Setting**). *Consider the two-period monopsony model with dynamic wage setting. The wage markdowns in the two periods satisfy the following system of two equations*

$$\left(\frac{A_1}{w_1} - 1\right) = \frac{1}{\epsilon^{D_1}_{w_1}(w_1, w_2)} - \left(\frac{A_2}{w_2} - 1\right) \cdot \frac{w_2}{w_1} \cdot \frac{\epsilon^{D_2}_{w_1}(w_1, w_2)}{\epsilon^{D_1}_{w_1}(w_1, w_2)} \cdot \frac{D_2(w_1, w_2)}{D_1(w_1, w_2)}, \tag{6}$$

$$\left(\frac{A_2}{w_2} - 1\right) = \frac{1}{\epsilon^{D_2}_{w_2}(w_1, w_2)} - \left(\frac{A_1}{w_1} - 1\right) \cdot \frac{w_1}{w_2} \cdot \frac{\epsilon^{D_1}_{w_2}(w_1, w_2)}{\epsilon^{D_2}_{w_2}(w_1, w_2)} \cdot \frac{D_1(w_1, w_2)}{D_2(w_1, w_2)}. \tag{7}$$

Evidently, $\epsilon^{D_1}_{w_1}(w_1, w_2)$ and $\epsilon^{D_2}_{w_2}(w_1, w_2)$—the elasticities of labor supply in a period with respect to that period's wage—are not a sufficient statistic for the wage markdown in that period. Hence, neglecting the additional terms on the right-hand side of (6)-(7) would lead to biased estimates of wage markdowns.

Generally, this bias has four components:[8]

- *The wage markdown(s) in the other period(s)*
- *The wage in the other period(s) relative to the current period's wage*
- *The elasticity of labor supply in the other period(s) with respect to the contemporaneous wage, relative to the current period's labor-supply elasticity*
- *The labor supply in the other period(s) relative to the current period's labor supply*

Note that Proposition 1 captures an important trade-off that firms face when determining wage markdowns. In the second period, workers have accumulated one year of tenure, which

[8]Note that the same logic applies to a multi-period monopsony model with dynamic wage setting. Let $w \equiv (w_1, w_2, \ldots, w_T)$ denote the sequence of wages (contract) offered by the monopsonist. The multi-period expression for the wage markdown in the first period is a forward-looking intertemporal condition

$$\left(\frac{A_1}{w_1} - 1\right) = \frac{1}{\epsilon^{D_1}_{w_1}(w)} - \sum_{j=2}^{T} \left[\left(\frac{A_j}{w_j} - 1\right) \cdot \frac{w_j}{w_1} \cdot \frac{\epsilon^{D_j}_{w_1}(w)}{\epsilon^{D_1}_{w_1}(w)} \cdot \frac{D_j(w)}{D_1(w)}\right]. \tag{8}$$

has increased their productivity. This puts upward pressure on margins and incentivizes firms to employ more workers in the second period. In the first period, then, firms face a trade-off between marking down wages to maximize static profits, or marking down wages less and increasing employment and profits in the more-productive second period. Likewise, a lower second-period wage increases the markdown there, but reduces firm size in the first (and hence second) period.

Proposition 1 also establishes that if firms take into account dynamic considerations when setting wages, then interpreting the inverse of the short-run labor-supply elasticity as the markdown leads to bias. The magnitude of this bias depends on the characteristics of the labor market in question. For instance, the bias is likely to be less relevant in markets that operate more like spot markets—such as seasonal hiring in industries like tourism, agriculture, or hospitality—where employment relationships are short-lived and dynamic considerations play only a minimal role. By contrast, in settings where long-term employment relationships are salient and firm-specific human capital accumulates, dynamic forces are much more important, and the bias can become quantitatively significant.

In the next section, we extend the analysis to an oligopsony environment where firms offer wage-amenity bundles, adjust wages across periods, and where workers accumulate transferable and non-transferable human capital over time. We also allow workers to derive utility from firm amenities, so that firms with more attractive amenities have a comparative advantage in retaining employees throughout the employment relationship, and we allow for idiosyncratic preferences across jobs. The central insight developed here—that the labor-supply elasticity in a given period is not a sufficient statistic for the wage markdown—continues to hold in this richer framework, even though the quantitative implications may differ substantially

4. Model

This section develops a dynamic oligopsony model that we use to analyze wage markdowns. Workers face a dynamic discrete-choice labor-supply problem in which they choose whether and where to work among firms offering different non-wage amenities. Labor supply is dynamic because the decision of where to work depends on firm-specific human capital that accumulates with job tenure. In addition, we assume that firms provide wage-tenure contracts to attract and retain employees, fostering long-term employment relationships.

We derive closed-form expressions for workers' conditional choice probabilities and continuation values by assuming extreme-value Type I idiosyncratic shocks. Leveraging these results, we develop a tractable computational algorithm to solve the equilibrium dynamic wage-contracting

problem, in which heterogeneous firms compete for workers through wage-tenure contracts.

4.1. Dynamic Labor Supply

We model the labor supply of a cohort of workers entering the labor market. Let $t = 0, 1, \ldots, T$ denote the number of periods elapsed since entry. The cohort starts with a measure-one continuum of ex-ante identical workers indexed by $i \in [0, 1]$ and a finite number of firms indexed by $j = 1, 2, \ldots, J$, with $J \geq 2$. Firms are differentiated by amenities, which capture a host of job-related benefits that provide utility to the worker beyond monetary compensation, and by revenue productivity.

A worker's revenue productivity evolves over time for two reasons. First, workers accumulate non-transferable and firm-specific human capital. We capture this via a worker's tenure e_j. Tenure in firm j starts at $e_j = 0$, increases to $e_j + 1$ if the worker stays with the same firm in the next period, and resets to $e_{j'} = 0$ if the worker switches to another firm $j' \neq j$ or chooses not to work. Thus, a worker's maximum possible tenure is $T - 1$. Second, workers also accumulate transferable human capital over time, which we capture with worker age.[9]

As in Lamadon, Mogstad, and Setzler (2022), we specify that the utility of a worker at time t, conditional on working at firm j, with tenure e_j, depends on after-tax wages, $(1 - \tau_t)W_{j,t}(e_j)$, where $0 < \tau_t < 1$ is the proportional tax rate on labor income, and firm-specific non-monetary amenities G_j, as follows

$$\begin{aligned} \tilde{U}_{i,j,t}(e_j; W_{j,t}(\cdot), \epsilon_{i,j,t}) &= U_{j,t}(e_j; W_{j,t}(\cdot)) + \epsilon_{i,j,t}, \\ &= \lambda_1 \left((1 - \tau_t)W_{j,t}(e_j)\right)^{\lambda_2} + \alpha \ln G_j + \epsilon_{i,j,t}. \end{aligned} \tag{9}$$

Here, the parameter λ_1 governs the extent to which after-tax wages influence where workers choose to work, conditional on the realizations of the extreme-value Type I idiosyncratic shocks $\epsilon_{i,j,t}$ and firms' amenities G_j.[10] The lower λ_1, the less wages affect workers' sorting across firms. Similarly, α governs the extent to which vertical differentiation by amenities influences where workers choose to work. Finally, λ_2 governs the extent of the compensating differential between

[9] If tenure were transferable, workers would become more attractive to other firms, lowering the cost of switching and thereby reducing wage markdowns.

[10] We simplify our analysis by not modeling unobserved heterogeneity in workers' preferences—an important factor that may affect how workers sort across firms (see Volpe, 2024). Our goal, however, is to show to what extent alternative wage-setting protocols—specifically, dynamic contracts versus static wage determination—affect the quantification of markdowns. From this perspective, it is not obvious that abstracting from preference heterogeneity is a first-order concern, since the main differences between dynamic contracts and static wages arise from firms' distinct incentives to extract rents over time—whether immediately or through the continuation value of the employment relationship. That being said, we view building on our model by incorporating preference heterogeneity as a valuable extension for future work.

monetary and non-monetary compensation. For a linear utility function ($\lambda_2 = 1$), the increase in amenities required to compensate a worker for an additional euro in wages is independent of the wage level. However, if the utility function is nonlinear, the compensating differential depends on the wage level.

Later, when we apply the model to the data, we expand workers' choice set by incorporating an outside option representing nonparticipation. This additional choice can be quantitatively significant, as firms compete not only with each other for workers but also with the option of nonparticipation or market exit.[11]

Workers' Dynamic Problem Let $\mathbf{W}_t \equiv \left\{ \{W_{j,\tau}(\cdot)\}_{j=1}^{J} \right\}_{\tau=t}^{T}$ denote the wage sequences offered by all firms to workers at time t. Workers solve a dynamic labor supply problem by choosing the best alternative among firms, taking the wages as given. A worker's value of working at firm j with tenure $0 \leq e_j < t$ is

$$v_{j,t}(e_j; \mathbf{W}_t) = U_{j,t} - \psi \cdot \mathbb{1}(e_j = 0) + \beta EV_{t+1}(e_j + 1; \mathbf{W}_{t+1}), \tag{10}$$

where ψ is a switching cost that is incurred in the first year that you work at a particular firm (i.e., when $e_j = 0$). As is evident from (10), workers are aware of the dynamic impact of their current choices on the continuation value $EV_{t+1}(\cdot)$, recognizing that staying with their current firm will increase their tenure to $e_j + 1$, which in turn affects their wage.

Given that tenure resets when a worker changes firm, the vector of tenures across all firms, $\overline{e}$, is a vector with a positive entry e_j for the current employer j and zeros for all other firms $j' \neq j$, $\overline{e} = [0, \ldots, e_j, \ldots, 0]$. Thus, the current continuation value $EV_t(\cdot)$ is calculated by integrating over the extreme value Type I random utility shocks across all options, $\epsilon_{i,j,t}$, that are i.i.d. across workers, firms, and time

$$EV_t(\overline{e}; \mathbf{W}_t) = \int \max_j \{v_{j,t}(e_j; \mathbf{W}_t) + \epsilon_{i,j,t}\} \, dF(\epsilon_{i,j,t}) \tag{11}$$

$$= \log\left(\sum_k exp(v_{k,t}(e_k; \mathbf{W}_t)) \right) + \gamma \,, \tag{12}$$

where γ is Euler's constant. The fraction of workers working in firm j for each vector of tenure profiles is

$$S_{j,t}(\overline{e}; \mathbf{W}_t) = \frac{exp(v_{j,t}(e_j; \mathbf{W}_t))}{\sum_k exp(v_{k,t}(e_k; \mathbf{W}_t))} \equiv S\big(\overline{e}, v_{j,t}(e_j; \mathbf{W}_t), \mathbf{v}_{-j,t}(\mathbf{e}_{-j}; \mathbf{W}_t)\big) \,, \tag{13}$$

[11] A similar idea is explored in Kahn and Tracy (2023), which examines how high-skilled workers can bypass the wage penalties of local monopsonistic labor markets by relocating to more competitive areas.

which depends on the value of working for the firm j ($v_{j,t}(\cdot)$), as well as on the vector of values of starting to work for one of the other firms ($\mathbf{v}_{-j,t}(\cdot)$).

4.2. Dynamic Labor Demand and Wage-Tenure Contracts

On the labor demand side, firms are assumed to generate revenue with a linear technology. The revenue production function of a single worker at time t in firm j with tenure e_j is given by

$$y_{e,j,t} = A_{j,t} \cdot (1 + e_j)^{\gamma_1}, \tag{14}$$

where $A_{j,t} \equiv A_j t^{\gamma_2}$ is the marginal revenue product of labor of firm j at time t. Parameters γ_1 and γ_2 capture, respectively, how firm-specific tenure and time since labor market entry affect worker productivity.

Firms' Dynamic Problem Firms solve the problem of offering contracts to workers they do not currently employ—i.e., for whom $e_j = 0$.[12] Once new workers are hired at firm j, the evolution of tenure e evolves in terms of the periods of continuous employment within that firm. Given the linear technology, firm j's expected profits of hiring D workers (with zero tenure) at time t is

$$\begin{aligned}\Omega_{j,t,0}(D, \mathbf{W}_j; \mathbf{W}_{-j}) = &(A_{j,t}D - (1 + \tau_t^p)W_{j,t}(0)D) + \\ &\sum_{e=1}^{T-t} \underbrace{P_{j,t+e}(e, \mathbf{W}_j; \mathbf{W}_{-j})}_{\text{Retention Probability}} \times \Big[A_{j,t+e} \cdot (1 + e)^{\gamma_1} \cdot D - (1 + \tau_{t+e}^p)W_{j,t+e}(e) \cdot D\Big],\end{aligned}$$

where $P_{j,t+e}(e, \mathbf{W}_j; \mathbf{W}_{-j})$ represents the probability of retaining a worker, and is given by the solution of the dynamic labor supply model given the sequence of workers' future wages. We define as τ_t^p the payroll tax that firms are required to pay to the government. We allow this tax to vary over time to capture institutional features specific to our empirical setting (the Netherlands). The precise details of these fiscal rules are provided in Section 6.

To solve for the optimal contract, the firm maximizes expected profits subject to delivering at least a given value to the worker. The dynamic problem corresponding to the optimal contract can be expressed in recursive form, with the promised value to the worker serving as the state variable and the continuation value as the control variable. Intuitively, the promised value V summarizes the history of the employment relationship up to the current period, while the continuation value V' becomes the state variable in the next period. Similarly, V_{-j} and V'_{-j}

[12] Although workers can be hired from other firms, the relevant labor supply curve faced by firm j consists only of workers with zero tenure at that firm ($e_j = 0$).

denote the promised values and the continuation values for all other remaining competing firms.

In this recursive problem, firms optimally allocate resources between today wage and future workers' value, internalizing the lack of commitment from the workers' side. We assume that firms commit to the contract, which fits the Dutch context with strong employment protection laws well. The entire recursive problem can then be characterized by the per-worker dynamic problem of the firm as follows

$$\pi_{j,t,e}(V_{j,t,e};\mathbf{V}_{-j}) = \max_{w,\,V'}\Big\{A_{j,t}(1+e)^{\gamma_1} - w(1+\tau_t^p) + \rho_{j,t+1,e+1}(V';\mathbf{V}'_{-j})\pi_{j,t+1,e+1}(V';\mathbf{V}'_{-j})\Big\}, \tag{15}$$

$$\text{s.t. } V_{j,t,e} \leq U_{j,t}(w) + \beta E V_{t+1}(V',\mathbf{V}'_{-j}), \tag{16}$$

$$\rho_{j,t+1,e+1}(V';\mathbf{V}'_{-j}) = S\big([e_j = e+1,\overline{0}],V',\mathbf{V}'_{-j}\big), \tag{17}$$

$$w \geq 0, \tag{18}$$

where $\pi(\cdot)$ represents the discounted stream of future per-worker profits of the firm, and $\rho(\cdot)$ represents the retention probability of that worker as a function of the future promised utility. The first constraint is the promise-keeping constraint, which guarantees that the workers obtain a value from working at firm j which is at least as large as the value that the firm already promised. In the final period of a worker's employment, the problem simplifies to a static problem for the firm

$$\pi_{j,T,e}(V_{j,T,e}) = \max_{w}\Big\{A_{j,T}\cdot(1+e)^{\gamma_1} - w(1+\tau_T^p)\Big\}, \tag{19}$$

$$\text{s.t. } V_{j,T,e} \leq U_{j,T}(w), \tag{20}$$

where the optimal wage policy is pinned down by the final promised value: $w_{j,t,e}(V_{j,T,e}) = U_{j,T}^{-1}(V_{j,T,e})$. The initial promised value $V_{j,t,0}$ is pinned down by the maximization problem of total expected profits. In this case, when choosing the initial value to attract new workers, the firm faces a trade-off between higher (lower) per-unit profit or lower (higher) mass of new workers. A higher initial promised value would attract more workers at the expenses of lowering the per-worker firm's profits. On the other hand, a lower initial promised value would increase the firm's per-worker profits but would deteriorate the firm's desirability from the workers' perspective. This problem can be characterized as follows

$$\Omega_{j,t,0}(\mathbf{V}_{-j}) = \max_{V_{j,t,0}}\Big\{\pi_{j,t,0}(V_{j,t,0};\mathbf{V}_{-j})D_{j,t}(0,V_{j,t,0};\mathbf{V}_{-j})\Big\}$$

$$\text{s.t. } D_{j,t}(0,V_{j,t,0};\mathbf{V}_{-j}) = \begin{cases} S\big(\overline{0},V_{j,t,0},\mathbf{V}_{-j}\big) & \text{if } t=1 \\ \sum_{h\neq j}\sum_{e_h=1}^{t-1} S\big([e_h,\overline{0}],V_{j,t,0},\mathbf{V}_{-j}\big)\phi_{h,t-1}(e_h-1) & \text{if } t>1 \end{cases},$$

where $\phi_{h,t-1}(e_h - 1)$ represents the end-of-last-period equilibrium mass of workers working at firm h with tenure $e_h - 1$ at time $t - 1$, who are currently deciding where to work. This equilibrium object is assumed to be taken as given by firm j.[13]

Definition 1 (Optimal Wage-Tenure Contract)**.** *An optimal wage-tenure contract is the solution to the firm's dynamic problem, which consists of*

- *Initial promised value at time of hiring:* $V^*_{j,t,0}$.
- *Promised values over a worker's tenure in the firm:* $V'_{j,t,e}(\cdot)$.
- *Wage policy over a worker's tenure in the firm:* $w_{j,t,e}(\cdot)$.

Definition 2 (Labor Market Equilibrium)**.** *The labor market equilibrium is defined as a set of wage-tenure contracts* $\left\{V^*_{j,t,0}, V'_{j,t,e}(\cdot), w_{j,t,e}(\cdot)\right\}_{j,t,e}$ *and firm-choice probabilities for workers* $\left\{S_{j,t}(\cdot)\right\}_{j,t}$ *such that*

- *The optimal wage contracts* $\left\{V^*_{j,t,0}, V'_{j,t,e}(\cdot), w_{j,t,e}(\cdot)\right\}_{j,t,e}$ *solve the firm's dynamic problem, given workers' choice probabilities* $\left\{S_{j,t}(\cdot; \mathbf{W}^*_t)\right\}_{j,t}$.
- *Firm choice probabilities for workers* $\left\{S_{j,t}(\cdot; \mathbf{W}^*_t)\right\}_{j,t}$ *solve their dynamic programming problem, given the set of optimal wage contracts of all the firms* $\left\{W^*_{j,t}(e) = w_{j,t,e}(V^*_{j,t,e})\right\}_{j,t,e}$.

5. Data and Setting

In this section, we introduce the Dutch matched employer-employee panel dataset and describe the temporary tax cut for migrant workers which we use to identify the model parameters. While we could estimate our model and quantify wage markdowns using only the matched employer-employee dataset, the tax cut provides a particularly attractive setting because the cost of compensating workers with wages rather than amenities changes predictably over workers' tenure. This allows us to highlight the interplay between dynamic wage contracts and firm heterogeneity, and to show how these factors influence estimates of wage markdowns.

[13] We assume that the firm cannot discriminate in the wage contract based on the firm of origin of workers. As a consequence of this assumption, workers who entered the labor market at the same time receive identical contracts upon switching to a new employer.

5.1. Migrants' Tax Exemption Policy

The Migrants' Tax Exemption Policy (MTEP) is a partial tax exemption allowing eligible high-skilled migrants to receive up to 30% of their gross salary tax-free for a limited duration.[14] The policy is intended to offset extraterritorial costs—such as relocation expenses and differences in the cost of living across countries—by providing a monthly lump-sum benefit. The MTEP aims to attract foreign workers with scarce skills, foster a more competitive and welcoming business environment in the Netherlands, and reduce administrative burdens for employers and employees (Dialogic, 2017).

To qualify for the exemption, an individual must satisfy several criteria. First, they must possess specific expertise that is scarce in the Netherlands—more on that below. Second, they must have been recruited from abroad. Third, within four months of the start of the employment relationship, the employer and employee must jointly apply for the employee's tax exemption by completing a three-page form. Once the exemption has been granted, an annual evaluation determines whether the specific-expertise criterion remains satisfied and whether the exemption can be maintained (Algemene Rekenkamer, 2016).[15]

The maximum duration of the exemption and the criteria for assessing specific expertise have changed over time. For workers whose eligibility began in 2012 or later, this duration was reduced from ten to eight years. In 2019, the maximum duration for new applicants was reduced to 5 years; employees who applied for the exemption prior to 2012 were not eligible from 2021 onward; and employees who applied for the exemption between 2012 and 2019 were eligible after 2020 only if they had used the exemption for less than five years in total. Previous periods of legal residence in the Netherlands are deducted from the maximum duration (Belastingdienst, 2019). Since additional features of the MTEP program were modified in subsequent reforms, our estimation sample is restricted to workers who were beneficiaries under the 10-year MTEP regime.[16]

Prior to 2012, specific expertise was evaluated based on education, experience, and whether the migrant would lose income from working in the Netherlands. In practice, however, the assessment was based on broad job categories, and applications were typically approved. For instance, a list of occupations—including "other"—was presumed always to satisfy the require-

[14] While preferential tax treatment for expats in the Netherlands dates back to the 1950s, the 30% cap has been in place only since 2001. There is no objective reason for an employer not to grant an eligible employee the maximum tax-free payment. On the contrary, employer-side social security contributions are also waived for the tax-free payment, and the costs of applying for the exemption are negligible. Therefore, net wages are highest for a given total labor cost when the full 30% payment is granted.

[15] An employee using the tax exemption loses the exemption when separating from her employer, but can reapply for the remaining years of eligibility with a new employer.

[16] The 2012 reform also introduced the requirement that migrants must have lived at least 150 kilometers from the Dutch border before moving.

ment. Moreover, the Ministry of Finance did not evaluate whether users of the tax benefits possessed specific and scarce expertise. Indeed, when in 2012 the specific-expertise test was replaced by an (annually indexed) taxable-income requirement of 35,000 euros, this reform was primarily aimed at saving 78 million euros per year by tightening the specific-expertise test (Algemene Rekenkamer, 2016). Employees who obtained the exemption before 2012 were also subject to the income requirement if they had used the tax benefits for less than five years prior to 2012 and at least five years in total (Belastingdienst, 2012).

The 2012 and 2019 reforms were discussed in Parliament only several months before implementation, making significant anticipation effects unlikely. The 30 May 2011 parliamentary meeting notes mention potential measures to restrict the use of the 30% rule but do not mention an income requirement or a reduction in the maximum duration (Tweede Kamer der Staten-Generaal, 2012).[17]

5.2. Matched Employer-Employee Data

We construct a matched employer-employee dataset for 2006 to 2021 using Dutch administrative data on tax records from Statistics Netherlands.[18] These tax records allow us to track employees and their income, including all employer- and employee-side taxes and social security contributions, as they switch employers. Our wage measure corresponds to the total labor cost of an employee. It is defined as the sum of that employee's taxable income, all mandatory employer-side income taxes and social security contributions, and, if applicable, the tax-free payment due to the 30% rule.[19] It is this measure, not take-home pay, that should equal the marginal revenue product of labor in a competitive labor market. We merge these data with employee-level identifiers for beneficiaries of the MTEP provided by the Dutch Ministry of Finance (available from 2009 to 2019). In addition, we observe employees' year and country of birth, as well as their sex.

In our empirical analysis, we are interested in how wages change around the anticipated end of the tax cut, as this is when the relative cost of compensating workers with wages versus non-wage amenities changes. Therefore, we restrict attention to employees whom we can classify according to their behavior after the benefits end; specifically, employees who either leave the Netherlands

[17] Although the October 2017 coalition agreement mentions an eventual further reduction in duration in passing—one sentence of 70 pages deals with the MTEP—only in the 20 April 2018 Parliament meeting notes are the 2019 changes discussed (Tweede Kamer der Staten-Generaal, 2018).

[18] Under certain conditions, these data are accessible for research. See https://www.cbs.nl/en-gb/our-services/customised-services-microdata for details.

[19] The tax-free payment, as well as income taxes and social security contributions, are calculated based on the employee's taxable income. If this income exceeds the minimum required to pass the specific expertise test, the maximum tax-free payment amounts to 3/7th of the taxable income.

or remain employed at a Dutch firm after they stop using the MTEP. We further restrict attention to employees aged 25 to 59 working full-time for non-government, non-financial, for-profit employers. In addition, we focus on the 10-year regime when constructing moments, as only this regime has sufficient observations post-MTEP. Finally, we limit the analysis to firms that employ at least one individual who does not switch jobs or leave the Netherlands when the tax benefits expire—a requirement for characterizing firms by the change in wages their employees experience when the tax cut expires. Appendix A provides a detailed overview of the sample selection procedure.

In estimating the model, part of the variation we exploit comes from the wage dynamics of workers as they approach and move past the point of benefit expiration. Since a worker's productivity is unlikely to change abruptly just before or after benefits expire, this variation is informative about the extent to which firms can smooth or redistribute monetary compensation over the course of a worker's tenure, and about how this ability differs across firms characterized by different non-monetary amenities.

6. Firm Heterogeneity and Wage-Tenure Contracts

In this section, we apply our theoretical model to the MTEP context to build the intuition behind the moments we use to estimate the model in Section 7. We develop two central insights. First, a worker's valuation of a firm hinges on a specific bundle of characteristics: the firm's revenue productivity and non-wage amenities, denoted by (A_j, G_j). Heterogeneity in this bundle across firms generates variation in the optimal design of wage-tenure contracts, as firms face different incentives in attracting and retaining workers. Second, wage markdowns evolve differently across firms over the course of the employment relationship because of differing incentives to attract and retain workers. Specifically, the more amenities a firm has and the more productive it is, the more it optimally reallocates wages over the employment relationship as the relative cost of compensating workers with wages changes. As a result, the timing and magnitude of wage markdowns can vary substantially across firms, reflecting differences in their productivity levels, amenity offerings, and strategic considerations in designing long-term compensation structures.

To illustrate how firm heterogeneity shapes incentives in wage-tenure contract design, we examine how firms adjust wages in response to the end of MTEP eligibility for their workers, particularly as they seek to retain affected workers. In our model, the MTEP implies a tax rate τ_t on labor income that depends on the number of years t a worker has been employed in the

Netherlands. Specifically, the tax rate takes the following form

$$(1-\tau_t) = \begin{cases} \varphi + (1-\varphi)(1-\tau) & \text{if } t < \overline{B} \\ (1-\tau) & \text{if } t \geq \overline{B} \end{cases}, \tag{21}$$

where $\varphi \in (0,1]$ is the tax-exemption parameter that governs the tax base before benefit expiration. After $\overline{B}$ years of employment in the Netherlands, as specified by the MTEP and anticipated by workers and firms, the entirety of labor income is taxed at the statutory tax rate τ.

For firms, payroll taxes apply only to the taxable portion of workers' earnings. Consequently, the effective payroll tax rate at any point in time depends on a worker's MTEP eligibility. Specifically, the payroll tax τ_t^p is given by:

$$(1+\tau_t^p) = \begin{cases} (1+(1-\varphi)\tau^p) & \text{if } t < \overline{B} \\ (1+\tau^p) & \text{if } t \geq \overline{B} \end{cases}, \tag{22}$$

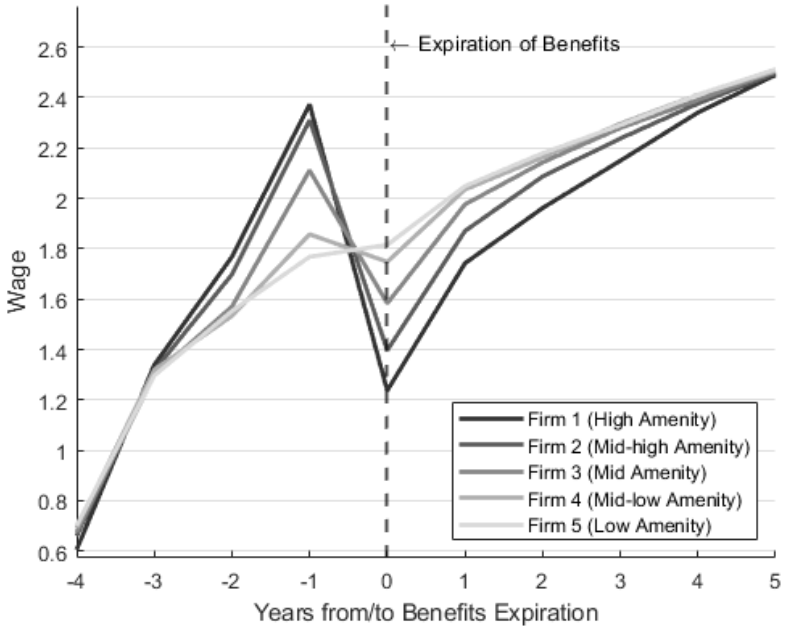
where τ^p is the statutory payroll tax rate and φ is the MTEP tax-exempt rate. This formulation captures the idea that payroll taxes are levied only on taxable labor income—hence, the firm's effective tax rate varies with the worker's eligibility for the exemption.

Here, we use a parameterized version of the model introduced in Section 4 to illustrate the theory's quantitative implications. This exercise highlights the model's ability to generate wage dynamics that underscore the role of dynamic contracts. Specifically, in the model, the firms' incentive to retain workers over time interacts with the MTEP in two key ways. First, for a given euro of wage paid, the firm recognizes that a worker takes home a larger fraction of that euro while the tax benefit remains active. Second, since payroll taxes apply only to the taxable portion of wages, the expiration of the worker's tax benefit mechanically raises the firm's payroll tax burden for a given gross wage. This mechanism further encourages firms to adopt dynamic rent-sharing (wage-setting) profiles that minimize tax incidence on their presented discounted value of profits.

Concretely, we focus on five firms and ten periods ($J = 5$ and $T = 10$). The expiration date of the tax exemption is set at $\overline{B} = 5$, with $\tau = 0.45$ and $\varphi = 0.3$. The firms differ only in amenities G_j, decreasing by 0.2 from $G_1 = 1$ to $G_5 = 0.2$, while revenue productivity is equal to $A_j = 1$ for all firms. Therefore, variation in wages over tenure reflects heterogeneity across firms solely in the amenities they provide. To prevent the misleading impression that the following quantitative results are driven by time variation in the payroll tax τ_t^p, we set it to zero in all periods (i.e., $\tau_t^p = 0$ for all t) for this example. In other words, in this quantitative exercise, the firm's incentive for dynamic wage setting stems solely from the recognition that

such wage–tenure contracts generate value for workers. By enhancing workers' perceived value of staying, dynamic contracts increase retention and, consequently, raise firms' expected profits. In the full estimation of the model, we set $\tau^p = 0.2$ to reflect the payroll tax rate implied by the fiscal framework of the Netherlands, thereby fully capturing firms' wage-setting incentives.[20]

Figure 1: Wage-Tenure Contracts by Firms' Amenities



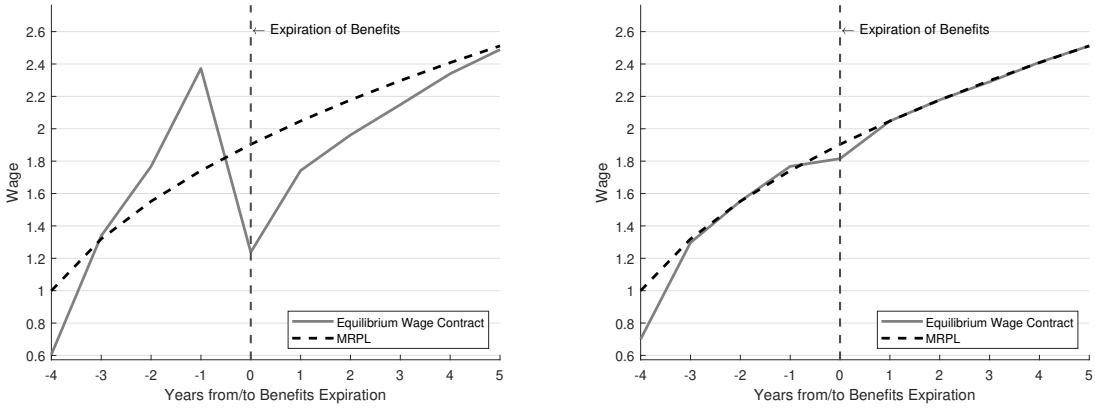
Notes: This figure plots simulated yearly wages five years before and after the end of MTEP benefits for employees who work for the same employer in all ten periods, broken down by the non-wage amenities of their employer. Simulation includes five firm types ($J = 5$) and ten periods ($T = 10$). We set the labor income tax rate to forty-five percent ($\tau = 0.45$), set the tax-exemption parameter to thirty percent ($\varphi = 0.3$), and the expiration date to five periods ($\overline{B} = 5$). On the worker side, we set the preference parameters to $(\lambda_1, \lambda_2, \alpha, \beta) = (5, 0.95, 1, 1)$ and the amenity values to $\{G_j\}_{j=1}^{5} = (1, \frac{4}{5}, \frac{3}{5}, \frac{2}{5}, \frac{1}{5})$. On the production side, we set the technological parameters to $(A_j, \gamma_1, \gamma_2) = (1, 0.2, 0.2)$, which implies no heterogeneity in A across firms.

Figure 1 depicts the wage–tenure profiles across firms for employees who remain with the same employer for all ten periods. The figure highlights that even when firms share identical productivity, heterogeneity in time-invariant amenities alone can generate substantial differences in wage dynamics. High-amenity firms optimally reallocate wage payments over the tenure profile—for instance, front-loading compensation in anticipation of the expiration of temporary tax benefits—to sustain workers' perceived value of employment. By doing so, they increase the likelihood of longer retention, greater accumulation of firm-specific human capital, and continued attachment beyond the benefit period. This strategy is optimal precisely because such firms

[20]More information on Dutch payroll taxes can be found in the official definition by the Netherlands Tax and Customs Administration (Belastingdienst): payroll taxes consist of wage tax, national insurance contributions, employee insurance contributions and the income-dependent health care insurance contribution. Source: belastingdienst.nl (accessed October 22, 2025).

can rely on both higher worker utility from nonpecuniary amenities and the productivity gains associated with tenure. In contrast, low-amenity firms, facing weaker attachment and stronger outside competition later in the employment relationship, find such anticipatory wage-setting less profitable and therefore adopt flatter wage profiles.

Figure 2: Wage Markdowns by Firms' Amenities



(a) Firm 1 (High Amenity) (b) Firm 5 (Low Amenity)

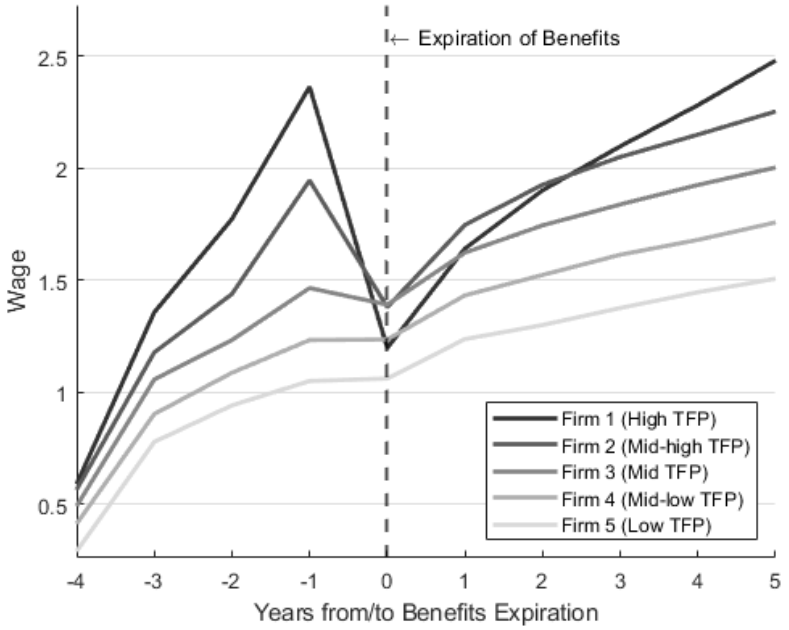
Notes: The figure reproduces the wage-tenure contracts for the highest-amenity firm (Panel a) and the lowest-amenity firm (Panel b) from Figure 1, and adds the marginal revenue product of labor (MRPL). As productivity is constant across firms and the figures are constructed for workers that do not switch employer, the MRPL is equal in both panels.

Moving to wage markdowns, Figure 2 reproduces the wage-tenure contracts for the highest-amenity and lowest-amenity firms, alongside the marginal revenue product of labor (MRPL). We include this figure to illustrate two key points. First, it highlights the mechanisms underlying the results shown in Figure 1. High-amenity firms retain workers by paying wages above the MRPL early in the employment relationship, but later reduce wages below the MRPL, as workers have accumulated substantial firm-specific skills and are willing to stay with the firm to enjoy the high amenity value. In contrast, low-amenity firms, which face (albeit imperfect) competition in retaining workers, are compelled to pay wages that closely track their revenue contribution. This limits their ability to reallocate compensation over time and reduces their flexibility to shift wages in anticipation of the expiration of the tax exemption. In the absence of random utility shocks, all workers prefer to be employed by the highest-amenity firm. As a result, lower-amenity firms face "vertical" competition and must raise wages—potentially close to the marginal revenue product of labor—to attract and retain workers.

Second, Figure 2 illustrates that a researcher interested in comparing wages to the marginal revenue product of labor (i.e., estimating wage markdowns) should account for the dynamic nature of the employment relationship. The wage markdown in a particular year is not representative of the wage markdown over the employment relationship, and the degree to which this

is the case depends on the non-wage amenities a firm offers. Therefore, a dynamic approach provides a more appropriate framework than a static approach. In Section 8, we compare markdowns estimated under static and dynamic wage setting and show that the static wage markdowns follow a very different and linear pattern, although even there the markdowns vary somewhat over time and across firms.[21]

Figure 3: Wage-Tenure Contracts by Firms' Productivity



Notes: This figure plots simulated yearly wages five years before and after the end of MTEP benefits for employees who work for the same employer in all ten periods, broken down by employer revenue productivity (TFPR). Simulation includes five firm types ($J = 5$) and ten periods ($T = 10$). We set the labor income tax rate to forty-five percent ($\tau = 0.45$), set the tax-exemption parameter to thirty percent ($\varphi = 0.3$), and set the expiration date to five periods ($\overline{B} = 5$). On the worker side, we set the preference parameters to $(\lambda_1, \lambda_2, \alpha, \beta) = (5, 0.95, 1, 1)$ and the TFPR values to $\{A_j\}_{j=1}^{5} = (1, 0.9, 0.8, 0.7, 0.6)$. Moreover, we set the technological parameters to $(\gamma_1, \gamma_2) = (0.2, 0.2)$, and fix amenities at $G_j = 1$ for all firms j.

Figure 3 repeats the exercise of Figure 1, but holding amenities fixed across firms at $G_j = 1$ and varying revenue productivity instead. Hence, variation in wages over tenure now reflects differences solely in the revenue-generating capacity of firms. As with amenities, more productive firms are better able to front-load wages before the expiration of the tax benefits. While productivity does not directly enter worker preferences, it enables firms to offer higher wages, making them more attractive to workers. However, unlike heterogeneity in amenities, productivity differences generate parallel wage–tenure profiles across firms, rather than convergence: more productive firms offer higher present values of wages, and less productive firms lower ones.

[21] Appendix Figure B-1 shows that absent the discontinuity induced by the expiration of the MTEP, the wage contracts for firm one and five are monotonically back-loaded.

By contrast, variation in amenities alters the timing of payments while leaving the present value largely unchanged (Figure 1). Thus, in the model, productivity primarily governs wage levels, whereas amenities primarily govern the intertemporal allocation of wages.

Guided by these results, our identification strategy—discussed in the next section—leverages a rich set of moments focused on wage profiles in continuing employer-employee relationships and on the dynamics immediately surrounding MTEP expiration, broken down by firm heterogeneity.

7. Estimation

In this section, we bring the model to the data by matching it to a combination of cross-sectional and longitudinal moments from the Dutch matched employer–employee panel dataset. We describe the estimation procedure and conclude with an assessment of the model's fit, providing a heuristic discussion of identification that maps key parameters of the model to the specific moments used in estimation. Beyond our specific application, the model can also be identified and estimated in other empirical contexts that feature different tax systems or rely exclusively on matched employer-employee data without any tax changes.

7.1. Mapping the Model to the Data

Firm Groups An important first step in bringing the model to the data involves modeling firm heterogeneity parsimoniously, which reduces the curse of dimensionality arising from the large number of firms in the Dutch data. To address this issue, we divide firms into groups, which capture heterogeneity in both revenue productivity and non-wage amenities across firms. For our baseline estimation, we select five firm types.[22]

As discussed in Section 4 and Section 6, a worker's valuation of an employment relationship depends on both wages and amenity values. Accordingly, workers rank firms based on the total promised value they offer—whether through higher wages, more attractive amenities, or a combination of both. In anticipation of the expiration of tax benefits, firms have an incentive to front-load wage payments, thereby reducing the present value of workers' tax liabilities. While this incentive applies to all firms, those with relatively high amenities are particularly well-positioned to take advantage of the intertemporal trade-off. These firms can offer a higher

[22]We think of a firm type as a large number of firms with similar revenue productivity and non-wage amenities—which makes them comparable in terms of how attractive they are to workers—and we model a representative firm for each group. All else equal, a greater number of firm types increases competition in the labor market, which in turn reduces wage markdowns.

overall value to workers and are better able to shift compensation toward the pre-expiration period. After the tax benefits expire, these firms can reduce pre-tax wages while maintaining retention via the appeal of amenities and the value of accumulated firm-specific human capital.

We use these model-based insights to classify firms into five distinct groups using the data. Specifically, we rely on the quintiles of the distribution of average labor cost changes around the expiration of tax benefits. This variable includes the firm's payroll taxes associated with the employment relationships, in addition to the wage payments. This is the appropriate measure of labor cost in the presence of payroll taxes from the firm's perspective, as it captures the correct counterpart to the marginal revenue product of labor and the associated profits that firms derive from hiring workers. Moreover, by incorporating wage information, the average labor cost helps identify firm types—the larger the average wage drop at a given firm, the greater the bundle of non-wage amenities and revenue productivity that the firm provides.[23]

Importantly, this approach captures how firms differ in their wage responses to changes in workers' tax rates. Reassuringly, this classification aligns ex post with other indicators of firm attractiveness derived from worker behavior—namely, observed patterns of worker mobility across firm groups, which reflect revealed preferences in the spirit of Sorkin (2018).

An attractive feature of our firm ranking is that it allows non-wage amenities to be unobserved by the researcher, thereby accommodating a wide variety of such amenities (Lamadon, Mogstad, and Setzler, 2022). In Appendix C, we investigate what type of amenities our measure captures using data from the Netherlands Employers Work Survey (NEA), a nationally representative survey that systematically records working conditions, employment relations, work-life balance, and employee well-being in the Netherlands. Changes in average labor costs associated with wage changes at the expiration of MTEP benefits—the statistic used to form the groups—are correlated with multiple measures of worker satisfaction, in particular with measures of job flexibility and work-life balance. This external validation is especially valuable because direct measures of amenities are typically unavailable.

Tax Rates In the model, we include two tax rates. First, wages are taxed according to the tax function in (21). The tax exemption parameter φ is set to 0.3 to align with the 30% tax exemption policy, while the average income tax rate τ is set to 0.45.[24] In addition, firms are subject to a payroll tax as defined in (22), with a statutory rate of $\tau^p = 0.2$.

[23] In the previous numerical exercise in Section 6, average wages are equal to the average labor cost, as payroll taxes were set to zero.

[24] In 2013, for example, the income tax in the Netherlands consisted of three brackets: income up to €19,645 was taxed at 37%; income from €19,646 to €55,911 at 42%; and income over €55,911 at 52%.

Workers' Outside Option Another element of the model that requires discussion is the worker's outside option, which—given our sample of migrants—corresponds in the data to the value of leaving the Netherlands altogether. Note that the tax exemption also applied to Dutch citizens returning to the Netherlands after a period of work abroad. However, the share of this group is negligible.[25]

To proceed, we expand the workers' choice set to include the outside option of leaving the Netherlands, a choice exercised with notable frequency in the data—approximately 15 percent annually on average. Specifically, the outside option depends on the worker's tenure e, a time trend t, and corresponding quadratic terms:

$$v^o_{j,t}(e_j) = \phi_0 + \phi_1 \cdot e_j + \phi_2 \cdot e_j^2 + \phi_3 \cdot t + \phi_4 \cdot t^2. \tag{23}$$

We assume that the outside option is an absorbing state, so once a worker chooses it, there is no way to return to the Netherlands. This assumption is dictated by the observation that the share of leavers who later return is very small in our data.[26]

Two features of the specification in (23) are worth mentioning. The tenure terms define heterogeneity in the value of the outside option for workers based on the accumulated tenure at the firm. This feature implies that firm-specific endogenous tenure profiles generate heterogeneity in the value of the outside option across firms. Second, time-related terms are common across workers in different firms and capture factors influencing the decision to leave the Netherlands that go beyond wage-related considerations tied to tenure.

7.2. Estimation and Model Fit

We estimate the model using the Simulated Method of Moments (SMM), which involves selecting a set of informative moments from the data and estimating the structural parameters as the set of parameters that make the corresponding simulated moments closely match their empirical counterparts.[27]

We rely on 75 moments to identify the model's parameters. Average labor costs for a balanced sample of employees who work at the same firm three years prior to and after MTEP expiration, broken down by firm group (30 moments). The same moments constructed for the lower half of the tenure distribution at expiration (30 moments). Exit probabilities and the

[25]For the workers for whom we have country of birth information (92.9% of our sample), 98.47% are foreign nationals, and missing country of birth itself is an indication that the individual is not Dutch.

[26]Of the 14,092 workers who at some point leave, 959 return; of those, 697 leave again.

[27]In practice, all wage-related moments are expressed in units of 100,000 euros. We rescale the probability-related moments by a factor of 100 to ensure they are of comparable magnitude to the wage-related moments. In the appendix, we report results obtained without this rescaling.

Table 1: Parameter Estimates

	Estimates	Standard Errors
Panel A: Preferences		
λ_1	1.8707	0.0174
λ_2	0.6714	0.0117
α	0.9388	0.0352
ϕ_0	22.5387	0.2689
ϕ_1	0.3276	0.0148
ϕ_2	-0.0649	0.0026
ϕ_3	1.3925	0.0113
ϕ_4	-0.2415	0.0022
ψ	4.0318	0.0233
Panel B: Technology		
γ_1	0.0279	0.0040
γ_2	0.2456	0.0085
A_1	0.8773	0.0223
A_2	0.7594	0.0214
A_3	0.7646	0.0180
A_4	0.7419	0.0200
A_5	0.8099	0.0220
Panel C: Amenities		
G_1	1.0853	0.0198
G_2	1.0000	(-)
G_3	0.7824	0.0208
G_4	0.7813	0.0292
G_5	0.6066	0.0251

Notes: This table displays the parameter estimates of the model described in Section 4. Asymptotic standard errors are reported.

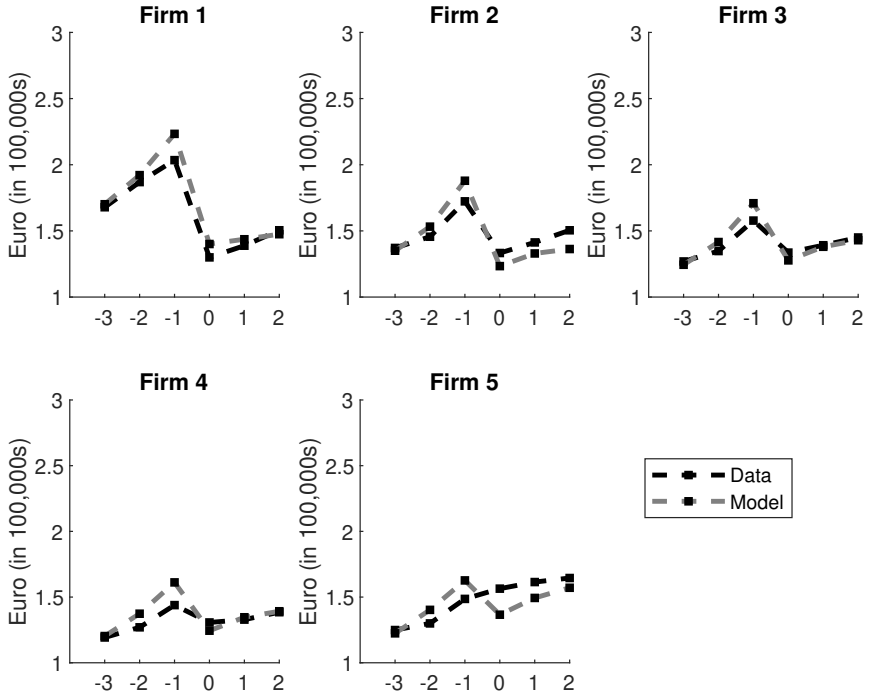
associated average labor costs for employees who leave the Netherlands, by tenure (6 moments). The probability that an employee leaves the Netherlands by firm group (5 moments), and the probability that a worker switches from the highest-ranked firm in terms of non-wage amenities and productivity to each of the other firm groups (4 moments). Table 1 presents the estimated parameters. Instead of commenting on each estimate separately, we first evaluate the model's overall fit. We then discuss, for each group of fitted moments, the parameters that play the key role in matching those moments and discuss their corresponding estimates.

As discussed, in the model, the dynamics of labor costs around the expiration of the tax benefit—driven by strategic wage adjustments—and their heterogeneity across firms contain key information about firms' revenue productivity and amenity values. The model predicts that high-amenity firms shift rents forward in anticipation of the benefit expiration for workers, resulting in a larger wage drop at the time of expiration than at low-amenity firms. In addi-

tion, higher firm-specific productivity raises the overall wage level throughout the employment relationship. Thus, both the level of labor costs and their changes around benefit expiration serve as important sources of identification for the model.

We compare the data and model on an equal footing. In both the data and the model, to account for selection and compositional effects on wage variation over time, we focus on a balanced panel of workers who remain with the same firm throughout the observation window. In particular, data and model generated moments centered on the expiration of tax benefits are calculated based on the criterion that workers stay in the same firm for six consecutive years, spanning three years before the benefit expiration, the year of expiration, and two years after the expiration. Given this criterion, the sample naturally includes workers with different tenures. More precisely, at the time of expiration, all workers by construction have at least four years of tenure in the firm, with some having more than four.

Figure 4: Sample Fit for Average Labor Cost Before and After the Expiration of the Tax Exemption



Notes: This figure shows average labor costs in the estimated model and data before and after the expiration of the tax benefit by firm type. Appendix Figure B-2 reports results for the same moments using equal weights across moments in the SMM estimator. The corresponding parameters are reported in Appendix Table B-1. Appendix Figure B-3 illustrates the wage dynamics after aggregating across all firms.

Figure 4 shows that the model replicates the dynamic profile of average labor costs before and after the tax benefit expiration closely. In both the model and the data, for the four

firms with the most non-wage amenities, average labor costs increase in the years leading up to expiration, drop in the year the benefit ends, and rise again afterward. These dynamics are driven by changes in worker wage compensation set by firms. The magnitude of the wage drop varies by firm type. While firms were ranked based on the size of these drops, the model's ability to capture their magnitude reflects a good fit. Firms with higher amenity values display greater reductions in average labor costs relative to firms with lower amenity values. For the lowest-ranked firm (Firm 5), the model reproduces the overall upward wage trajectory and its magnitude, though it captures less of the specific pattern around benefit expiration. Given the normalization of Firm 2's amenity to $G_2 = 1$, our estimated firm-specific amenities, $[G_1, \ldots, G_5] = [1.0853, 1, 0.7824, 0.7813, 0.6066]$, are scaled by the estimated preference parameter $\alpha = 0.9388$, which captures workers' valuation of firm amenities (in logs). This yields a natural monotonic ranking of firms from the highest-amenity (Firm 1) to the lowest-amenity (Firm 5).

In terms of firm-specific productivity, our estimated TFPs closely mirror the overall differences in average labor costs across firms, which are driven by differences in average paid wages: $[A_1, \ldots, A_5] = [0.8773, 0.7594, 0.7646, 0.7419, 0.8099]$. These values display a U-shaped pattern, in contrast to the monotonically decreasing profile of firm-specific amenities. This distinction aligns with the intuition that amenities primarily drive the behavior of wages around the MTEP expiration, while TFP mainly determines the overall level of wage profiles over a worker's tenure. For instance, the estimated value $A_1 = 0.8773$ indicates that a newly hired worker at Firm 1 in the Dutch labor market produces an output equivalent to roughly €87,700 in value.

In addition, Figure 4 provides valuable information for the model's estimation by illustrating average labor cost dynamics prior to the expiration of tax benefits. This pattern is informative because, as firms optimally smooth wage compensation over time to retain workers, the timing of rent sharing depends on the curvature of the worker's utility function (λ_2) with respect to wages. Specifically, this parameter governs how workers trade off wages against non-monetary amenities in a compensating differential framework, making the wage level itself a key determinant of workers' willingness to trade-off amenities for wages. Our estimates imply that preferences are non-linear in wages, with $\lambda_2 = 0.6714$, indicating concavity. The quantitative magnitude of compensating differentials also depends on the estimated $\lambda_1 = 1.8707$. Given these preference parameters $(\lambda_1, \lambda_2, \alpha)$, we find that replacing Firm 1's amenities with those of Firm 5 would reduce workers' ex-ante welfare by an amount equivalent to a 17 percent wage cut across all firms.[28]

[28]This result is obtained by solving the estimated model to quantify the welfare loss experienced by workers when Firm 1's amenities are reduced and by identifying the wage reduction that generates an equivalent welfare loss. The analysis implies that a uniform 17 percent decrease in wage–tenure contracts across all firms and

To identify the role of tenure on wage profiles, we use a set of moments characterizing the dynamics of average labor costs associated with a subsample of workers who reached the tax exemption expiration with shorter firm tenure than those who stayed at the same firm for all years of the tax exemption.[29] This set of moments aims to inform the model about how within-firm tenure interacts with wage dynamics just before and just after the benefit expiration. From the perspective of our model, the same firm faces distinct wage-setting incentives for early-tenure workers at the time of benefit expiration. In particular, the structure of wage–tenure contracts is shaped by the extent to which tenure raises revenue productivity. Moreover, the firm's incentives to front-load wages in anticipation of future benefit expiration critically depend on its ability to retain workers with varying tenure levels after the benefits expire. Figure 5 shows that the model, with an estimated value of the return to tenure $\gamma_1 = 0.0279$, successfully reproduces the additional wage profiles for this sample of workers who experience the termination of tax benefits early in their tenure at the current firm. Before-tax wages increase in the years leading up to the benefit expiration, drop at the expiration of the tax advantage, then start increasing again afterward. The estimated value of γ_1 indicates that, within a given period, a worker's productivity is approximately 2.8 percent higher than that of a coworker with one year less tenure in the same firm. Finally, fitting the overall wage trends in Figures 4 and 5—beyond the two-year window comparing periods just before and after the tax benefit expiration—provides information about the value of the time trend component in productivity ($\gamma_2 = 0.2456$).

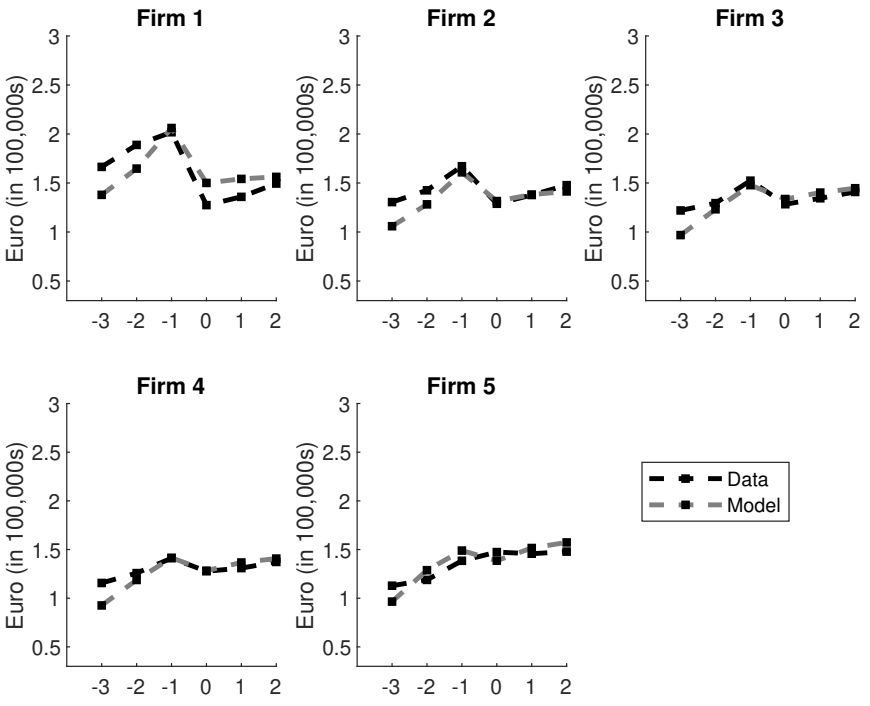
Another source of valuable information is the workers' decisions to stay or leave the Netherlands. To capture this information, we consider a number of data moments, measuring the frequency with which workers take the outside option of leaving the Netherlands, as well as their wages at exit. Specifically, we target the sample probabilities of exit by firm type and tenure, as well as the average labor costs paid by firms for leavers by tenure. These moments are informative about the conditional choice probability of taking the outside option, given a worker's tenure and current employer, and about the characteristics of leavers, as revealed by their wage prior to departure.

Appendix Figures B-4 and B-5 show the model's ability to reproduce the profiles of the probability of leaving the Netherlands by firm ranking and workers' tenure. In the model as in the data, the probability of leaving the Netherlands increases as we move from the highest-ranked to the lowest-ranked firm. The model also produces the hump-shaped profile of the frequency with which workers leave the Netherlands, thus exercising the outside option in the model, by tenure. In addition, Appendix Figure B-6 shows that the model produces the correct selection among those who leave the Netherlands. Among workers who take the outside option, pre-exit wages rise with tenure, alongside corresponding increases in average labor costs, consistent with

periods would yield the same welfare loss as the reduction in Firm 1's amenities.

[29] At expiration, median tenure is seven years in our sample.

Figure 5: Sample Fit for Average Labor Cost of Low-Tenure Workers Before and After the Expiration of the Tax Exemption



Notes: This figure shows average labor costs in the estimated model and data before and after the expiration of the tax benefit by firm type (as in Figure 4). However, it focuses on workers who reached the expiration of the tax benefit with tenure of at most seven years, the sample median at expiration.

the data. The parameters that are critical for the model to replicate the moments related to workers leaving the Netherlands are, naturally, those governing the attractiveness of the outside option. Our estimates, $[\phi_0, \ldots, \phi_4] = [22.5387, 0.3276, -0.0649, 1.3925, -0.2415]$, imply a positive outside option that increases with tenure and time but at a decreasing rate—that is, it is increasing and concave in both dimensions.

Finally, we target the frequency with which workers currently in Firm 1, the highest-ranked firm given its amenity value and revenue productivity, switch to each lower-ranked firm. These four moments are informative about the switching cost for workers, as well as about the relative ranking of firms along the amenity–productivity dimension. Appendix Figure B-7 shows that the model, with an estimated switching cost of $\psi = 4.0318$, reproduces well the monotonically decreasing profile of the probability of switching in the data from the highest-ranked firm to each of the other firm. The model's ability to reproduce the declining profile of switching probabilities is consistent with the theory, which provides further validation that the ranking of firms we identified from the data is plausible. Recall that we use the wage dynamics around benefit expiration to rank firms. It is reassuring that the model is also consistent with an

alternative natural ranking based on revealed preferences, workers' choices of where to work, akin to the empirical approach in Sorkin (2018).

8. Quantification

This section discusses the main quantitative implications of the estimated model. We begin by evaluating the role of wage-tenure contracts in wage markdowns. Using several experiments, we show that estimating markdowns based on the static notion of labor-supply elasticity—that is, the elasticity of firm-specific labor supply with respect to the same-period wage—commonly used in empirical work leads to highly misleading conclusions about the magnitude of markdowns. As discussed, in the presence of dynamic wage setting, the static (short-run) elasticity of labor supply is *not* a sufficient statistic for markdowns. In particular, a short-run inelastic labor supply in our framework is not evidence of labor market power but rather of firms' ability to retain workers by promising future value in the employment relationship.

8.1. Model-Based Wage Markdowns

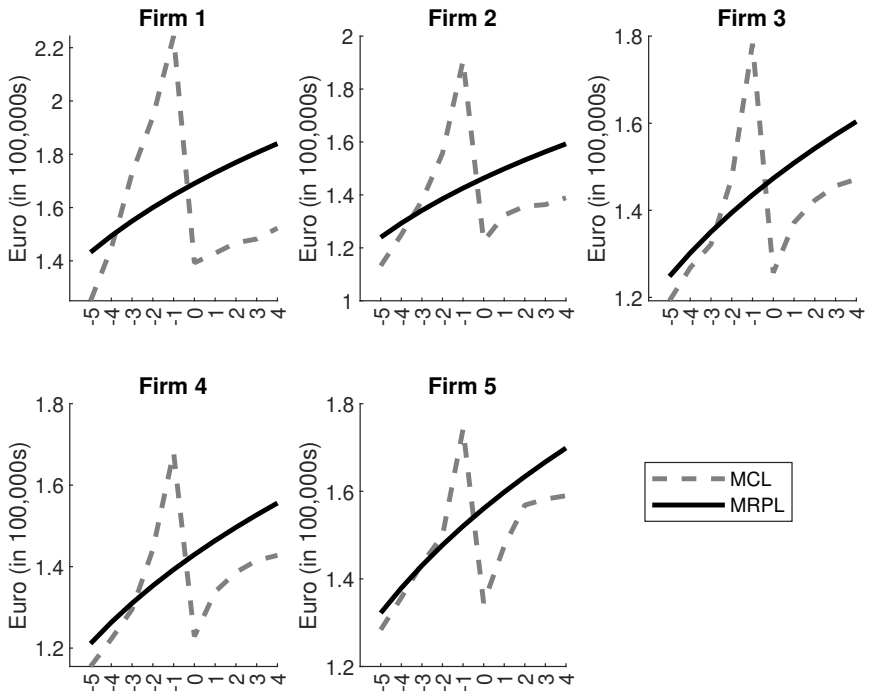
This subsection uses the estimated model to quantify labor market rents, defined as markdowns of workers' marginal labor costs relative to their marginal revenue product of labor (MRPL). To sharpen our insight about the role of wage-tenure contracts, we compare the implications of the estimated model with those of a static counterpart in which firms are not allowed to set wages dynamically; instead, we force firms to set the wage period-by-period to maximize current profits rather than the present discounted value of profits as in the wage-contracting model.

Wage Markdowns in the Estimated Model with Wage-Tenure Contracts Figure 6 displays the marginal cost of labor per worker (MCL) implied by the wage–tenure contracts, alongside the corresponding marginal revenue product of labor (MRPL) by firm type in the estimated model. The MRPL varies across firm types because we allow for heterogeneity in exogenous revenue productivity—captured by firm-specific technological parameters $A_{j}{}_{j=1}^{5}$—in addition to heterogeneity in amenities.

The estimated model produces significant variation in wage markdowns over a worker's tenure. Note that marginal costs of labor are *above* the marginal revenue product of labor for a few years before the expiration of the tax benefits. During this period, wages are marked up rather than marked down. However, this pattern is only temporary: wages fall below the

MRPL at the tax-benefit expiration date and thereafter increase toward the MRPL over the remainder of tenure. On average, wages are marked down, but firms use wage contracts to shift compensation intertemporally to provide static and dynamic incentives, enhancing worker retention and thereby firm size and profitability.

Figure 6: Wage-Tenure Contracts—MCL vs. MRPL



Notes: This figure compares the marginal cost of labor (MCL) with the marginal revenue product of labor (MRPL) implied by the optimal dynamic wage–tenure contracts for each firm type in the estimated model. Estimates of underlying model parameters are in Table 1.

As the figure illustrates, the average markdown is modest, averaging around 10% over the course of tenure. Under the standard static sufficient-statistics logic, this would imply a high firm-specific labor-supply elasticity. To test whether that logic holds in our setting, we compute a short-run labor-supply elasticity by marginally increasing the wage in a single model period while keeping all other contract wages fixed, doing so for workers at different tenures. We find an average short-run elasticity of 0.27—consistent with prior firm-level estimates of short-run labor supply elasticities—thereby breaking the tight link between markdowns and elasticities typically assumed in the literature. While this elasticity may appear small at first glance, it is natural in our environment, as it is calculated from a one-period, transitory change in current wages, whereas workers' labor supply decisions are shaped by dynamic considerations related to human capital accumulation (Imai and Keane, 2004). Put differently, the elasticities we estimate are not primitive parameters but rather endogenous outcomes of firms' use of dynamic

contracts to retain workers. The relatively muted responsiveness of labor supply to temporary wage changes is itself evidence that firms successfully smooth rents intertemporally to maintain worker attachment and preserve profitability.

Wage-Contracts vs. Static Wages A static version of our model has firms set wages period-by-period to maximize static profits:

$$\pi_{j,t,e} \equiv \max_{w \geq 0} \left(A_{j,t}(1+e)^{\gamma_1} - w\right) D_{j,t}(e,w), \tag{24}$$

where $D_{j,t}(e,w)$ represents the upward-sloping labor supply of workers faced by firm j with tenure e at time t.

Figure 7 reproduces the wage contracts for the estimated model shown in Figure 6 and the wages implied by the static-wage version given by equation (24). In both settings, the figure reports wages rather than average labor costs, thereby excluding the payroll tax rate. This choice highlights the differences in wage levels between the two wage-setting specifications. Here, we do not re-estimate the model with static wages; instead, we temporarily keep the same baseline parameter values. This is done to illustrate how, holding parameter values fixed, static wage setting can alter the main conclusions about the magnitude and dynamics of wage markdowns over workers' tenure.

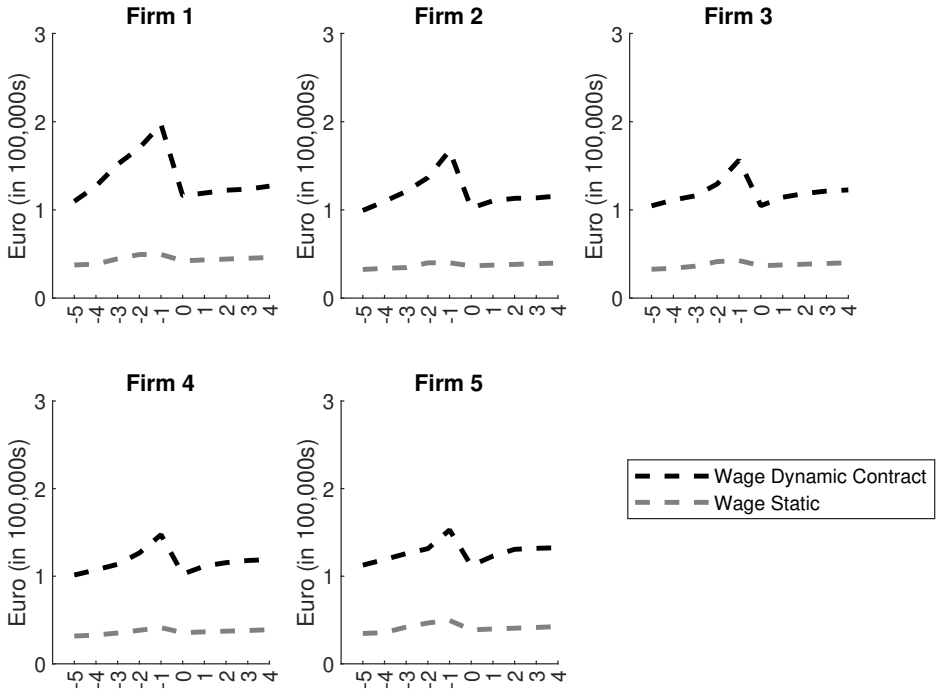
Two stark differences emerge. First, wages in the static model are substantially lower than in the estimated model with dynamic contracts. This gap is not especially informative, however, because the static model has not yet been disciplined by the data; re-estimating the productivity parameters would shift wage levels upward in the manner discussed above.

Second—and more salient—the static model largely eliminates wage variation at the tax-benefit expiration date: except for one firm type, there is no local adjustment at expiration. Instead, wage profiles are back-loaded over tenure, rather than exhibiting the discrete expiration-date adjustment characteristic of dynamic contracts. Importantly, these qualitative features persist when we estimate the static version of the model, to which we now turn.

In Figure 8, we place the two models on comparable footing by estimating the static-wage model using the same set of moments as the dynamic-contract model; parameter estimates are reported in Appendix Table B-2. To enable a meaningful comparison of firm profitability—both within and across model specifications—with respect to labor and its implications for labor market rents, we plot the marginal cost of labor (MCL) and the marginal revenue product of labor (MRPL), as in Figure 6.

The figure shows that the static model implies larger wage markdowns than the wage contract model—evidenced by the wider vertical gap between the marginal revenue product of labor and

Figure 7: Wage-Tenure Contracts vs. Static Wages



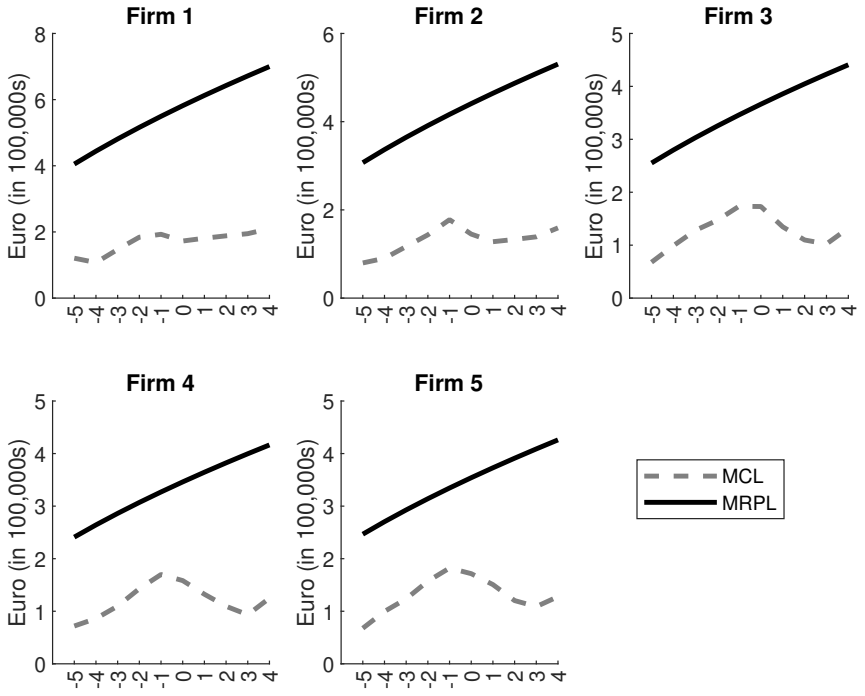
Notes: The figure reproduces the wage-tenure contracts from the estimated model and the wages implied by a static version of the estimated model in which we keep the parameters at their baseline estimated values in Table 1 but with a different wage-setting protocol whereby firms sets the wage period-by-period to maximize current profits.

equilibrium MCL. Although markdowns in the static model still vary over tenure, they lack the pronounced, piecewise linear pattern generated by dynamic wage contracts. Consistent with the preceding intuition, dynamic contracts allow firms to flexibly reallocate compensation over a worker's tenure, whereas the static-wage model, by construction, precludes such flexibility. While this comparison does not provide direct evidence of the specific wage–tenure contracts we model, it supports our central claim that firms engage in dynamic wage setting—an approach far less restrictive than the static wage-setting mechanism often assumed in the literature.

8.2. Share of Efficient Terminations

In the model, some terminations of existing employment relationships are inefficient: they could be avoided if one redistributes rents from the firm to the worker. Of course, the model also features efficient terminations in the sense that the firm's rents are insufficient to make the worker stay even if all firm's profits were to be given to the worker. The latter case occurs when the worker's utility gain from terminating the current employment relationship with

Figure 8: Static Wages—MCL vs. MRPL



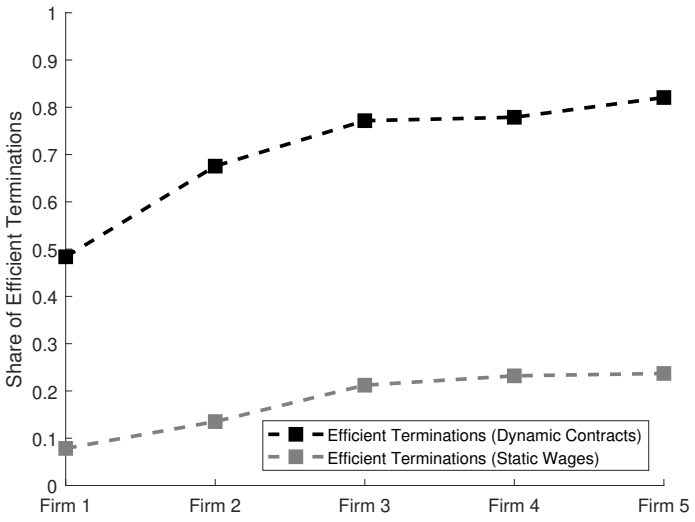
Notes: This figure compares the marginal cost of labor (MCL) with the marginal revenue product of labor (MRPL) as implied by the static-wage version of the model where we re-estimate the parameters to reproduce the same set of data moments used for the model with wage-tenure contracts. Estimated parameters for the static-wage model are in Appendix Table B-2.

tenure e and firm j and starting a new one with reset tenure $e = 0$ at firm j', calculated as $v_{j',t}(0) + \epsilon_{i,j',t} - v_{j,t}(e) - \epsilon_{i,j,t}$, exceeds the profits $\pi_{j,t}(e)$ the current employer can possibly render to the worker.

Figure 9 shows the share of total employment relationships that are efficient as defined above by firm types in the estimated model with wage-tenure contracts and in its static-wage version without contracts. Note that the static-wage model is re-estimated with the same moments used for the model with contracts so that we can compare the two models on an equal footing.

As is evident, the share of efficient terminations in the model with wage-tenure contracts is significantly larger than in the static-wage model. In fact, most terminations are efficient in the model with contracts, whereas the opposite is true for the model with static wages where most terminations are instead inefficient. This result underscores the critical role of the wage-tenure contracts as devices that greatly enhance the firms' ability to retain workers in a way that exploits most of the available gains from trade.

Figure 9: Share of Efficient Employment Terminations
Wage Contracts vs Static Wages



Notes: The figure shows the share of efficient terminations by firm types in the estimated model with wage-tenure contracts and its static-wage version without contracts. The static-wage model is re-estimated to the same data moments used for the model with the wage-tenure contracts. Estimated parameters for the model with contracts are in Table 1 and those for the static-wage model are in Appendix Table B-2.

9. Conclusion

This paper develops a model of oligopsonistic competition in the labor market in which firms use wage-tenure contracts to attract and retain workers. Firms differ by non-wage amenities and productivity, allowing the model to reproduce the observed variation of wages over tenure and across firms. We estimate the model parameters using matched employer-employee data from the Netherlands, coupled with administrative tax records. We then use the estimated model to quantify labor market power, gauging how wage markdowns vary over workers' tenure.

The key lesson is that quantifying markdowns critically depends on how the labor-supply elasticities firms face vary over a worker's tenure at the firm. Firms do not take these elasticities as given; rather, they use wage-tenure contracts to increase retention, thereby reducing the elasticities they face. However, low firm-specific labor-supply elasticities are neither evidence of, nor a sufficient statistic for, high wage markdowns—contrary to the inference from static models in which firms set wages period-by-period to maximize current profits. Allowing firms to use wage contracts instead of static wages tilts the firms' trade-off toward larger firm size via improved retention, at the cost of lower per-worker markdowns.

References

Algemene Rekenkamer. 2016. "Fiscale Tegemoetkoming Voor Experts uit het Buitenland: De 30%-Regeling." .

Ashenfelter, Orley C, Henry Farber, and Michael R Ransom. 2010. "Labor Market Monopsony." *Journal of Labor Economics* 28 (2):203–210.

Autor, David, David Dorn, Lawrence F Katz, Christina Patterson, and John Van Reenen. 2020. "The Fall of the Labor Share and the Rise of Superstar Firms." *Quarterly Journal of Economics* 135 (2):645–709.

Azar, José and Ioana Marinescu. 2024. "Monopsony Power in the Labor Market." In *Handbook of Labor Economics, Volume 5*, edited by C. Dustmann and T. Lemieux. Elsevier, 761–827.

Azar, José and Xavier Vives. 2021. "General Equilibrium Oligopoly and Ownership Structure." *Econometrica* 89 (3):999–1048.

Azar, José A, Steven T Berry, and Ioana Marinescu. 2022. "Estimating Labor Market Power." *NBER Working Paper no. 30365* .

Azariadis, Costas. 1975. "Implicit Contracts and Underemployment Equilibria." *Journal of Political Economy* 83 (6):1183–1202.

Bagga, Sadhika. 2023. "Firm Market Power, Worker Mobility, and Wages in the US Labor Market." *Journal of Labor Economics* 41 (S1):S205–S256.

Bagger, Jesper, François Fontaine, Fabien Postel-Vinay, and Jean-Marc Robin. 2014. "Tenure, Experience, Human Capital, and Wages: A Tractable Equilibrium Search Model of Wage Dynamics." *American Economic Review* 104 (6):1551–1596.

Bagger, Jesper and Rasmus Lentz. 2019. "An Empirical Model of Wage Dispersion with Sorting." *Review of Economic Studies* 86 (1):153–190.

Balke, Neele and Thibaut Lamadon. 2022. "Productivity Shocks, Long-Term Contracts, and Earnings Dynamics." *American Economic Review* 112 (7):2139–2177.

Belastingdienst. 2012. "Handboek Loonheffing 2012." .

———. 2019. "Handboek Loonheffing 2019." .

Berger, David, Kyle Herkenhoff, Andreas R Kostøl, and Simon Mongey. 2024. "An Anatomy of Monopsony: Search Frictions, Amenities, and Bargaining in Concentrated Markets." *NBER Macroeconomics Annual* 38 (1):1–47.

Berger, David, Kyle Herkenhoff, and Simon Mongey. 2022. "Labor Market Power." *American Economic Review* 112 (4):1147–1193.

———. 2025. "Minimum Wages, Efficiency, and Welfare." *Econometrica* 93 (1):265–301.

Burdett, Ken and Melvyn Coles. 2003. "Equilibrium Wage-Tenure Contracts." *Econometrica* 71 (5):1377–1404.

Cahuc, Pierre, Fabien Postel-Vinay, and Jean-Marc Robin. 2006. "Wage Bargaining With On-the-Job Search: Theory and Evidence." *Econometrica* 74 (2):323–364.

Card, David. 2022. "Who Set Your Wage?" *American Economic Review* 112 (4):1075–1090.

Card, David, Ana Rute Cardoso, Joerg Heining, and Patrick Kline. 2018. "Firms and Labor Market Inequality: Evidence and Some Theory." *Journal of Labor Economics* 36 (S1):S13–S70.

Chen, Kuan-Ming, Ning Ding, John A List, and Magne Mogstad. 2025. "Reservation Wages and Workers' Valuation of Job Flexibility: Evidence From a Natural Field Experiment." *Journal of the European Economic Association* :jvaf022.

Deb, Shubhdeep, Jan Eeckhout, Aseem Patel, and Lawrence Warren. 2024. "Walras–Bowley Lecture: Market Power and Wage Inequality." *Econometrica* 92 (3):603–636.

Delabastita, Vincent and Michael Rubens. 2025. "Colluding Against Workers." *Journal of Political Economy* 133 (6):1796–1839.

Dialogic. 2017. "Evaluatie 30%-Regeling." .

Erdem, Tülin, Susumu Imai, and Michael Keane. 2003. "Brand and Quantity Choice Dynamics Under Price Uncertainty." *Quantitative Marketing and Economics* 1 (1):5–64.

Gottfries, A and G. Jarosch. 2025. "Wage Fixing." *NBER Working Paper no. 33501* .

Gottfries, Axel and Gregor Jarosch. 2023. "Dynamic Monopsony with Large Firms and Non-competes." *NBER Working Paper no. 31965* .

Han, Qingyang. 2024. "Human Capital Transferability and Employer Monopsony Power." *Mimeo* .

Hendel, Igal and Aviv Nevo. 2006. "Measuring the Implications of Sales and Consumer Inventory Behavior." *Econometrica* 74 (6):1637–1673.

Holmstrom, Bengt. 1981. "Contractual Models of the Labor Market." *American Economic Review* 71 (2):308–313.

———. 1983. "Equilibrium Long-Term Labor Contracts." *Quarterly Journal of Economics* :23–54.

Huggett, Mark, Gustavo Ventura, and Amir Yaron. 2011. "Sources of Lifetime Inequality." *American Economic Review* 101 (7):2923–2954.

Imai, Susumu and Michael P. Keane. 2004. "Intertemporal Labor Supply and Human Capital Accumulation." *International Economic Review* 45 (2):601–641.

Jarosch, Gregor, Jan Sebastian Nimczik, and Isaac Sorkin. 2024. "Granular Search, Market Structure, and Wages." *Review of Economic Studies* 91 (6):3569–3607.

Jungerman, William. 2023. "Dynamic Monopsony and Human Capital." *Mimeo* .

Kahn, Matthew E and Joseph Tracy. 2023. "A Human Capital Theory of Who Escapes the Grasp of the Local Monopsonist." *NBER Working Paper no. 31014* .

Kline, Patrick. 2024. "Labor Market Monopsony: Fundamentals and Frontiers." In *Handbook of Labor Economics, Volume 5*, edited by C Dustmann and T. Lemieux. Elsevier.

Kroft, Kory, Yao Luo, Magne Mogstad, and Bradley Setzler. 2025. "Imperfect Competition and Rents in Labor and Product Markets: The Case of the Construction Industry." *American Economic Review* 115 (9):2926–2969.

Lamadon, Thibaut, Magne Mogstad, and Bradley Setzler. 2022. "Imperfect Competition, Compensating Differentials, and Rent Sharing in the US Labor Market." *American Economic Review* 112 (1):169–212.

Lentz, Rasmus. 2014. "Optimal Employment Contracts with Hidden Search." *NBER Working Paper no. 19988* .

Maestas, Nicole, Kathleen J Mullen, David Powell, Till Von Wachter, and Jeffrey B Wenger. 2023. "The Value of Working Conditions in the United States and Implications for the Structure of Wages." *American Economic Review* 113 (7):2007–2047.

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

Menzio, Guido and Shouyong Shi. 2010. "Block Recursive Equilibria for Stochastic Models of Search on the Job." *Journal of Economic Theory* 145 (4):1453–1494.

Robinson, Joan. 1933. *The Economics of Imperfect Competition*. MacMillan Press.

Rosen, Sherwin. 1986. "The Theory of Equalizing Differences." In *Handbook of Labor Economics, Volume 1*, edited by O.C. Ashenfelter and R. Layard. Elsevier.

Rubens, Michael, Yingjie Wu, and Mingzhi Xu. 2025. "Exploiting or Augmenting Labor." *American Economic Review: Insights* .

Sorkin, Isaac. 2018. "Ranking Firms Using Revealed Preference." *Quarterly Journal of Economics* 133 (3):1331–1393.

Thomas, Jonathan and Tim Worrall. 1988. "Self-Enforcing Wage Contracts." *Review of Economic Studies* 55 (4):541–554.

Treuren, Leonard. 2022. "Wage Markups and Buyer Power in Intermediate Input Markets." *KU Leuven Discussion Paper DPS22.06* .

Tweede Kamer der Staten-Generaal. 2012. "Wijziging van Enkele Belastingwetten en Enige Andere Wetten (Belastingplan 2012)." *Kamerstukken II 2011/12, 33003, nr. 10.* .

———. 2018. "Wijziging van Enkele Belastingwetten en Enige Andere Wetten (Belastingplan 2018)." *Kamerstukken II 2017/18, 34785, nr. 83.* .

Volpe, Oscar. 2024. "Job Preferences, Labor Market Power, and Inequality." *Mimeo* .

Yeh, Chen, Claudia Macaluso, and Brad Hershbein. 2022. "Monopsony in the US Labor Market." *American Economic Review* 112 (7):2099–2138.

Appendix

A. Data and Sample Selection

Data This paper uses Dutch administrative data for the years 2006 to 2021 obtained from Statistics Netherlands.[30] We construct a matched employer-employee dataset that allows us to track beneficiaries of the MTEP over time and across firms, drawing on four sources. The ABR is an annual registry of all Dutch firms based primarily on anonymized data from the Dutch Chamber of Commerce and the Dutch Tax Authority. It includes the firm's 4-digit NACE industry. Employee-level data on remuneration, income taxes, social security contributions, and contract hours are obtained from SPOLISBUS, which contains anonymized monthly information on all employees who work for a firm in the Netherlands for at least four hours in that month. We aggregate all SPOLISBUS data to the annual level, as variation in MTEP usage is available only at this level, and tax-exemption eligibility is determined based on taxable annual income. All monetary values are converted to 2015 euros. Information on employees' country and date of birth, as well as their sex, is obtained from GBAPERSOONTAB.

Annual information from 2009 to 2019 on the employees who use the MTEP is obtained from the Dutch Ministry of Finance. This information allows us to identify, at the worker–year level, who benefits from the tax exemption. Pre-2009 use of the tax exemption is imputed based on occasionally observed start dates and the fact that migrants must sign up for the exemption with their employer no later than four months after their employment relationship starts and must be recruited from abroad. To avoid misclassifying the first year based on only a few months, we drop workers who appear to use the agreement for more than the maximum allowed duration. Our main compensation measure is the wage, which equals the total labor cost to the firm for a given employee-year. We calculate the tax-free portion assuming it is as high as possible—given that taxes are waived on this portion, there is no reason for a firm to choose a lower amount. That is, the tax-free amount is 3/7 of taxable income. In practice, the MTEP is implemented via an addendum to a labor contract that specifies the tax-free income. This addendum may, but need not, specify that taxable income (that we observe) is adjusted for as long as the MTEP lasts. Tenure is the number of consecutive years a worker has worked at a given firm.

Sample Selection We focus on a sample of workers who use the MTEP at some point and who work at a nonfinancial, for-profit firm—i.e., we exclude government organizations, financial

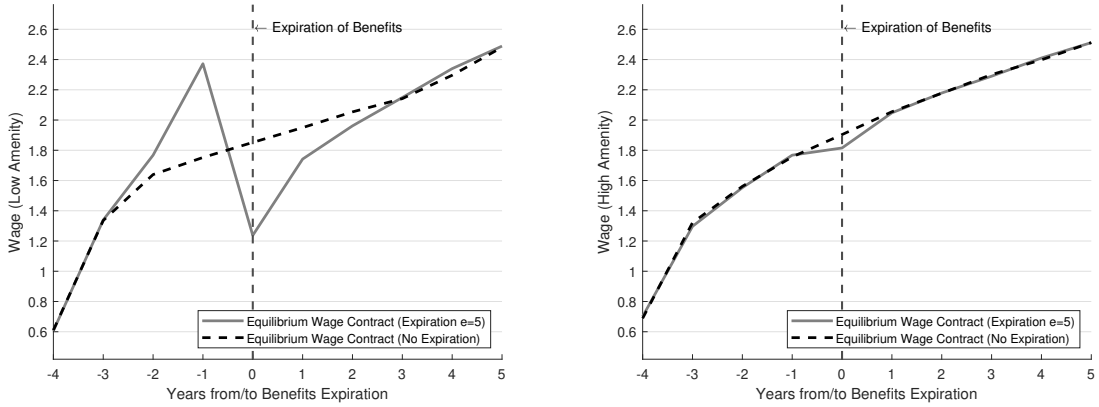
[30] See https://www.cbs.nl/en-gb/our-services/customised-services-microdata for access conditions and data documentation.

firms, nonprofits, and entrepreneurs. As we rely on wage variation around the final year of the tax exemption, we restrict attention to employees whom we can classify by their behavior after the benefits end; that is, employees who either leave the Netherlands or remain employed at a Dutch firm after they stop using the MTEP. Post-2019 data for employees who should still be eligible for the tax exemption in 2020 are dropped as we cannot entirely exclude the possibility that these workers stop using the agreement for unobserved reasons. We drop data from workers who start using the MTEP in 2019 or for whom the 2019 changes retroactively applied.

We further restrict attention to employees aged 25-59 working full-time. In the rare cases where an employee holds two full-time jobs, we drop that employee. To classify firms by the change in wages their employees experience when the tax cut expires, we limit the analysis to firms that employ at least one individual who remains with the firm when the tax benefits expire. Finally, we drop employees in the top and bottom percentiles of the distributions of average wage and average year-on-year wage change, as well as the top and bottom percentiles of employees who stay at the same firm after MTEP expiration, based on the change in wage around expiration and their wage in the final year of MTEP use. At the firm level, we drop the top and bottom two percentiles of firms based on the average change in wages at MTEP expiration and the average wage in the final year of MTEP use. The final sample contains 179,500 observations covering 47,499 workers and 2,546 firms. When constructing specific moments, we make further restrictions as discussed in the main text—for instance, focusing on the 10-year regime or on a balanced panel of workers who do not switch firms after MTEP expiration.

B. Additional Tables and Figures

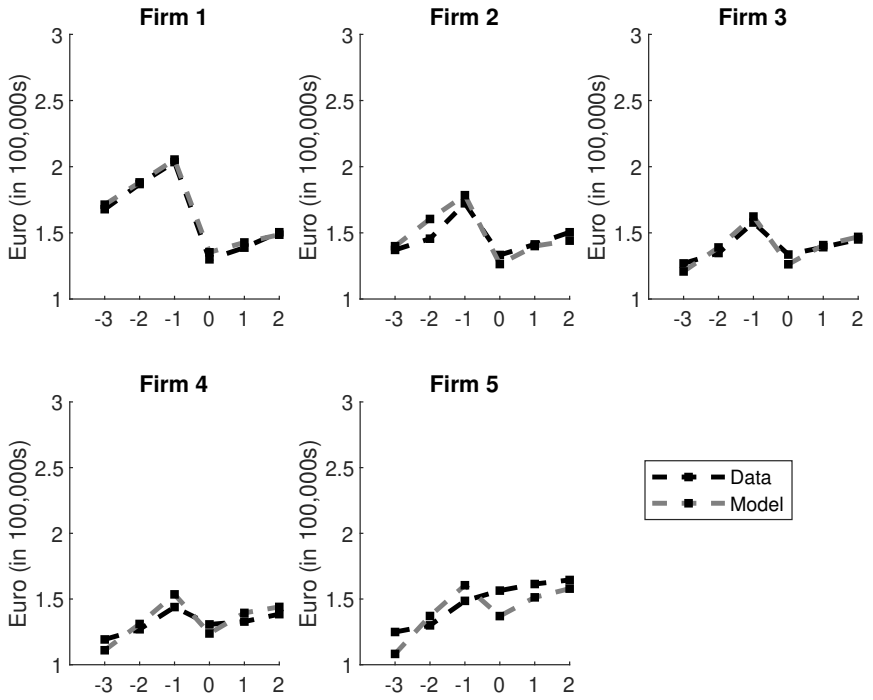
Figure B-1: Wage–Tenure Contracts With and Without the Tax Benefit Policy



(a) Firm 1 (High Amenity) (b) Firm 5 (Low Amenity)

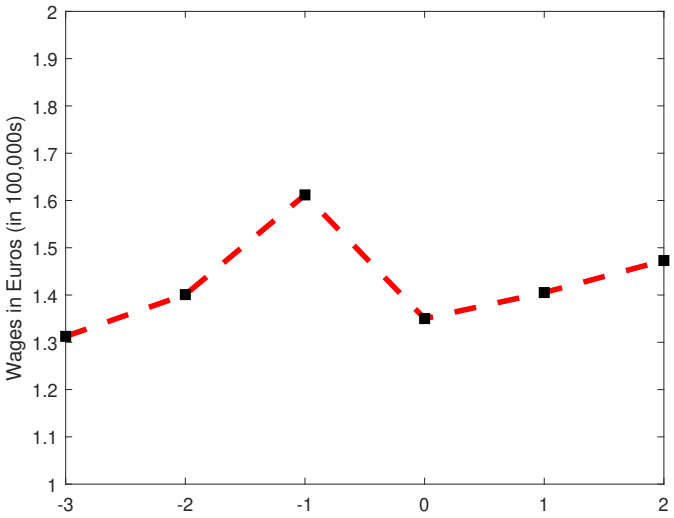
Notes: This figure shows the wage–tenure contracts for the highest-amenity firm (left panel) and the lowest-amenity firm (right panel) with and without the tax benefit. We simulate an economy with five firm types ($J = 5$) and ten periods ($T = 10$). We set the labor income tax rate to forty-five percent ($\tau = 0.45$), set the tax-exemption parameter to thirty percent ($\varphi = 0.3$), and set the expiration date to five periods ($\overline{B} = 5$). On the worker side, we set the preference parameters to $(\lambda_1, \lambda_2, \alpha, \beta) = (5, 0.95, 1, 1)$ and the amenity values to $\{G_j\}_{j=1}^{5} = (1, \frac{4}{5}, \frac{3}{5}, \frac{2}{5}, \frac{1}{5})$. On the production side, we set the technological parameters to $(A_j, \gamma_1, \gamma_2) = (1, 0.2, 0.2)$, which implies no heterogeneity in A across firms.

Figure B-2: Sample Fit for Average Labor Costs Before and After the Expiration of the Tax Exemption



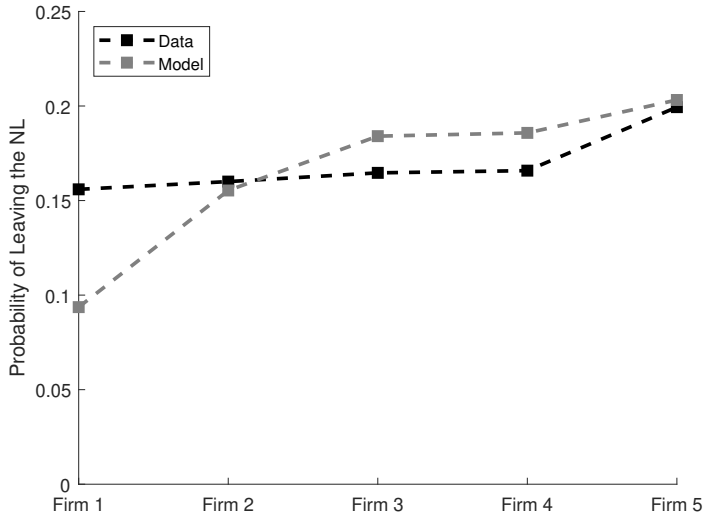
Notes: This figure shows average labor costs in the estimated model and in the data before and after the expiration of the tax benefit, by firm type. It replicates Figure 4, but the model is estimated using equal weights across moments in the SMM estimator.

Figure B-3: Average Labor Costs Before and After the Expiration of the Tax Benefit



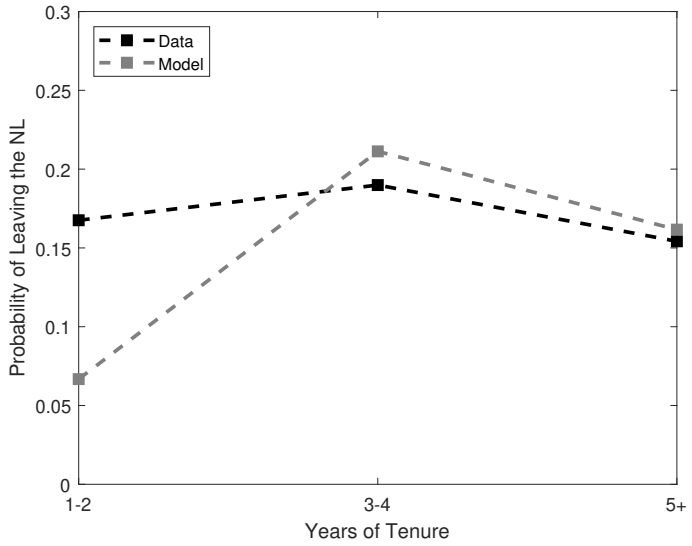
Notes: This figure plots the yearly average labor cost three years before and after the end of MTEP benefits for employees who remain with the same employer in all six years.

Figure B-4: Sample Fit for the Probability of Leaving the Netherlands, by Firm Type



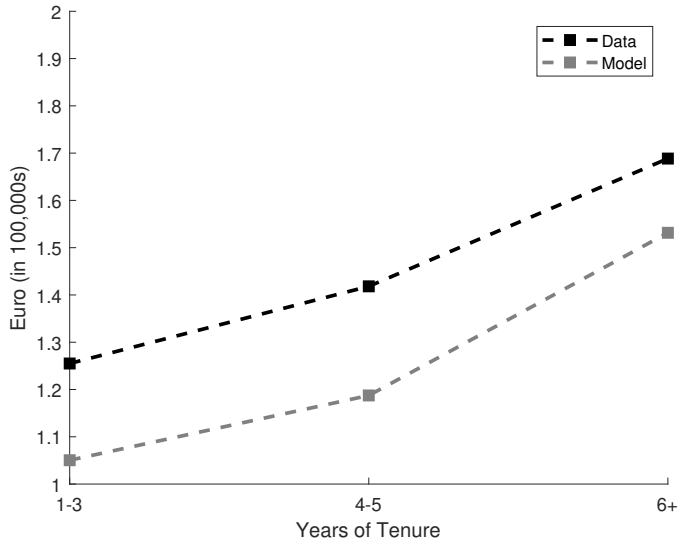
Notes: This figure shows the probabilities of leaving the Netherlands in the estimated model with wage–tenure contracts and in the data, by firm type.

Figure B-5: Sample Fit for the Probability of Leaving the Netherlands, by Tenure



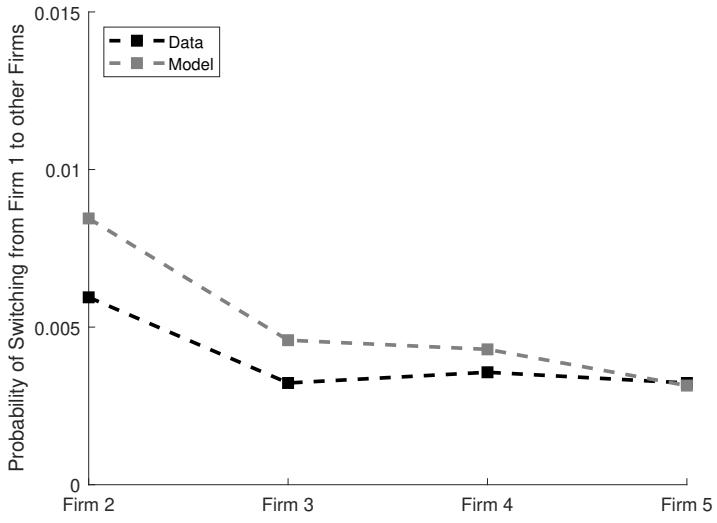
Notes: This figure shows the probabilities of leaving the Netherlands in the estimated model with wage–tenure contracts and in the data, by tenure.

Figure B-6: Sample Fit for Average Labor Costs—Leavers



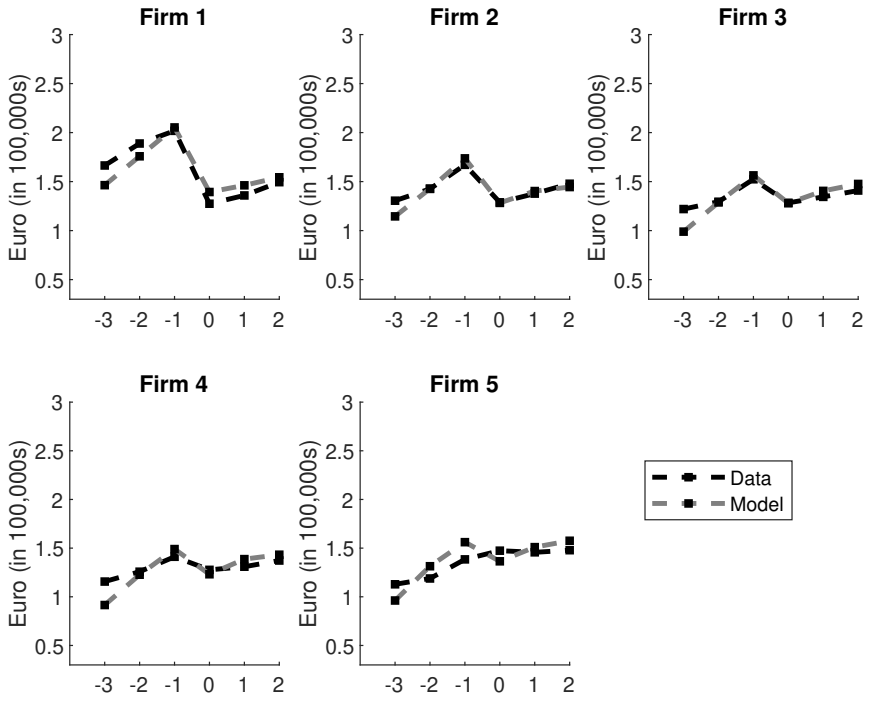
Notes: This figure shows the average labor costs paid by firms for workers who leave their current employer, by tenure, in the estimated model with wage–tenure contracts and in the data.

Figure B-7: Sample Fit for the Probability of Switching from Firm 1 to Other Firms



Notes: This figure shows the probabilities of switching from Firm 1 to the other firms.

Figure B-8: Sample Fit for Average Labor Costs of Low-Tenure Workers Before and After the Expiration of the Tax Exemption



Notes: This figure shows average labor costs in the estimated model and in the data before and after the expiration of the tax benefit, by firm type. It replicates Figure 5, but the model is estimated using equal weights across moments in the SMM estimator.

Table B-1: Parameter Estimates (Equally Weighted SMM)

	Estimates	Standard Errors
Panel A: Preferences		
λ_1	2.0094	0.0397
λ_2	0.6600	0.0106
α	0.9433	0.0413
ϕ_0	22.4507	0.7379
ϕ_1	0.3490	0.0155
ϕ_2	-0.0616	0.0027
ϕ_3	1.3582	0.0153
ϕ_4	-0.2463	0.0049
ψ	3.0315	0.0872
Panel B: Technology		
γ_1	0.0290	0.0015
γ_2	0.2539	0.0045
A_1	0.8778	0.0120
A_2	0.7816	0.0142
A_3	0.7625	0.0128
A_4	0.7422	0.0154
A_5	0.8068	0.0155
Panel C: Amenities		
G_1	1.1168	0.0284
G_2	1.0000	(-)
G_3	0.8131	0.0280
G_4	0.7696	0.0232
G_5	0.6053	0.0164

Notes: This table replicates the estimates in Table 1, but the model is estimated using equal weights across moments in the SMM estimator. Asymptotic standard errors are reported.

Table B-2: Parameter Estimates (Static-Wage Model)

	Estimates	Standard Errors
Panel A: Preferences		
λ_1	2.0061	0.0144
λ_2	0.8413	0.0192
α	0.9564	0.0197
ϕ_0	23.7680	0.1194
ϕ_1	0.4094	0.0093
ϕ_2	-0.0231	0.0017
ϕ_3	0.4874	0.0085
ϕ_4	-0.1780	0.0009
ψ	4.4117	0.0937
Panel B: Technology		
γ_1	0.0368	0.0028
γ_2	0.5598	0.0052
A_1	1.3910	0.0125
A_2	1.0543	0.0133
A_3	0.8760	0.0096
A_4	0.8276	0.0112
A_5	0.8466	0.0196
Panel C: Amenities		
G_1	0.8460	0.0090
G_2	1.0000	(-)
G_3	0.7250	0.0185
G_4	0.7103	0.0129
G_5	0.6141	0.0178

Notes: This table reports estimates for the static-wage model, in which firms set wages period-by-period to maximize static profits as in Equation (24). Asymptotic standard errors are reported.

C. Non-Wage Amenities

In this appendix, we compare the measure that we use to rank firms into groups based on non-wage amenities and revenue productivity—the average change in wages for employees that stay at the same firm after MTEP expiration—to a series of survey questions on non-wage amenities to validate our measure and infer what type of amenities our analysis is picking up. We show that our theory-based measure aligns most strongly with flexibility and work–life balance—greater feasibility of working from home, more varied work, and less fatigue—and with higher engagement and job satisfaction, even within sectors and conditional on wage and age. It is largely unrelated to perceived job security or the social climate and shows only modest links to career dynamics. While there is limited direct evidence on amenities in the broader literature, our measure aligns well with earlier work based on survey data or field experiments (e.g., Maestas et al. (2023); Chen et al. (2025)).

Before introducing the survey data, we document how firms in our five groups are distributed across sectors. Amenities might be correlated with sector due to the nature of jobs differing across sectors. For instance, offering the possibility to work from home to workers engaged in production activities at manufacturing firms might not be possible. Table C-1 reports, for each of the five groups, counts and column percentages across six sectors following Autor et al. (2020): manufacturing, retail trade, wholesale trade, services, utilities and transport, and other. To test whether the sector distribution is the same across firm groups, we perform a joint Pearson χ^2 test of homogeneity on the 6×5 table, which rejects equality of the column distributions ($\chi^2(20) = 50.45$, $p = 0.000$). In terms of magnitude, however, the differences are small—Cramér's V equals 0.073—and only in manufacturing and services is there a (near-)monotonic relationship between firm groups and sector shares. Based on these results, we include sector fixed effects in the regressions that follow and show that our amenities measure primarily captures within-sector variation.

We use the Netherlands Employers Work Survey (NEA), obtained from Statistics Netherlands, which has been conducted annually since 2003.[31] The NEA is a large nationally representative survey designed to provide systematic information on working conditions, employment relationships, work–life balance, and well-being of employees in the Netherlands. It is based on a stratified random sample drawn from the Dutch population register, with stratification ensuring representativeness by age, gender, and sector. Data are collected using a mixed-mode design, combining online questionnaires with paper-and-pencil surveys to maximize accessibility and participation. The NEA is designed as a repeated cross-sectional survey, meaning that a

[31] See https://www.cbs.nl/en-gb/our-services/customised-services-microdata for access conditions and data documentation.

Table C-1: Distribution of Firms Across Sectors and Firm Groups

Sector	Firm 1	Firm 2	Firm 3	Firm 4	Firm 5	p_{χ^2}
Manufacturing	43 (9.66%)	85 (17.60%)	82 (16.57%)	84 (17.61%)	98 (19.76%)	0.001
Retail	10 (2.25%)	6 (1.24%)	17 (3.43%)	10 (2.10%)	21 (4.23%)	0.033
Wholesale	80 (17.98%)	85 (17.60%)	82 (16.57%)	83 (17.40%)	104 (20.97%)	0.437
Services	276 (62.02%)	267 (55.28%)	264 (53.33%)	249 (52.20%)	225 (45.36%)	0.000
U+T	19 (4.27%)	21 (4.35%)	33 (6.67%)	23 (4.82%)	28 (5.65%)	0.405
Other	17 (3.82%)	19 (3.93%)	17 (3.43%)	28 (5.87%)	20 (4.03%)	0.368
Total	445 (100%)	483 (100%)	495 (100%)	477 (100%)	496 (100%)	—

Notes: Cells display counts and, in parentheses, column percentages. p_{χ^2} reports row-wise Pearson χ^2 p-values (unadjusted) from 2×5 tests of equal proportions for each sector ("this sector" vs. "not this sector")..

new random sample of respondents is drawn each year rather than tracking the same individuals over time. While the survey is large, the probability that an individual worker or firm appears is low. Each annual wave surveys approximately 1 in every 850 employees, corresponding to about 0.12% of the Dutch workforce. Only 369 of the workers from our sample appear in the survey, for a total of 373 observations.[32] Moreover, questions can change in different versions of the survey, and response rates are question-specific. Therefore, the vast majority of questions in the NEA is unusable for our purposes. We focus on a set of 30 questions that are connected to non-wage amenities and for which sufficient data is available.

We correlate our measure that captures amenities and revenue productivity—the negative of the average change in wages for employees who remain with the same firm after MTEP expiration—with indicators based on the survey questions. For interpretability, we standardize our firm-level measure, $\widetilde{Amenity}_j$, to have mean zero and unit standard deviation. For each survey outcome, we report specifications without controls and with controls for log wage, age, and sector fixed effects. That is, for a survey-based outcome measure of respondent i in year t, Y_{it}, we run the following regression

$$Y_{it} = \beta_0 + \beta_{amen}\widetilde{Amenity}_j + \mathbf{X}'_{it}\boldsymbol{\beta} + \varepsilon_{it}, \tag{C-1}$$

where i, j, and t refer to employee, firm, and year, respectively. Standard errors are clustered at the firm level to account for possible dependencies across observations within a single firm.

Table C-2 correlates our measure of amenities to a series of questions relating to job security and switching. Columns 1 to 6 reveal that the non-wage amenities that a firm offers are unrelated to employees' perceived likelihood of losing one's job, and the perceived ease of getting a different job or position within the same company or at another company—in all

[32]These observations are split over firm groups 1 to 5, respectively, 20, 87, 104, 115, and 47.

Table C-2: Survey Outcomes: Job Switching and Security

Dependent variable:	Risk of job loss		Easily get job internally		Easily get job externally	
	(1)	(2)	(3)	(4)	(5)	(6)
Amenities (standardized)	-0.000	0.007	0.009	0.001	0.013	0.017
	(0.037)	(0.036)	(0.036)	(0.036)	(0.029)	(0.029)
Controls	no	yes	no	yes	no	yes
R^2	0.000	0.037	0.000	0.022	0.001	0.029
Observations	365	365	352	352	353	353
Dependent variable:	Stay despite equal offer		Considered job switch		Attempted job switch	
	(7)	(8)	(9)	(10)	(11)	(12)
Amenities (standardized)	0.060	0.068	-0.073	-0.072	-0.063	-0.058
	(0.035)*	(0.034)**	(0.039)*	(0.039)*	(0.039)	(0.040)
Controls	no	yes	no	yes	no	yes
R^2	0.009	0.025	0.010	0.045	0.008	0.020
Observations	317	317	365	365	364	364

Notes: This table reports output from regressions relating a binary indicator based on a survey question to our (standardized) amenities measure; Full questions and coding of the indicators is in Table C-7; All regressions include an unreported constant term; Controls consist of logged age and wage, and sector fixed effects; Standard errors clustered at the firm level; ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively.

cases the point estimate is close to zero. Nonetheless, regardless of whether we condition on wage, age, and sector, employees at high amenity firms are more likely to report no willingness to switch to a different employer if they could get a job at equal wage (Columns 7 and 8), and are less likely to have considered searching for a different job in the past year (Columns 9 and 10). They are also less likely to have actually attempted to switch jobs, although this difference is not significant at conventional levels (Columns 11 and 12). Summing up, firms that we classify as having more non-wage amenities have workers that appear less inclined to switch jobs than those at low-amenity firms, a fact that cannot be explained by their wage of perceived employability and job risk.

Table C-3 reports results on survey questions relating to education and career path. The level of amenities a firm provides is unrelated to it's workers highest level of attained education (Columns 1 and 2). However, employees of high-amenity firms are more likely to have received on-the-job training the past two year, and report less desire for further training—although the final difference is not statistically significant (Columns 3 to 6). While employees' tenure does not vary with the amenities that employers offer (Columns 7 and 8), employees at high-amenity firms report having spent less time in their current job than those at low-amenity firms (Columns 9 and 10). This might be partially explained by them reporting more often to have

Table C-3: Survey Outcomes: Education and Career Path

Dependent variable:	Attained higher education		Received on-the-job training		Desire on-the-job-training	
	(1)	(2)	(3)	(4)	(5)	(6)
Amenities (standardized)	-0.000	0.002	0.058	0.067	-0.043	-0.035
	(0.032)	(0.030)	(0.039)	(0.038)*	(0.043)	(0.045)
Controls	no	yes	no	yes	no	yes
R^2	0.000	0.037	0.007	0.051	0.003	0.032
Observations	356	356	328	328	331	331
Dependent variable:	Tenure		Years in current job		Recently promoted	
	(7)	(8)	(9)	(10)	(11)	(12)
Amenities (standardized)	0.000	0.163	-0.577	-0.573	0.038	0.051
	(0.489)	(0.443)	(0.283)**	(0.259)**	(0.035)	(0.033)
Controls	no	yes	no	yes	no	yes
R^2	0.000	0.323	0.011	0.142	0.003	0.051
Observations	363	363	356	356	370	370

Notes: The table reports output from regressions relating a binary indicator or integer based on a survey question to our (standardized) amenities measure; Full questions and coding of the indicators is in Table C-7; All regressions include an unreported constant term; Controls consist of logged age and wage, and sector fixed effects, and sector fixed effects; Standard errors clustered at the firm level; ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively.

received a promotion in the past two years, albeit this difference is insignificant at conventional levels (Columns 11 and 12). Overall, while amenities do not correlate with prior education or experience, there is some evidence that high-amenity firms offer more opportunities for training and promotion.

Table C-4 correlates our measure of amenities to survey questions on workload and job flexibility. The more amenities a firm provides, the less likely are its workers to report neglecting family due to work, and experiencing end-of-day exhaustion, although no major differences in reported work intensity exist (Columns 1 to 6). One explanation is greater flexibility and task variety at high-amenity firms—the more amenities, the more likely employees can work from home and report creative, varied work (Columns 7–10). Experiencing autonomy in how you organize your job, however, does not correlate with amenities. We conclude that our measure of amenities appears to coincide with lower work-induced fatigue and a larger job flexibility.

Table C-5 considers a set of questions related to physical and cognitive job demands, and wellbeing. Employees at high and low-amenity firms do not differ in the degree to which they feel that they can meet the physical and psychological demands of their jobs (Columns 3 to 6), and also do not rate the cognitive demands of their jobs differently (Columns 1 and 2). Nevertheless, the more amenities a firm provides, the less likely are employees to report feeling emotionally

Table C-4: Survey Outcomes: Workload and Flexibility

Dependent variable:	Neglect family for work		High work intensity		End-of-day exhaustion	
	(1)	(2)	(3)	(4)	(5)	(6)
Amenities (standardized)	-0.086	-0.087	-0.024	-0.019	-0.091	-0.088
	(0.036)**	(0.036)**	(0.041)	(0.041)	(0.033)***	(0.034)**
Controls	no	yes	no	yes	no	yes
R^2	0.019	0.043	0.001	0.041	0.020	0.061
Observations	366	366	372	372	371	371
Dependent variable:	Can work from home		Creative and varied work		Work autonomy	
	(7)	(8)	(9)	(10)	(11)	(12)
Amenities (standardized)	0.094	0.107	0.100	0.109	0.006	0.012
	(0.034)***	(0.033)***	(0.060)*	(0.063)*	(0.030)	(0.032)
Controls	no	yes	no	yes	no	yes
R^2	0.016	0.089	0.019	0.037	0.000	0.025
Observations	356	356	206	206	373	373

Notes: The table reports output from regressions relating a binary indicator based on a survey question to our (standardized) amenities measure; Full questions and coding of the indicators is in Table C-7; All regressions include an unreported constant term; Controls consist of logged age and wage, and sector fixed effects; Standard errors clustered at the firm level; ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively.

exhausted due to their work (Columns 7 and 8), the more likely they are to frequently feel work engagement (Columns 9 and 10), and the higher is their overall job satisfaction (Columns 11 and 12), even conditional on wage and age. Thus, even though self-assessed aptitude does not vary with amenities, employees are more engaged and satisfied in higher-amenity firms.

Finally, Table C-6 presents outcomes related to the social environment at work. The degree to which employees perceive their boss to care about their wellbeing does not correlate with amenities (Columns 1 and 2), and neither does the perception of how well your boss listens to what you say – if anything, bosses are perceived to listen slightly less well at high amenity firms, although not significantly so (Columns 3 and 4). Likewise, employees' assessment of how friendly their colleagues are, and of whether they show interest, is unrelated to amenities (Columns 5 to 8). Finally, although unwanted behavior – both by colleagues and customers – is slightly less likely to occur at high-amenity firms, these differences are not statistically significant (Columns 9 to 12). We conclude that our measure of non-wage amenities does not measure how employees perceive the social environment at their workplace.

Table C-5: Survey Outcomes: Job Demands and Wellbeing

Dependent variable:	High cognitive demands		Meet physical demands		Meet psychological demands	
	(1)	(2)	(3)	(4)	(5)	(6)
Amenities (standardized)	-0.003	0.002	0.005	0.003	0.025	0.021
	(0.031)	(0.030)	(0.013)	(0.014)	(0.018)	(0.018)
Controls	no	yes	no	yes	no	yes
R^2	0.000	0.044	0.000	0.023	0.004	0.023
Observations	372	372	355	355	354	354
Dependent variable:	Emotional exhaustion		Work engagement		Job satisfaction	
	(7)	(8)	(9)	(10)	(11)	(12)
Amenities (standardized)	-0.067	-0.069	0.121	0.168	0.087	0.086
	(0.033)**	(0.032)**	(0.062)*	(0.055)***	(0.029)***	(0.031)***
Controls	no	yes	no	yes	no	yes
R^2	0.010	0.038	0.027	0.175	0.018	0.022
Observations	372	372	90	90	364	364

Notes: The table reports output from regressions relating a binary indicator based on a survey question to our (standardized) amenities measure; Full questions and coding of the indicators is in Table C-8; All regressions include an unreported constant term; Controls consist of logged age and wage, and sector fixed effects; Standard errors clustered at the firm level; ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively.

Table C-6: Survey Outcomes: Social Environment at Work

Dependent variable:	Boss cares for wellbeing		Boss listens		Colleagues show interest	
	(1)	(2)	(3)	(4)	(5)	(6)
Amenities (standardized)	0.005	0.006	-0.033	-0.037	0.021	0.019
	(0.032)	(0.030)	(0.029)	(0.028)	(0.028)	(0.028)
Controls	no	yes	no	yes	no	yes
R^2	0.000	0.013	0.004	0.038	0.002	0.017
Observations	354	354	357	357	355	355
Dependent variable:	Colleagues are friendly		Unwanted behavior colleagues		Unwanted behavior customers	
	(7)	(8)	(9)	(10)	(11)	(12)
Amenities (standardized)	-0.010	-0.015	-0.031	-0.039	-0.025	-0.026
	(0.016)	(0.016)	(0.029)	(0.027)	(0.024)	(0.022)
Controls	no	yes	no	yes	no	yes
R^2	0.001	0.036	0.003	0.021	0.004	0.075
Observations	360	360	373	373	373	373

Notes: The table reports output from regressions relating a binary indicator based on a survey question to our (standardized) amenities measure; Full questions and coding of the indicators is in Table C-8; All regressions include an unreported constant term; Controls consist of logged age and wage, and sector fixed effects; Standard errors clustered at the firm level; ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively.

Table C-7: Survey Questions and Coding for Tables C-2 to C-4

Label	Full Question	Coding
Attained higher education	What is your highest completed level of education or training?	1: Higher professional (HBO) or academic training (WO), 0: All else
Attempted job switch	In the past year, have you taken any action to find other work?	1: yes, 0: no
Can work from home	Can you perform part of your work from home?	1: yes, 0: no
Considered job switch	In the past year, have you considered looking for other work than your job with your current employer?	1: yes, 0: no
Creative and varied work	Is your work varied and does it require creativity and learning new skills?	1: often or always, 0: rarely or never
Easily get job externally	I could easily get a new job/position with another employer.	1: (strongly) agree, 0: (strongly) disagree
Easily get job internally	I could easily get a new job/position with my current employer.	1: (strongly) agree, 0: (strongly) disagree
End-of-day exhaustion	Do you feel exhausted at the end of a workday?	1: yes, 0: no
High work intensity	Do you have to work very fast, very much, or extra hard?	1: often or always, 0: rarely or never
Neglect family for work	Do you miss or neglect family activities because of your work?	1: often or always, 0: rarely or never
Received on-the-job training	In the past 2 years, have you taken a course or training for your work?	1: yes, 0: no
Recently promoted	In the past 2 years, have you changed positions within your company?	1: yes, 0: no
Risk of job loss	Are you at risk of losing your job?	1: yes, 0: no
Stay despite equal offer	I would continue working for my current employer, even if I could do the same work for the same salary somewhere else.	1: (strongly) agree, 0: (strongly) disagree
Tenure	How many years have you been working for your current employer?	Integer (years)
Wants on-the-job training	Do you want to take a course or training for your work?	1: yes, 0: no
Work autonomy	Do you experience autonomy in your work?	1: often or always, 0: rarely or never

Table C-8: Survey Questions and Coding for Tables C-5 and C-6

Label	Full Question	Coding
Meet physical demands	I can easily meet the physical demands my job places on me.	1: (strongly) agree, 0: (strongly) disagree
Meet psychological demands	I can easily meet the psychological demands my job places on me.	1: (strongly) agree, 0: (strongly) disagree
Emotional exhaustion	I feel emotionally exhausted by my work.	1: (strongly) agree, 0: (strongly) disagree
High cognitive demands	Does your work require intense thinking, sustained focus, or a lot of attention?	1: often or always, 0: rarely or never
Job satisfaction	Overall, how satisfied are you with your work?	1: satisfied, 0: not satisfied
Work engagement	I experience work engagement.	1: at least several times per week, 0: never or weekly
Boss cares for wellbeing	My boss cares about the well-being of employees.	1: (strongly) agree, 0: (strongly) disagree
Boss listens	My boss pays attention to what I say.	1: (strongly) agree, 0: (strongly) disagree
Colleagues show interest	My colleagues have personal interest in me.	1: (strongly) agree, 0: (strongly) disagree
Colleagues are friendly	My colleagues are friendly.	1: (strongly) agree, 0: (strongly) disagree
Unwanted behavior colleagues	Have you ever experienced unwanted behavior from bosses or colleagues?	1: yes, 0: no
Unwanted behavior customers	Have you ever experienced unwanted behavior from customers?	1: yes, 0: no