TI 2026-091/III
Tinbergen Institute Discussion Paper

The Time-Varying Impact of Sectoral and Aggregate Labor Market Tightness on Wage Growth

Marente Vlekke [1]
Matthijs Katz [2]
Siem Jan Koopman [3]
Rixt Swierstra [4]

1 CPB Netherlands Bureau for Economic Policy Analysis
2 CPB Netherlands Bureau for Economic Policy Analysis
3 Vrije Universiteit Amsterdam
4 De Nederlandsche Bank

The Time-Varying Impact of Sectoral and Aggregate Labor Market Tightness on Wage Growth *

Marente Vlekke† Matthijs Katz ‡ Siem Jan Koopman§
Rixt Swierstra ¶

November 10, 2026

Abstract

We examine the time-varying relationship between labor market tightness and wage growth using a panel model of collective bargaining agreement outcomes in the Netherlands. Cross-sectional differences in sectoral tightness explain little of the variation in wage growth. Because time fixed effects absorb all aggregate variation, standard panel models identify this relationship only from cross-sectional variation: the so-called missing intercept problem. We therefore model the time fixed effect as a time-varying intercept, estimated with state space methods, and decompose its variation by relating it to aggregate labor market tightness and other macroeconomic variables. This allows us to disentangle the effects of sectoral and aggregate tightness. Aggregate tightness has a substantially larger impact on wage growth than sectoral tightness, but only when labor markets are exceptionally tight, consistent with a nonlinear wage Phillips curve.

Keywords: wage Phillips curve, labor market tightness, missing intercept problem, collective bargaining, state space methods, nonlinearities, principal component analysis.

JEL classification: E24, J31, J52, C33, C32, E31.

*We thank Freek van Sambeek and Luciënne Verbeek-Damshuizer for their excellent research assistance, and Guido Ascari, Robbert Coenmans, Sabrina Genz, Maarten Goos, Laurens Harteveld, Jan de Kramer, Christiaan van der Kwaak, Geoffrey Minne, Jante Parlevliet, Pim Paulusma, Emilie Rademakers, Anna Salomons, Marcel Timmer, Wiljan van den Berge, Jena de Wit and several participants at seminars at CPB, DNB, Vrije Universiteit Amsterdam and EUROFRAME '25 for valuable comments and suggestions.

†Corresponding author. CPB Netherlands Bureau for Economic Policy Analysis; Department of Econometrics, Vrije Universiteit Amsterdam. E-mail: m.vlekke@cpb.nl.

‡CPB Netherlands Bureau for Economic Policy Analysis. E-mail: m.d.b.katz@cpb.nl.

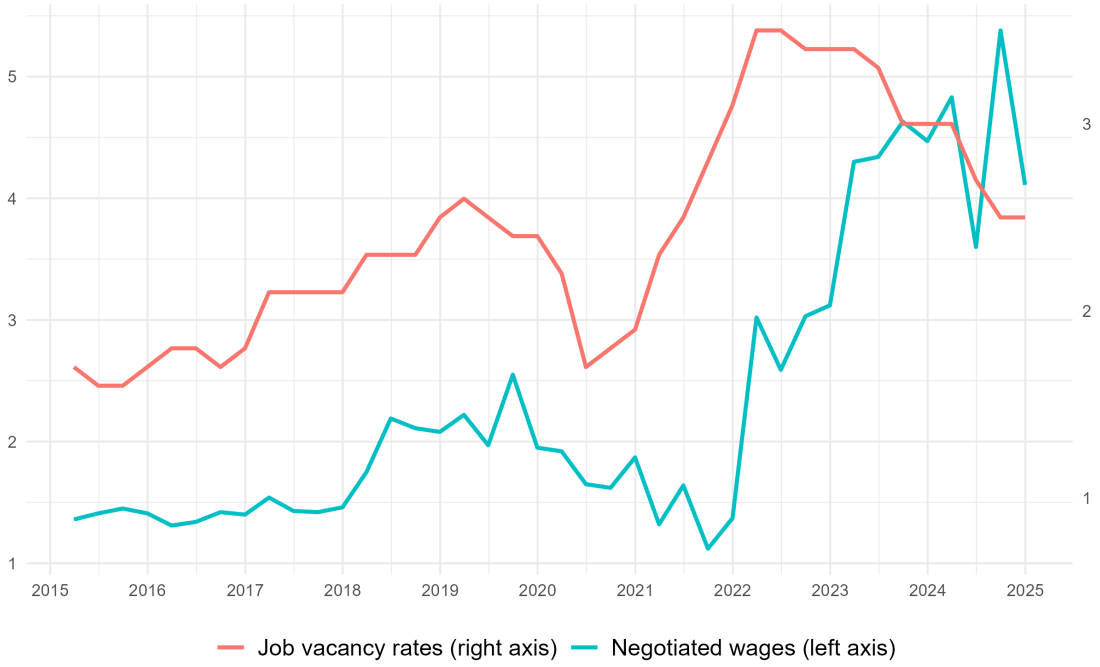
§Department of Econometrics, Vrije Universiteit Amsterdam; Tinbergen Institute, The Netherlands. E-mail: s.j.koopman@vu.nl.

¶Dutch National Bank (DNB), The Netherlands. E-mail: r.a.swierstra@dnb.nl

1 Introduction

Labor markets have been exceptionally tight in the euro area in the aftermath of the COVID-19 pandemic. After a period of relatively low wage growth in the euro area, these tight labor markets have gone hand in hand with sharp increases in nominal wages. Figure 1 shows how negotiated wages rose sharply after vacancy rates increased in the euro area. However, the empirical relationship between the two is less clear-cut, as wage growth only accelerated after vacancy rates reached historically unprecedented levels. This suggests that there is a nonlinear relationship between tightness and wages, where wages only increase when labor markets are extraordinarily tight.

Figure 1: Job vacancy rates and negotiated wages in the euro area, 2015-2025Q3



Source: European Central Bank

The Netherlands, the euro area's fifth-largest economy, offers a stark case study: both tightness and wages soared post-pandemic, but this coincided with an unprecedented surge in inflation following Russia's invasion of Ukraine and the subsequent energy price shock. This sharp increase in inflation makes it difficult to assess the relationship between tightness and wages using macroeconomic data, as inflation expectations are notoriously hard to measure (Mavroeidis, Plagborg-Møller and Stock, 2014). The recent Phillips curve literature has proposed a solution for this measurement problem: using disaggregated data and panel data methods makes it possible to control for inflation expectations and other confounding macroeconomic developments through the use of time fixed effects (McLeay and Tenreyro, 2020; Hazell, Herreño, Nakamura and Steinsson, 2022).

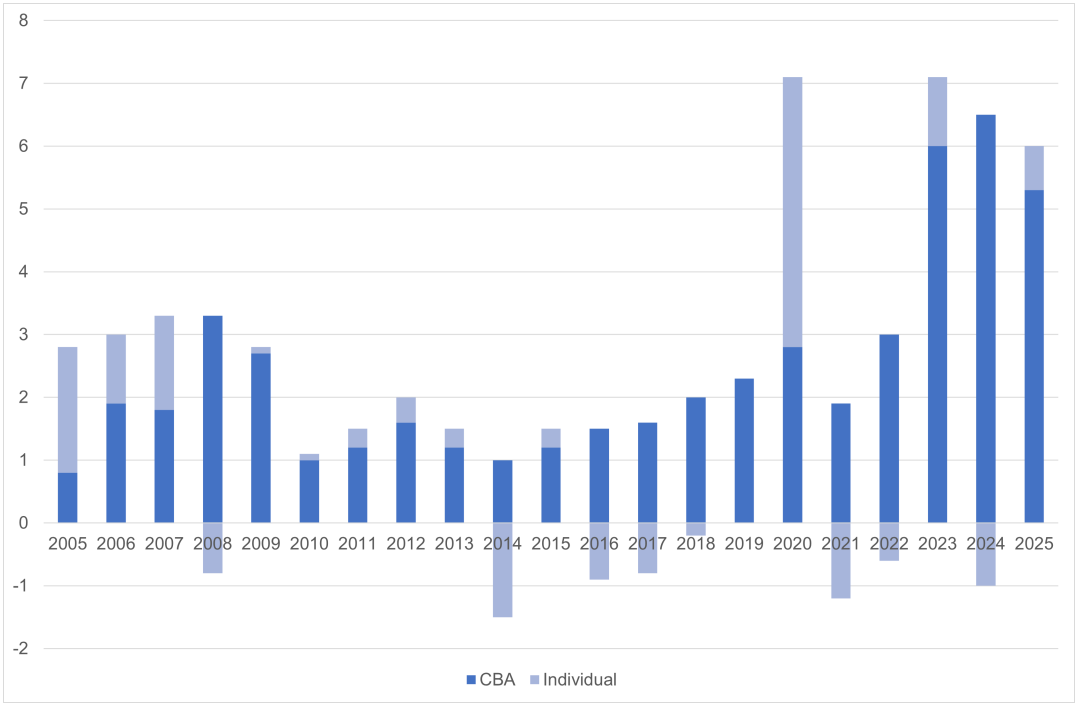
Building on these insights, we investigate the potentially time-varying relationship between labor market tightness and wage growth using panel data. We do so by estimating a dynamic panel data model of collective bargaining agreements (CBA) in the Netherlands from 2005 to 2025Q2. We relate CBA wage growth outcomes to changes in the vacancy rate at the sectoral level. We estimate our model using state space methods and estimate the time-varying effect of sectoral labor market tightness on CBA wage growth. In addition, we control for macroeconomic developments by estimating a time-varying intercept. Essentially, our model is a wage Phillips curve for a panel of CBAs, where inflation and inflation expectations are absorbed by the time-varying intercept.

We focus on CBA wage growth rather than total wage growth, as total wage growth also includes wage drift: increases in wages that result from individual negotiations between employees and employers. Our rationale is threefold. First, the determinants of CBA wages differ markedly from those of wage drift (Bhuller, Moene, Mogstad and Vestad, 2022); consequently, their underlying data-generating processes are distinct. The collective bargaining process is staggered and centralized, and involves professional negotiators who align their positions with the latest macroeconomic forecasts and data. Conversely, individual employees may renegotiate salaries at any time and often lack the specialized resources to interpret economic data at a professional level. Second, Figure 2 demonstrates that at the aggregate level, CBA wage growth accounts for the largest share of year-on-year hourly wage increases in the Netherlands during our sample period.[1] Third, by focusing on CBA wage growth, we gain a deeper understanding of how CBAs shape wage outcomes, an area of study which has been neglected in favor of models of individual wage setting. Given that most OECD economies are characterized by institutionalized collective bargaining, investigating the relation between CBA wage growth and labor markets potentially improves our understanding of wage growth dynamics in these countries. Since CBA wages are also used by the European Central Bank to monitor and forecast inflationary pressures, a better understanding of CBA wage developments can also aid the conduct of monetary policy (Górnicka and Koester, 2024).

To the best of our knowledge, we are the first to examine the stability of the parameter on labor market tightness by estimating a dynamic panel data model with time-varying parameters. For this purpose, we adopt state space methods, as these are well-suited to assess time variation in coefficients. State space methods estimate parameters as weighted local averages, for which the time window and the observation weights are determined as part of the model-based analysis and estimation. This integrated approach gives state space methods a distinctive advantage over alternative approaches such as rolling regressions, for which the researcher has to predetermine the window length and associated weights. While state space methods such as the Kalman filter are widely used in time series analysis, their

[1]The lockdown year of 2020 remains a notable exception, primarily due to temporary COVID-19 measures that subsidized firms for maintaining salaries despite reduced working hours.

Figure 2: Average hourly wage growth decomposed into CBA and wage drift growth



Source: CPB Netherlands Bureau for Economic Policy Analysis (2026)

application to panel data models has been limited, see van de Werve, Koopman, Weerman and Blokland (2026). Applying them to a large, unbalanced panel is a further contribution of this paper.

Since we estimate our model using panel data, we can use time fixed effects to capture the effect of variations in expected inflation and other macroeconomic variables on wage growth. Furthermore, using sectoral data allows us to use variation in labor market tightness at the sector level to better identify the effects of changes in tightness on wage growth. However, estimating models using panel data does come with a downside: while including entity and time fixed effects helps control for many types of bias, it also means that the effect of the independent variable on the dependent variable can only be identified through cross-sectional deviations from common fluctuations in labor market tightness. This makes it difficult to identify a relationship between two variables if the key independent variable has little cross-sectional variation and the relevant dynamic variation is absorbed by the time fixed effect. As a result, models estimated on *sectoral* data, like ours, are potentially uninformative about *aggregate* effects. This is called the missing intercept problem.

We solve the missing intercept problem in a novel way: we propose a method to decompose the time fixed effect. This is made possible by the state space framework, in which the fixed effect is estimated as a time-varying intercept. We subsequently decompose this intercept. We do so by using principal component analysis to obtain components or

factors that capture the main dynamics in business cycles, labor markets, productivity, inflation, economic sentiment, and union activity in the Netherlands and add these to our model. In addition, we include data pertaining to the Dutch institutional setting of staggered collective wage bargaining. We use these principal components to reduce the variation over time of the time-varying intercept. Its reduction to a constant would signify that our covariates have fully captured the common shocks and aggregate trends previously latent in the fixed effect. A successful decomposition hence effectively shifts the explanatory power from the unobserved intercept to observable variables. This approach allows us to leverage both the cross-sectional and the aggregate dimensions of our data, such that we can simultaneously estimate the effects of both cross-sectional variation and aggregate fluctuations in labor market tightness. Our approach therefore effectively integrates panel data techniques with time series analysis.

Our analysis yields three key findings. First, we find that cross-sectional variation in sectoral labor market tightness explains little of the variation in CBA wage growth: over the entire sample, the relationship between sectoral labor market tightness and CBA wage growth is not statistically significant. We find similar results using standard panel OLS techniques. Our results suggest that sector-specific labor demand conditions play only a limited role in central wage bargaining outcomes and that this role does not become larger when labor markets become tighter.

Second, we find that the time-varying intercept strongly relates to wage growth, suggesting that macroeconomic developments play an important role in determining CBA wages. Through our decomposition of the time fixed effect, we find that incorporating macroeconomic variables substantially reduces the volatility of the time-varying intercept in our state space model. These results indicate that macroeconomic variables explain a substantial share of common wage dynamics. Hence, macroeconomic conditions, rather than sectoral idiosyncrasies, are the primary drivers of centrally negotiated wages in the Netherlands.

Third, we show that aggregate labor market tightness, measured via the principal components that relate most strongly to vacancies and employment, has a significant and time-varying effect on CBA wage growth. This effect becomes particularly pronounced in the post-2021 period, consistent with a nonlinear wage Phillips curve that steepens during episodes of extreme tightness.

Taken together, our results suggest that the post-pandemic increase in CBA wages in the Netherlands was not caused by labor shortages in some sectors relative to others. This suggests that insufficient sectoral reallocation of labor, in the form of relative labor shortages across sectors, is unlikely to have been an important source of any inflationary effects of faster CBA wage growth. Taken a step further, our results also suggest that collective bargaining institutions are not flexible enough to facilitate adjustments in the labor market after sectoral shocks, since we find that changes in labor demand in

a single sector do not significantly influence wage growth in that sector. However, a macroeconomic shock that does not have a large impact on labor demand in a single sector, but does influence labor demand in the aggregate, does lead to higher wage growth in that single sector. As a consequence, collective bargaining that predominantly focuses on macroeconomic instead of sectoral developments can cause wages in some sectors to grow faster than warranted by fundamentals, while in others the reverse might be true.

Literature review. We follow the recent literature that uses panel data to estimate the slope of the Phillips curve. Mavroeidis et al. (2014) argue that the lack of variation in aggregate time series data precludes accurately identifying the separate effect of unemployment and expected inflation on current inflation. Furthermore, McLeay and Tenreyro (2020) assert that Phillips curve estimates will have the "wrong" sign when a central bank can successfully offset demand shocks, as the residual movements in inflation will be caused by supply shocks, which induce positive co-movement between unemployment and inflation. McLeay and Tenreyro (2020) and Hazell et al. (2022) tackle this issue by leveraging regional data for the United States: the national central bank does not offset the effects of regional demand shocks using its national monetary policy instrument.[2] Hazell et al. (2022) show that in such a setting, the monetary policy response to aggregate shocks and long-run inflation expectations are absorbed by time fixed effects. In line with this literature, we use panel data in the form of sectoral data. Furthermore, by applying state space methods to panel data, we aim to solve both the identification and stability issues highlighted by Mavroeidis et al. (2014), McLeay and Tenreyro (2020) and Hazell et al. (2022) in a novel way.

Other work has also proposed solutions to the missing intercept problem. Matthes, Nagasaka and Schwartzman (2025) tackle the missing intercept problem by combining microeconometric methods at the level of individual units with time series methods to estimate effects at the aggregate level. Related but different is Wolf (2023), who uses a structural model to identify the conditions under which local effects and macroeconomic estimates of general equilibrium effects can be summed up to identify the total effect of a shock. Hazell et al. (2022) tackle the missing intercept problem in a different way. Using a structural model of a monetary union with two regions with tradable and non-tradable sectors, Hazell et al. (2022) show that the slope of the regional Phillips curve for inflation in the non-tradable sector and the slope of the aggregate Phillips curve are identical. Hence, the local and aggregate coefficients are the same in their setup. We tackle the missing intercept problem by estimating the time fixed effect as a time-varying intercept, and subsequently decomposing the intercept using a large number of time series. In addition, our use of state space methods allows us to estimate time-varying relationships at both

[2]Other papers that leverage regional data to estimate the Phillips curve are Beraja, Hurst and Ospina (2019), Hooper, Mishkin and Sufi (2020), and Fitzgerald, Jones, Kulish and Nicolini (2024).

the sectoral and the aggregate level.

Conceptually, our paper is related to the literature on the wage Phillips curve. Galí (2011) shows that introducing staggered wage setting frictions in the vein of Calvo (1983) following Erceg, Henderson and Levin (2000) and Schmitt-Grohé and Uribe (2005) in a New Keynesian business cycle model effectively introduces a forward-looking wage Phillips curve. This New Keynesian wage Phillips curve relates current wage growth to deviations of unemployment from its natural rate and expected inflation. Galí (2011) estimates a New Keynesian wage Phillips curve with adaptive inflation expectations for the US, and finds a strong, statistically significant negative relationship between unemployment and nominal wage growth.

The interaction between the state of the labor market and wage growth has also been studied in models of labor markets with search frictions in the Diamond-Mortensen-Pissarides tradition, e.g. Mortensen and Pissarides (1994). Wages in these models are typically determined through bilateral period-by-period Nash bargaining, while Gertler, Sala and Trigari (2008) and Gertler and Trigari (2009) introduce staggered wage bargaining in the vein of Calvo (1983). In these models, labor market tightness, defined as the ratio between vacancies and job seekers, is an important determinant of wage dynamics.

Our work is also related to the literature on the time-varying relationship between labor market tightness and the steepness of the Phillips curve. Benigno and Eggertsson (2023) estimate a reduced-form Phillips curve for the US using aggregate data, and show that inflation accelerates when labor market tightness exceeds a threshold. Benigno and Eggertsson (2023) rationalize this relationship in a New Keynesian model where wages increase with expected inflation when the labor market is tight. Gitti (2024) builds on this work by estimating a New Keynesian Phillips curve on regional US data, where the Phillips curve becomes steeper as labor market tightness increases. We directly estimate the time-varying slope in a wage Phillips curve, without imposing a threshold level above which wages respond more strongly to changes in labor market tightness. Finally, Bonam, de Haan and Van Limbergen (2021) estimate a time-varying parameter wage Phillips curve for the five largest euro area countries using aggregate data and a survey-based labor market slack indicator, and find that the slope of the wage Phillips curve is indeed time-varying. We differ from their setup by using panel data for a single country, which allows us to address the issues with time series described above.

The remainder of the paper is structured as follows: Section 2 provides additional detail on the collective wage bargaining institutions in the Netherlands, Section 3 discusses the main models and the state space methodology, Section 4 describes the data, and Section 5 discusses the results. Section 6 concludes.

2 Collective wage bargaining in the Netherlands

To inform our empirical specification, we rely on descriptions of the collective wage bargaining process in van der Velden (2016), Visser (2021), Caloia, Parlevliet and Mastrogiacomo (2023) and Ministry of Social Affairs and Employment (2025), as well as interviews we conducted with representatives of the largest trade union federation (FNV), the employers' organization responsible for coordinating the collective wage bargaining process (AWVN), and a CBA expert at the Dutch Ministry of Social Affairs. Summaries of these interviews can be found in Appendix A. Below we summarize the key takeaways.

As is the case for most OECD countries, annual wage growth for the majority of employees in the Netherlands is determined through collective bargaining between employers and labor unions. The collective bargaining system in the Netherlands is most similar to those of Denmark, Norway, and Sweden, and, to a lesser extent, to those of Iceland, Germany, Austria, Finland, Japan, and Belgium (Bhuller et al., 2022). In 2024, 72.5% of employees were covered by a CBA.[3] This is comparable to other countries in continental Europe and Scandinavia (Bhuller et al., 2022). Around 65% of employees are covered by a CBA that holds for an entire sector (e.g. construction or hairdressing), while 7.5% of workers are covered by a firm-specific CBA (e.g. the CBA for Heineken). Wages for the remaining 27.5% are negotiated solely between individual employees and employers. In turn, 61.2% of workers covered by a CBA are employed in the private sector.

The Netherlands has multiple unions, of which FNV is by far the largest.[4] CBA negotiations involve negotiators from at least one union, employer representatives, and the employers' organization AWVN.

The wage bargaining season starts in September with the publication of the centrally demanded wage growth percentage (CDW) for the following year by the FNV. This typically occurs one week before the Dutch government publishes its budget and policy plans, which happens on the third Tuesday of September every year. The CDW serves as the FNV's primary wage growth target. Figure 3 shows the dynamics of the CDW over the sample period.[5] Historically, the CDW has been set to at least match CPB Netherlands Bureau for Economic Policy Analysis's (CPB, the Dutch independent fiscal institution) August forecast of CPI inflation at constant tax rates to ensure constant purchasing power. To further calibrate this demand, FNV evaluates the so-called wage margin: how much wages can reasonably grow given macroeconomic conditions, which depends on productivity growth and the GDP deflator. Interview evidence suggests that the CDW is also heavily influenced by the broader business cycle and labor market conditions, with

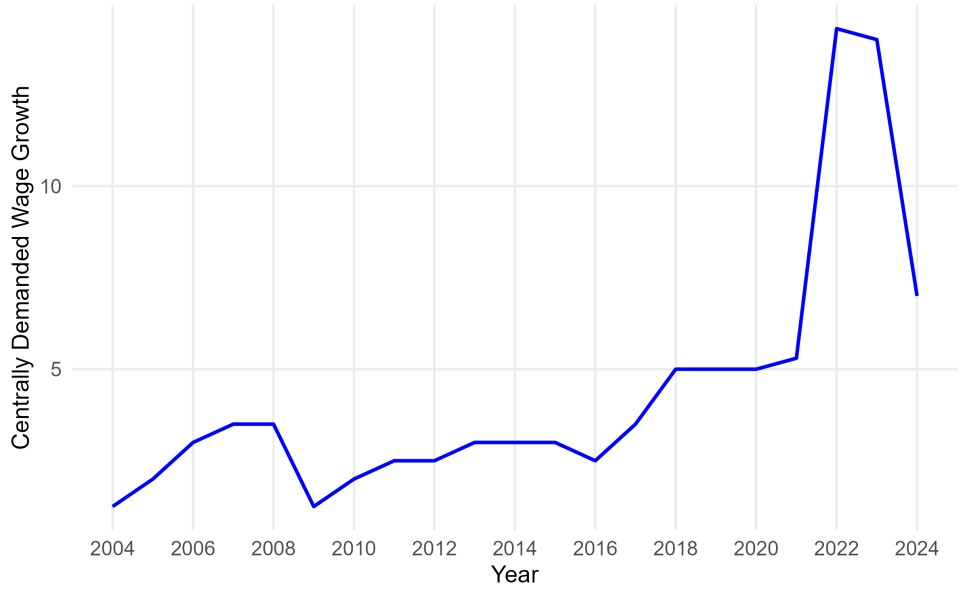
[3] This is down from 76.7% in 2010, see Statistics Netherlands (2025).

[4] Historically, approximately 60% of all union members have been members of FNV (Statistics Netherlands, 2026).

[5] The CDW announced in 2023 for 2024 was a range (5-14%), rather than a single number, to differentiate between CBAs that had and had not been already compensated for the inflation surge of 2022. In Figure 3 we show the maximum of that range.

economic booms and aggregate labor market tightness exerting upward pressure on the final CDW. While this claim was traditionally translated into sector-specific demands, in recent years it has evolved into a uniform national wage claim that applies across all industries.

Figure 3: The centrally demanded wage growth percentage of the largest union, FNV, announced in 2004-2024



Note: The graph shows the centrally demanded wage growth percentage published in year t for the year $t + 1$.

FNV prioritizes the ordering of CBA negotiations, starting early in the year with CBAs that are deemed the most likely to yield sizeable wage growth. These CBAs serve as a benchmark or set an example for subsequent negotiations in other sectors where union bargaining power might be lower. CBAs are negotiated on a staggered basis: the duration of CBA contracts differs between CBAs and over time, and the majority of CBA contracts are finalized retroactively, meaning that the start date of the contract precedes the agreement date (see Section 4.1 for more details).

In our analysis we focus on collective bargaining outcomes in the private sector. By law, maximum wage growth in the public sector (health care, public administration, and education) is limited by the projected average wage growth in the private sector.[6] To avoid adding further complexities to our model to accommodate this, we decided to leave an analysis of public sector CBA wage growth to future research.

The qualitative evidence we gathered from the literature as well as our interviews suggests that a Phillips curve-type model, which captures the relationship between labor market slack, inflation expectations and nominal wage growth, is particularly suited to

[6]See Ministry of the Interior and Kingdom Relations (2017).

explain outcomes of Dutch CBA wage setting institutions. Besides the prominent role of inflation forecasts in informing the CDW, the experts we interviewed acknowledged that actual CBA outcomes remain heavily contingent on the business cycle.

Furthermore, although CBA negotiations in the Netherlands feature both centralized and decentralized characteristics, from our interviews we gathered that the bargaining process is institutionalized in a way that strongly emphasizes macroeconomic data. This centralization is reinforced by FNV's strategy to negotiate high-profile agreements first to set a benchmark for the remaining CBAs. Furthermore, negotiators on both sides rely heavily on (CPB's) macroeconomic forecasts to establish a common macroeconomic reference point, largely because detailed sectoral forecasts are often unavailable or lack the same institutional authority. Consequently, the interview evidence suggests that *aggregate* labor market tightness and inflation expectations, rather than *sectoral* idiosyncrasies, are the primary drivers of negotiated outcomes.

3 Empirical models and methodology

In our empirical analysis, we adopt three different panel data models. First, we consider the textbook two-way fixed effects model. Second, we keep the standard model, but we replace the time fixed effect with a time-varying intercept and introduce time-varying regression coefficients. Finally, we include macroeconomic variables in the model with the aim of reducing the volatility, and hence the importance, of the time-varying intercept. Each of these three specifications is described in detail below.

3.1 The two-way fixed effects panel data model

Our dependent variable of interest is $\Delta w^p_{i,s,t}$, defined as a time-weighted average negotiated wage growth rate, annualized over the total duration of the contract period p, for CBA i, in sector s, and agreed during quarter t (the time index). Hence, a contract can encompass multiple years and hence multiple consecutive wage increases. A unique observation of an individual CBA contract can be identified either by the CBA-agreement date pair (i, t) or by the CBA-period pair (i, p), as contract periods for a given unit i do not overlap. Note that a sector s encompasses multiple CBAs. Our sample contains all the CBAs for which we observe at least two CBA contracts, see Section 4.1 for more details on the wage data structure.

We are primarily interested in how $\Delta w^p_{i,s,t}$ relates to $\Delta\theta_{s,t-2}$, which denotes the year-on-year change in sectoral labor market tightness of two quarters ago. We consider the standard fixed effects panel data model for quarterly data as given by

$$\Delta w^p_{i,s,t} = \delta_i + \tau_t + \beta_1 \Delta\theta_{s,t-2} + \varepsilon_{i,s,t}, \qquad \varepsilon_{i,s,t} \sim NID(0, \sigma^2_\varepsilon),$$

where δ_i is the CBA fixed effect, τ_t is the year-quarter time fixed effect, β_1 is the regression coefficient for the explanatory variable $\Delta\theta_{s,t-2}$, and $\varepsilon_{i,s,t}$ is the error which is assumed to be normally, independently, and identically distributed (NID) with mean zero and variance $\sigma_\varepsilon^2 > 0$.

Following the search and matching literature, we specify the model in differences rather than levels.[7] Hence, wage growth should be related to changes in labor market tightness.[8] Furthermore, we use year-on-year changes rather than annualized quarterly growth rates, as this aligns with the institutional setting that we have described in Section 2, where unions and employers typically benchmark their demands against annual trends in macroeconomic data. Our primary parameter of interest is β_1, which captures the impact of the year-on-year change in sectoral labor market tightness, $\Delta\theta_{s,t-2}$. While we observe average wage growth at the individual CBA level, vacancy data are only available at the sectoral level. This creates a nested panel structure. We employ CBA-level fixed effects rather than sectoral fixed effects, as the former capture all time-invariant structural differences at both the sectoral and the individual agreement level. We employ a two-quarter lag for our baseline specification to reflect the information set available to negotiators at time t. This choice is motivated by the fact that data from Statistics Netherlands for a given quarter typically become available one to two quarters later. We are hence interested in relating *recent* labor market tightness to the *current* average wage agreement made at time t, which usually consists of both retroactive wage growth and future wage growth. We describe our considerations regarding the operationalization of $\Delta\theta_{s,t-2}$ in more detail in Section 4.1.

Since we are primarily interested in the time-varying intercept and the regression parameters, we avoid explicitly estimating the individual fixed CBA effects δ_i by applying the within-transformation. We subtract the CBA-specific means over time from our dependent and independent variables. Hence, we obtain the simplified and parsimonious panel data model:

$$\Delta\hat{w}^p_{i,s,t} = \tau_t + \beta_1 \Delta\hat{\theta}_{s,t-2} + \varepsilon_{i,s,t}, \qquad \varepsilon_{i,s,t} \sim NID(0, \sigma_\varepsilon^2), \tag{1}$$

where $\Delta\hat{w}^p_{i,s,t} \coloneqq \Delta w^p_{i,s,t} - \mu^{\Delta w}_{i,s}$, and $\Delta\hat{\theta}_{s,t-2} \coloneqq \Delta\theta_{s,t-2} - \mu^{\Delta\theta}_{s}$. Here, $\mu^{\Delta w}_{i,s}$ and $\mu^{\Delta\theta}_{s}$ denote the entity-specific means for each CBA unit i in sector s.[9]

[7] When wages are determined through Nash bargaining in search and matching models, the level of the wage rate depends on the level of labor market tightness. Hence, the growth rate of wages depends on the growth rate of labor market tightness.

[8] In the Phillips curve literature, it is more common to relate wage inflation to the level of labor market tightness. As a robustness check, we estimated the model using levels, yielding broadly comparable results (see section 5.4).

[9] We demean the variables at the most disaggregated level possible, which in the case of our measure for labor market tightness is the sectoral level.

3.2 Introducing time-varying parameters

To decompose the time fixed effect, a procedure detailed in the next section, we replace it by a time-varying intercept τ_t. We specify τ_t as a stochastic variable that evolves over time using an additional sequence of NID errors or innovations as input. In particular, we adopt a random walk process to obtain a stochastically time-varying effect τ_t given by

$$\tau_t = \tau_{t-1} + \eta_{\tau,t}, \qquad \eta_{\tau,t} \sim NID(0, \sigma^2_\tau), \tag{2}$$

where we assume that the sequences $\{\varepsilon_{i,s,t}\}$ and $\{\eta_{\tau,t}\}$ are independent of each other for all i, s. The resulting model (1)–(2) constitutes a multivariate (panel) linear Gaussian state space model with a large set of measurement equations (1) and a single updating (or state) equation (2); see Durbin and Koopman (2012, Chapter 3) for the general discussion of state space representations and Jungbacker and Koopman (2015) for a possible treatment of panel state space models.

An additional level of flexibility is introduced by replacing the time-invariant regression coefficient β_1 by the time-varying parameter $\beta_{1,t}$ that we model in a way similar to the time-varying intercept τ_t. We obtain

$$\Delta\hat{w}^p_{i,s,t} = \tau_t + \beta_{1,t}\Delta\hat{\theta}_{s,t-2} + \varepsilon_{i,s,t}, \qquad \varepsilon_{i,s,t} \sim NID(0, \sigma^2_\varepsilon), \tag{3}$$

where $\beta_{1,t}$ follows a random walk:

$$\beta_{1,t} = \beta_{1,t-1} + \eta_{\beta_1,t}, \qquad \eta_{\beta_1,t} \sim NID(0, \sigma^2_{\beta_1}), \tag{4}$$

where we assume that all paired sequences $\{\varepsilon_{i,s,t}\}$, $\{\eta_{\tau,t}\}$ and $\{\eta_{\beta_1,t}\}$ are independent of each other for all i, s. The resulting model (1)–(4) still has a multivariate (panel) linear Gaussian state space representation, but now with two random walk updating equations (2) and (4).

Under what conditions can we interpret the estimate of $\beta_{1,t}$ as causal? Lacking an appropriate instrument for this setting, we list the conditions and assumptions under which we can make a case for a causal interpretation. First, specific features of the collective wage bargaining setting support the view that our estimate of $\beta_{1,t}$ is not likely to be affected by reverse causality or simultaneity bias: since labor market tightness information is *predetermined* (i.e. generated before the CBA is concluded at time t) it is not affected by $w^p_{i,s,t}$. Second, the within-CBA transformation and the time-varying intercept τ_t inoculate our estimates against omitted variable bias emanating from structural differences between CBAs and common fluctuations, respectively, as well as the selection bias that occurs as a result of the non-random start of bargaining rounds for different CBAs.

A more subtle concern is that $\beta_{1,t}$ may be biased by serial correlation in $\Delta\hat{w}^p_{i,s,t}$. On the

one hand, past CBA wage growth raises wage costs, which can depress labor demand and hence lead to lower sectoral tightness today. If wage growth is persistent, $\Delta\hat{\theta}_{s,t-2}$ becomes negatively correlated with the error term and $\hat{\beta}_{1,t}$ is biased downward. On the other hand, $\hat{\beta}_{1,t}$ is biased upward if current wage growth is instead negatively associated with past wage growth, as would be the case if current settlements compensate for perceived insufficient wage growth in the preceding contract. Two features of our setting mitigate both concerns. First, as discussed in Section 2, CBA wage setting in our institutional context is predominantly forward-looking, with claims anchored to macroeconomic forecasts rather than to past realized wage growth, which limits the persistence available to generate either form of bias in the first place. This is reinforced by the timing of our data: because individual CBA contracts often span a year or more (see section 4.1) and successive agreement dates for a given CBA are typically separated by irregular, multi-year gaps, whatever conditions shaped one settlement have considerable time to dissipate before the next round of forward-looking bargaining begins. Second, we explicitly account for the main backward-looking channel, which is catch-up for inflation that exceeded expectations under the previous contract, by including a variable that measures inflation shortfall (see Section 4.3.3). Because this variable is common across CBAs facing the same realized inflation surprise, it captures precisely the persistent catch-up dynamic that would otherwise bias $\beta_{1,t}$.

Finally, CBA-specific shocks that are neither absorbed by CBA nor time fixed effects can also be a potential source of omitted-variable bias. In particular, the literature points to productivity shocks as an important source of wage differentials and hence bias in our models. Therefore, as a robustness check, we control for sectoral productivity growth, $\Delta\hat{z}_{s,t}$. Furthermore, given the 2022 energy price shock in the Netherlands, the latter part of the sample is potentially sensitive to sector-specific energy price inflation due to differences in energy intensities between sectors. Therefore, as a robustness check, we include $\pi^{e}_{s,t-4}$, which measures the change in energy costs as a share of total costs per sector. Appendix B.4 describes these variables in more detail and section 5.4 discusses the empirical results.

3.3 Introducing macroeconomic variables

The time-varying intercept τ_t in equation (1) absorbs all common shocks, including aggregate fluctuations in labor market tightness. In a standard panel data model, such aggregate fluctuations are either differenced out or captured by time dummies, which in our setting are perfectly collinear with aggregate macroeconomic variables, so their separate effects cannot be estimated. As a consequence, only the idiosyncratic cross-sectional deviations from common fluctuations in tightness remain to inform $\beta_{1,t}$. Hence, while $\beta_{1,t}$ captures cross-sectional differences in labor market tightness, it is not informative about the impact of common changes in labor market tightness over time. Depending on how

prominent aggregate trends are over the sample, this construction may understate the full effect of labor market tightness on nominal CBA wage growth. The issue is a manifestation of the missing intercept problem in empirical macroeconomics: panel regressions with time fixed effects identify only *relative* effects, not aggregate ones.

To enable the estimation of the aggregate effects of labor market tightness on wage growth, a possible approach would be to include its aggregate counterpart. Hence, we consider the model

$$\Delta \hat{w}^{p}_{i,s,t} = \tau_t + \beta_{1,t}\Delta\hat{\theta}_{s,t-2} + \beta_{2,t}\Delta\theta_{t-2} + \varepsilon_{i,s,t}, \tag{5}$$

where $\Delta\theta_{t-2}$ denotes a measure of aggregate labor market tightness at time $t-2$, which varies over time but not across CBAs, and $\beta_{2,t}$ also follows a random walk as in (4). Using this approach, aggregate fluctuations in labor market tightness that would otherwise be collinear with τ_t are no longer absorbed deterministically, but instead attributed to either τ_t or $\beta_{2,t}\Delta\theta_{t-2}$ based on their respective contributions to the dynamic structure of the data. As a result, we can recover estimates of the aggregate relationship between labor market tightness and wages that are not available in a traditional fixed effects framework. Hence, this approach enables the simultaneous estimation of the effects of both cross-sectional variation in tightness, captured by $\Delta\hat{\theta}_{s,t-2}$, and aggregate fluctuations in labor market tightness over time, captured by $\Delta\theta_{t-2}$.

In the extended model (5), τ_t becomes a measure of those aggregate dynamics in CBA wage growth that are *not* explained by observed labor market tightness. Hence, if both $\Delta\hat{\theta}_{s,t-2}$ and $\Delta\theta_{t-2}$ equal 0, the conditional expectation of $\Delta\hat{w}^{p}_{i,s,t}$ equals τ_t. Therefore, the volatility of τ_t can be interpreted as a measure of unexplained aggregate trends. We can leverage this feature of the model to further decompose the time fixed effect τ_t by systematically incorporating additional macroeconomic variables: if including relevant macroeconomic controls explains more of the variation in τ_t, then its estimated volatility should decrease. In the limiting case where τ_t converges to a constant, this would suggest that all systematic aggregate variation in CBA wage growth is accounted for by the included observables.

We therefore estimate:

$$\Delta \hat{w}^{p}_{i,s,t} = \tau_t + \beta_{1,t}\Delta\hat{\theta}_{s,t-2} + \beta_{2,t}\Delta\theta_{t-2} + \boldsymbol{\gamma}\boldsymbol{M}_{t-2} + \varepsilon_{i,s,t}, \tag{6}$$

where $\boldsymbol{M}_{t-2}$ is an $m \times 1$ vector containing macroeconomic variables and the $1 \times m$ vector $\boldsymbol{\gamma}$ contains the corresponding coefficients. We discuss the macroeconomic variables that enter $\boldsymbol{M}_{t-2}$ in Section 4.3.

Since we do not instrument $\Delta\theta_{t-2}$, under what conditions can $\beta_{2,t}$ be interpreted causally? Reducing τ_t to a constant using lagged macroeconomic data shows only that the included covariates predict common fluctuations in CBA wage growth, and predictability

is not sufficient for causation. Two features of our setting nonetheless support a causal reading. First, as with $\beta_{1,t}$, aggregate tightness enters with a two-quarter lag and is therefore predetermined: it is realized before the agreement is concluded at time t and cannot be a consequence of $\Delta\hat{w}^p_{i,s,t}$, which rules out reverse causality. Second, we can invoke a narrative identification argument. Our interviews indicate which macroeconomic information matters most to negotiators, which renders it plausible that the variables entering $\boldsymbol{M}_{t-2}$ approximate the information set on which they base their wage claims. Under the assumption that $\boldsymbol{M}_{t-2}$ captures this information set fully, so that $\Delta\theta_{t-2}$ is conditionally unrelated to the remaining determinants of wage growth, $\beta_{2,t}$ admits a causal interpretation: conditional on the other macroeconomic variables, a one-unit increase in $\Delta\theta_{t-2}$ raises $\Delta\hat{w}^p_{i,s,t}$ by $\beta_{2,t}$. This conditional-independence assumption is strong, and the interview evidence makes it plausible but does not guarantee it. We therefore also appeal to theory: estimates consistent with the predicted relationship between labor market tightness and wage growth lend further, if indirect, support to a causal reading.

Only sectoral and aggregate labor market tightness enter the model with time-varying coefficients, since we are mainly interested in the potentially nonlinear relationship between wage growth and labor market tightness. For the remaining macroeconomic variables we assume that the parameters in $\boldsymbol{\gamma}$ are fixed in order to avoid overparameterizing the model, which would reduce statistical power.

We make one notable exception for inflation-related variables. Following the COVID-19 pandemic, inflation reached an unprecedented level in 2022. This surge was initially driven by post-pandemic supply frictions and amplified by the energy price shock following Russia's invasion of Ukraine in 2022. To allow for such a large structural change while keeping the model parsimonious, we interact the inflation-related variables with time dummies, see section 5.

3.4 Estimation methods

We estimate the standard two-way fixed effect (TWFE) panel data model using standard regression methods. We use the method implemented in `Stata`. When the panel data model is extended with time-varying parameters as in (3), (5) and (6), we represent the model in a linear Gaussian state space formulation. A more detailed discussion of treating time-varying parameters in the setting of a state space model using the Kalman filter is set out in Appendix C. Estimation of the time-varying intercept τ_t and the time-varying parameters $\beta_{i,t}$, for $i = 1, 2$, reduces to taking a locally weighted average of the data (corrected for the other effects of the model). The time window and observation weights are determined by the variances of the NID sequences in the model. These variances are estimated by the maximum likelihood method which requires numerical optimization of the log-likelihood function. The log-likelihood function is evaluated using the Kalman filter

which effectively carries out the prediction error decomposition of the model. We have implemented these methods in `Ox` using the `SsfPack` library (Doornik, 2021; Koopman, Shephard and Doornik, 1999).

4 Data and descriptive statistics

4.1 CBA wage growth data

We combine two datasets to measure average agreed upon wage growth for each CBA: one from Brightmine, an HR and compliance information provider, and one from employers' organization AWVN. In the Netherlands, CBAs are public records. After negotiations have ended, the full text of each agreement is published online. These documents provide the necessary details on the start and expiration date of each contract, which often spans multiple years, and the schedule of all agreed upon wage increases throughout that period. Both AWVN and Brightmine systematically codify this information and record the agreement date of each CBA. Both datasets also record the sector to which each CBA belongs following to the national industry standard classification system ('Standaard Bedrijfsindeling', abbreviated to SBI) set by Statistics Netherlands, which we need to link the CBA wage data to data on sectoral labor market tightness.[10]

While cleaning the data, we found that the Brightmine dataset exhibited fewer manual errors, fewer missing observations, and more detailed variables, leading us to select it as our primary source. However, Brightmine only started recording the dates of various negotiation phases from 2007Q3 onwards. We therefore append the AWVN agreement date variable to the Brightmine dataset. This merged dataset allows us to also use data on CBA wages during 2005-2007Q2, albeit for a very limited sample.[11] Rather than discard these observations, we use them to initialize the Kalman filter.[12]

Our final dataset on CBA wage growth spans the 2005-2025Q2 period and contains 7,182 observations for 949 unique CBAs in 16 sectors and for 2 CBA types. For each of the 949 CBAs, on average we observe 7 or 8 individual staggered contracts with consecutive starting and expiration dates. Table 1 shows the number of observations per sector and per CBA type.[13]

[10] At the first digit level, which is what we use in this paper, this classification system is identical to its international counterparts ISIC and NACE, see Statistics Netherlands (2024).

[11] In the early years of the AWVN sample, the CBA ID was often missing which prevents us from merging it with the Brightmine dataset.

[12] Since the Kalman filter is a recursive method over time, the initial phase of estimation relies on a low number of data points, leading to more uncertain estimates of the time-varying parameters (states). After several updating steps (learning) over time, when more observations become available, the state estimates become more accurate.

[13] As explained in Section 2, the CBAs for categories O (Government), P (Education) and Q (Health care) follow a different bargaining process from the other sectors. Therefore, we leave an analysis of these sectors for future research.

A substantial portion of our sample consists of CBAs within the manufacturing industry (Sector C). More broadly, the dataset is weighted toward goods-producing sectors (Sectors A–F, as detailed in Table 1). Furthermore, as industry-wide CBAs cover more employees, there are fewer of them in our dataset: the majority of CBAs (5,603 out of 7,182) are firm-specific. To ensure our findings are not driven exclusively by these features of the data, we perform several checks, see section 5.4.

Table 1: CBA contracts per sector and per CBA type

SBI	**Description**	**Industry**	**Firm**	**Total**	
		#	#	#	%
A	Agriculture, forestry, and fishing	100	44	144	2.0
B	Mining and quarrying	0	42	42	0.6
C	Manufacturing	419	3,102	3,521	49.0
D	Electricity, gas, steam, and air conditioning supply	18	53	71	1.0
E	Water supply; waste and wastewater management	32	30	62	0.9
F	Construction	104	39	143	2.0
G	Wholesale and retail trade; motor vehicle repair	325	425	750	10.4
H	Transportation and storage	76	833	909	12.7
I	Accommodation and food services	50	11	61	0.8
J	Information and communication	65	121	186	2.6
K	Financial institutions	44	242	286	4.0
L	Real estate activities	10	6	16	0.2
M	Professional, scientific, and technical services	46	310	356	5.0
N	Administrative and support services	97	137	234	3.3
R	Arts, entertainment, and recreation	118	68	186	2.6
S	Other service activities	75	140	215	3.0
Total		1,579	5,603	7,182	100.0

Notes: This table shows the number of observations for each sector (SBI) as well as CBA type. 'Industry' refers to industry-wide CBAs that cover more than one firm, such as bakeries or supermarkets. 'Firm' refers to CBAs that apply to only one firm.

Due to the frequency of our sectoral and macro labor market tightness measures, our analysis is at the quarterly level. Each quarter we observe on average 88 CBA contracts. Notable outliers are the 2005-2007Q2 period, which we keep in the sample for the Kalman filter initialization, and the first two quarters of 2020, which coincide with the first COVID-19 lockdown, when collective wage bargaining was temporarily suspended. Figure 5 also reveals distinct seasonal variation: the number of reached agreements peaks in the second quarter of each year. This is consistent with the seasonal pattern of collective bargaining discussed in Section 2.[14]

The data on CBA wage growth form an unbalanced panel. Our empirical models relate agreed upon average wage growth at time t to lagged measures of labor market

[14]This observation could have warranted adding a seasonal component to our empirical model in case the time-varying intercept τ_t or the time-varying parameters $\beta_{1,t}$ or $\beta_{2,t}$ exhibited a seasonal pattern, but we did not find a seasonal component.

tightness. Since CBAs often span multiple years and negotiations are staggered, we do not observe a new agreement for every unit i or even every sector s in each quarter t.[15] Fortunately, state space models estimated via the Kalman filter are particularly well-suited for unbalanced panels, because the filter can easily account for missing data points in the measurement equation. When an observation for a specific CBA is missing in a given quarter, the Kalman filter simply skips the update step for that unit while continuing to propagate the state.

Finally, our dependent variable, $\Delta w^p_{i,s,t}$, represents time-weighted average negotiated wage growth rate, annualized over the total duration of the contract. Because a contract period p often spans multiple years, it typically includes several consecutive wage increases. Conversely, a single calendar year may contain multiple intra-period increases. As Table 2 shows, the distribution is highly concentrated: durations of exactly 1 and 2 years are the most frequent, featuring an average of 1.3 and 2.3 wage increases, respectively. As such, simply using first-year negotiated increases would ignore that some CBAs have a longer duration and/or more scheduled increases in wages over the contract period. Our approach of constructing an indicator of time-weighted average CBA wage growth more accurately captures this heterogeneity between CBAs. In Appendix B.3 we show how we construct this indicator. The resulting variable, shown in Figure 4, shows how in 2022 CBA wages started to increase rapidly. This justifies our model setup with time-varying parameters, which can accommodate nonlinearities in the data.

[15] For our state space specification we reshape the dataset so that each CBA has its own row, and the columns represent all the quarters. Since we have 82 quarters in our dataset, and we on average observe in between 7 and 8 CBA contracts per CBA, each row contains more missing observations than actual observations. To illustrate this, Appendix B.1 gives an excerpt of the data for one firm and one industry-wide CBA in the original format and the state space format.

Figure 4: Time-weighted average annualized CBA wage growth, 2005-2025Q2

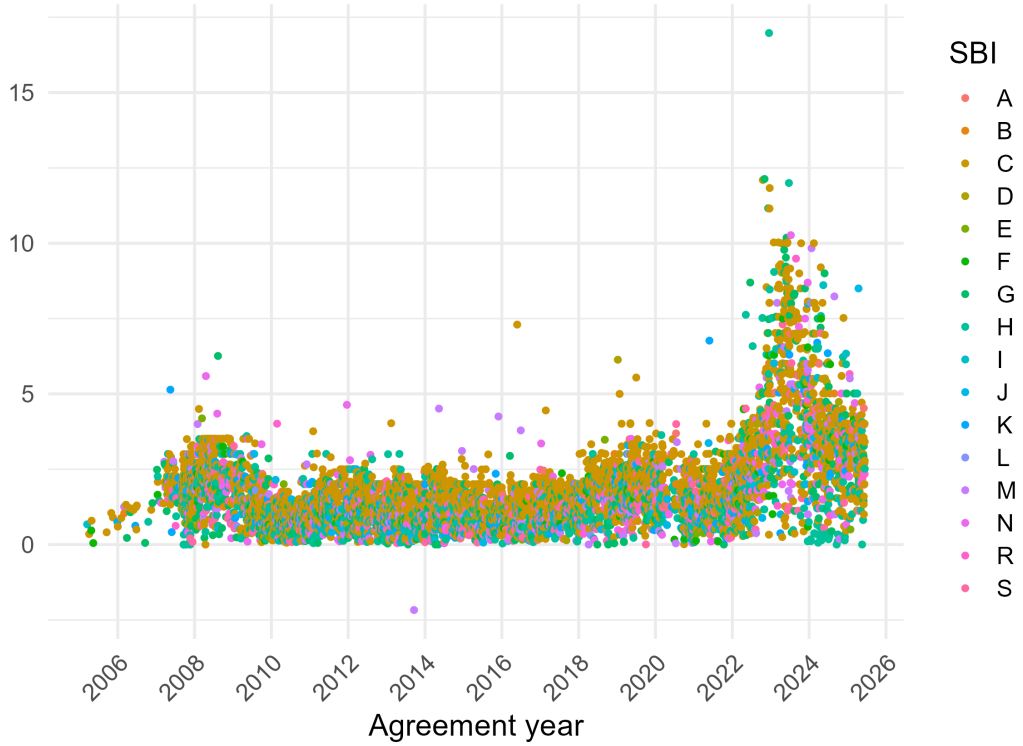


Figure 5: Total number of observations per quarter and sector, 2005-2025Q2

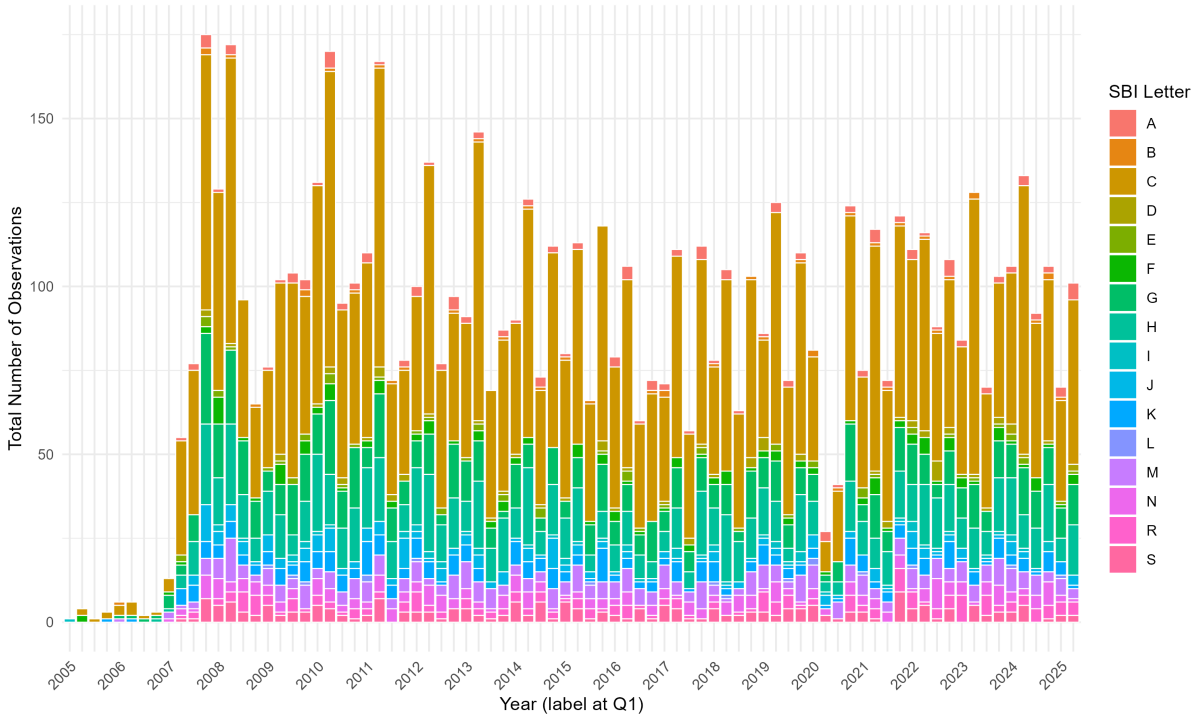


Table 2: Contract duration and frequency of increases

			Increases per contract	
Contract duration	n	%	Average	SD
Less than 1 year	210	2.9	1.1	0.4
Exactly 1 year	2,819	39.3	1.3	0.5
Between 1 and 2 years	1,329	18.5	1.8	0.7
Exactly 2 years	1,930	26.9	2.3	0.8
More than 2 years	894	12.4	2.8	1.2
Total	**7,182**	**100.0**	–	–

Notes: Percentages may not sum exactly to 100% due to rounding. *SD* denotes the standard deviation.

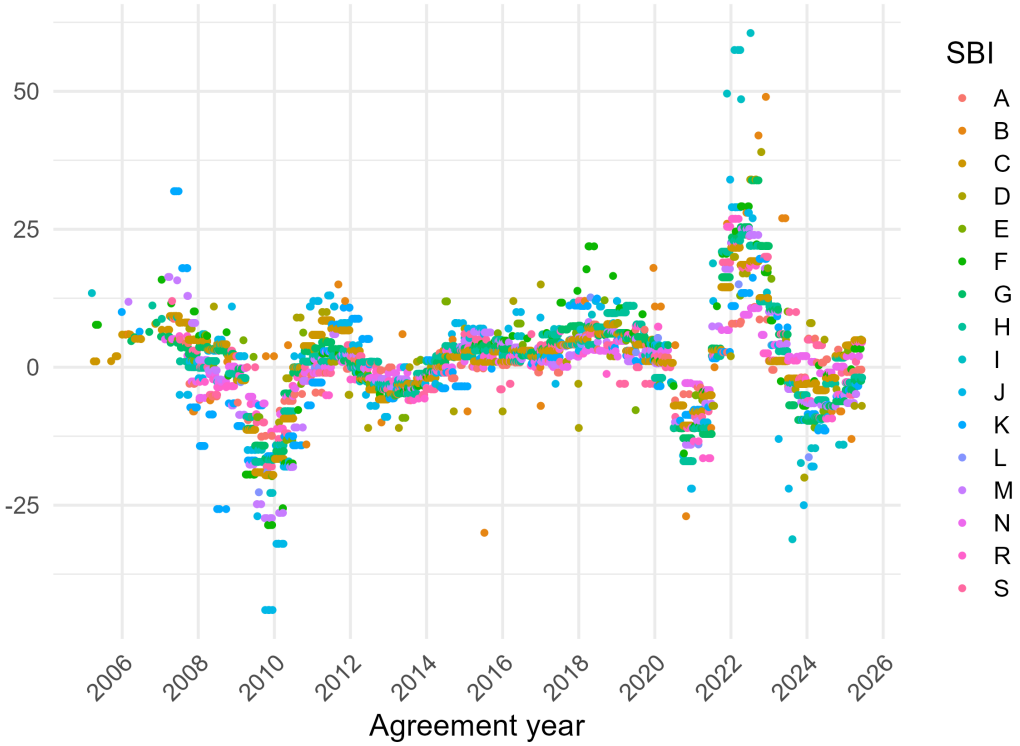
4.2 Sectoral vacancy rates

We measure sectoral labor market tightness, $\Delta\theta_{s,t-2}$, using the year-on-year change in vacancy rates as reported by Statistics Netherlands, lagged by two quarters.[16] The vacancy rate is defined as the number of outstanding vacancies per thousand jobs within a specific sector. Ideally, we would have used a more traditional labor market tightness measure, e.g. the ratio of the number of vacancies to unemployed workers. However, this information is not available at the sectoral level.

The dynamics of sectoral tightness, illustrated in Figure 6, are characterized by two pronounced structural shocks that dominate the sample period. Following the onset of the 2008 financial crisis, the growth of vacancy rates across all sectors became negative, reflecting a significant increase in labor market slack. A more dramatic shift occurred during and after the COVID-19 pandemic: an initial sharp contraction was followed by a rapid surge once lockdown restrictions were lifted. This post-pandemic surge pushed labor market tightness to historically high levels across the Dutch economy.

[16] One reason to use a two-quarter lag is that most CBAs are to some degree at least partially retroactive. We dive deeper into this issue in Appendix B.2.

Figure 6: Year-on-year vacancy rate changes per sector, 2005-2025Q2



4.3 Macroeconomic data and principal components

In this section, we discuss which macroeconomic variables enter $\boldsymbol{M}_{t-2}$ in equation (6).

4.3.1 The centrally demanded wage growth percentage

As a first candidate, we include the centrally demanded wage growth percentage (CDW) introduced in Section 2. As discussed before, it represents the FNV's national benchmark for wage increases and serves as a coordination device for sectoral negotiators. Given that FNV is by far the largest labor union in the Netherlands, its bargaining strategy and wage claim can play a potentially large role in determining aggregate CBA wage dynamics. From an econometric perspective, it captures aggregate inflation expectations and the perceived national bargaining power of labor union representatives, which are otherwise absorbed by standard time fixed effects.

The CDW is typically announced each September and remains the leading union objective until the subsequent year's announcement. To align this annual variable with our CBA panel, we map the claim to all agreements concluded within the particular period during which the CDW was relevant. Specifically, for each individual CBA, the variable takes the value of the most recent claim announced prior to the agreement date. An agreement concluded between the announcement in September of year $t-1$ and the day

preceding the announcement in year t is thus assigned the year t claim. We compiled the historical series by combining records from van der Velden (2016) with contemporary data from FNV's official publications.

In 2023, FNV published a range for the CDW between 5% and 14% in order to differentiate between CBAs which had already negotiated inflation compensation for the 2022 inflation spike and those that did not. In our analysis we decided to include the maximum of 14%, and control for idiosyncratic inflation compensation via our inflation shortfall variable (see Section 4.3.3).

4.3.2 Principal components of macroeconomic data

Further decomposing τ_t likely requires a richer set of macroeconomic controls than can feasibly be included directly in equation (6). To address this challenge, we adopt a dynamic factor modeling approach. Rather than entering dozens of macroeconomic indicators directly into the regression, we extract a small number of lagged latent principal components that capture the common variation across a broad set of observed macroeconomic variables.

While it is possible to estimate these factors jointly with equation (6), to keep the analysis tractable, we incorporate the principal components in the model as suggested by Stock and Watson (2002). Specifically, we first apply principal component analysis (PCA) to k macroeconomic variables to extract the leading m orthogonal components, which serve as proxies for the latent factors. We implement the PCA using `Stata` with varimax rotation to enhance interpretability of the resulting components.

After obtaining the principal components, we characterize each component by running the following $k \times m$ regressions:

$$y_{i,t} = \phi_{i,j} pc_{j,t} + \epsilon_{i,j,t}, \tag{7}$$

where $\phi_{i,j}$ captures the bivariate correlation between macroeconomic variable $y_{i,t}$ and principal component $pc_{j,t}$, and hence the power of $pc_{j,t}$ in predicting $y_{i,t}$, $i = 1, ..., k$, $j = 1, .., m$. We subsequently rank the resulting R^2-values for each principal component $pc_{j,t}$ to see which variables are best described by which component and whether these variables can be clustered in an economically meaningful way.

Based on our interviews with representatives from the employers' organization and the largest labor union, we collected a dataset of 52 macroeconomic variables, which were deemed the most relevant for the wage bargaining process. These include measures of real activity, such as GDP growth, unemployment and vacancy rates, inflation and its various subcomponents, productivity, economic sentiment, and union activity.[17] The majority of the data are obtained from Statistics Netherlands.[18]

[17] We include both the CPI and HICP.

[18] Union representatives mainly use data from Statistics Netherlands to inform the bargaining process.

We also include CPB forecasts for inflation and productivity, as these have historically been instrumental in the formation of the FNV's CDW. Because contract durations vary among CBAs signed at the same point in time, the relevant inflation forecast—calculated as an average over the specific contract period—is not uniform across the sample. Consequently, despite being derived from macroeconomic data, these variables introduce cross-sectional variation that may affect our estimates of the coefficient on sectoral labor market tightness. These variables hence serve as control variables both at the aggregate and CBA level, ensuring a more robust estimation of the coefficient on sectoral labor market tightness.

The raw data contain annual, quarterly, and monthly variables. Since our labor market tightness measures are available at the quarterly level, we took quarterly averages of the monthly variables to make them more comparable. To ensure stationarity, variables identified as I(1) enter the PCA in year-on-year growth rates (or first differences for variables already expressed as rates). A description of each variable is given in Table 13 in Appendix B.5.

From our macroeconomic dataset we extract 8 principal components, which are shown in Figures 15a-16d in Appendix B.5.[19] In Table 3 we list the variables which had the highest R^2 value in regressions (7). The regression results are shown in Table 15 in Appendix B.5. We can use these values to characterize each component.

We take a simple average of the two principal components associated with aggregate labor market tightness to measure $\Delta\theta_{t-2}$ in equations (5) and (6). In Table 3 these principal components are denoted by $\Delta\theta^{(1)}_{t-2}$ and $\Delta\theta^{(2)}_{t-2}$. The loadings of these two components partially overlap. Principal component $\Delta\theta^{(1)}_{t-2}$ is associated with (un)employment, GDP growth, and vacancy growth, while $\Delta\theta^{(2)}_{t-2}$ mainly captures other vacancy measures in the dataset. Strictly speaking, these two principal components do not only capture variations in aggregate labor market tightness, since they also load on other macroeconomic variables such as imports, GDP growth, and inflation. However, given that both $\Delta\theta^{(1)}_{t-2}$ and $\Delta\theta^{(2)}_{t-2}$ load strongly on vacancies and other labor market indicators (e.g. jobs and unemployment) that capture labor demand, we include both in our benchmark model. $\Delta\theta^{(1)}_{t-2}$ broadly captures macroeconomic variables that explain changes in employment, while $\Delta\theta^{(2)}_{t-2}$ does so for vacancies and labor demand. In a robustness test, we show that using either $\Delta\theta^{(1)}_{t-2}$ or $\Delta\theta^{(2)}_{t-2}$ as a proxy for aggregate labor market tightness does not alter our main conclusions.

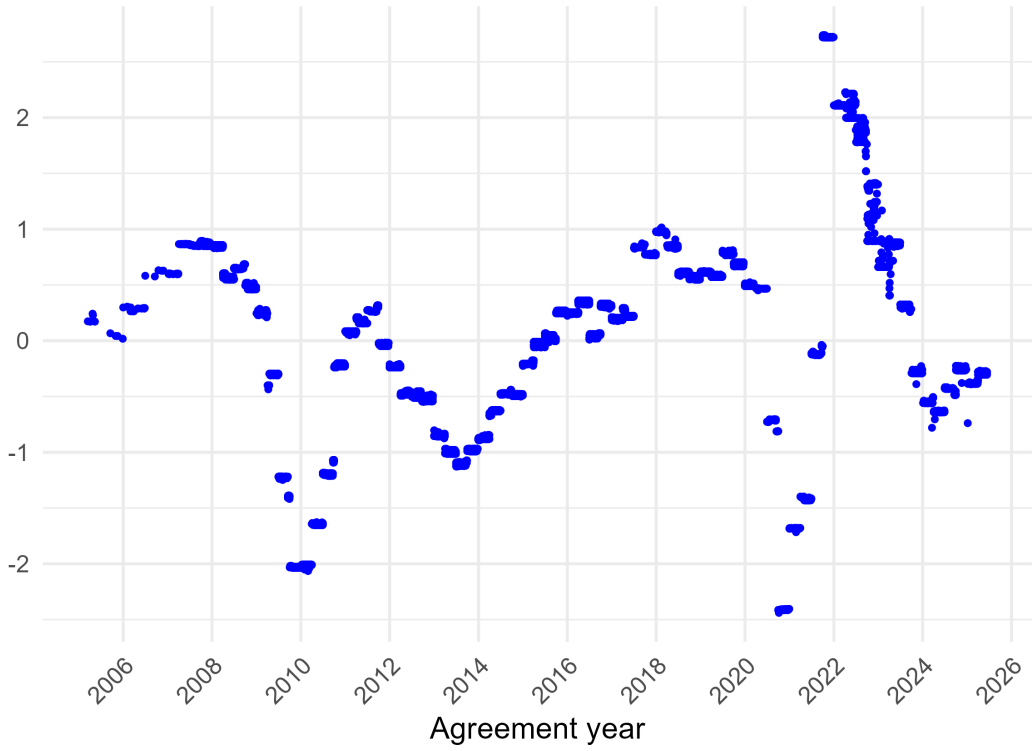
[19] We follow the Kaiser Criterion, retaining all components with an eigenvalue greater than 1.0. While the ninth component technically meets this threshold ($\lambda = 1.00268$), its marginal contribution to the total explained variance is negligible (1.93%). We therefore exclude it to maintain a more parsimonious factor structure, as the first eight components already capture approximately 87% of the total variation in the dataset. The eigenvalues and explained variance of the first 8 principal components are shown in Table 14 in appendix B.5.

Table 3: Principal component vectors and their associated economic indicators

PC	Selection of associated macroeconomic variables
$\Delta\theta^{(1)}$	Total employment, jobs, net labor force participation rate, jobs of employees, labor force growth, unemployment rate, GDP, hours worked, expected vacancy rate index, filled vacancies.
$\Delta\theta^{(2)}$	New vacancies, import index volumes, unfilled vacancies, vacancy rate, filled vacancies, expected vacancy rate index.
$\pi^{(1)}$	Goods CPI inflation, energy CPI inflation, import index inflation, headline HICP inflation, headline CPI inflation, HICP and CPI CPB forecast.
$\pi^{(2)}$	CPI inflation excluding energy and food, services CPI inflation, food, alcohol and tobacco CPI inflation, GDP deflator CPB August forecast, CPI net of taxes CPB August forecast.
Δs	Perceived economic climate, consumer sentiment, perceived economic conditions last 12 months, perceived economic climate next 12 months.
Δz	Productivity growth measures.
$\Delta u^{(1)}$	Union membership.
$\Delta u^{(2)}$	Hours and workers involved in union activity.

Note: This table lists a selection of the macroeconomic variables associated with each principal component (PC) derived from a principal component analysis. All variables are either y-o-y growth rates or differences, unless otherwise indicated.

Figure 7: Aggregate labor market tightness $\Delta\theta_{t-2}$ measured by the standardized average of principal components $\Delta\theta^{(1)}$ and $\Delta\theta^{(2)}$



The remaining components enter our models separately and are estimated with constant coefficients. The two components associated with inflation variables $\pi^{(1)}_{t-2}$ and $\pi^{(2)}_{t-2}$ capture different inflation measures and different types of inflation (e.g. energy and goods price inflation). Principal components Δs_{t-2} and Δz_{t-2} respectively mainly describe economic sentiment indicators and aggregate productivity growth. Lastly, we extracted two principal components measuring union membership ($\Delta u^{(1)}_{t-2}$) and hours and workers involved in union activity ($\Delta u^{(2)}_{t-2}$).

4.3.3 Inflation shortfall

In addition to include two principal components reflecting inflation measures, we construct a measure of *inflation shortfall* at the CBA level, capturing deviations of *realized* inflation from *expected* inflation. Since CBA wage setting is staggered and often partly retroactive, information on inflation available to negotiators during bargaining consists of both realized inflation and inflation forecasts. This is especially relevant for the post-pandemic period and after the Russian invasion of Ukraine, when forecasted and realized inflation deviated substantially. To quantify this inflation shortfall, we construct

$$\tilde{\pi}^{p}_{i,t} = \max\Big(\overline{\pi}^{p}_{i,t} - \overline{\pi}^{p,e}_{i,t},\, 0\Big)\,. \tag{8}$$

Here, $\overline{\pi}^{p}_{i,t}$ denotes the average realized inflation during period p, as observed at agreement date t for CBA i, while $\overline{\pi}^{p,e}_{i,t}$ denotes the average expected inflation rate for that same period and CBA. Although this variable is created using macroeconomic data, it is CBA-specific: p differs per individual CBA and over time, and hence $\tilde{\pi}^{p}_{i,t}$ can differ for different i observed at the same time t. As with the CPB inflation forecasts, this introduces cross-sectional variation that may influence our estimates of the coefficient on sectoral labor market tightness.

Because the majority of agreements is partially retroactive, this expectation measure itself combines realized inflation (for the elapsed part of the contract period) and inflation forecasts (for the remaining part). We construct this variable using the CPB CPI inflation forecast released most recently prior to agreement date t. If the contract period lies entirely in the past ($p_e < t$), $\overline{\pi}^{p,e}_{i,t}$ consists entirely of realized inflation data and $\tilde{\pi}^{p}_{i,t}$ is 0. We take the maximum of the difference between $\overline{\pi}^{p}_{t}$ and $\overline{\pi}^{p,e}_{i,t}$ and 0, as the inflation shortfall is expected to only play a role in subsequent negotiations if realized average inflation turned out to be higher than what was expected at time t. We include the lagged value of this measure, computed for the previous contract period $p-1$, in our empirical specification, as inflation shortfalls realized under the previous agreement constitute the relevant reference point in subsequent wage negotiations.

5 Empirical results

In this section, we present our empirical findings in three stages. We begin by estimating a fixed parameter specification using standard panel OLS methods to provide a benchmark for our results. We then transition to a state space framework, allowing for time-varying parameters and replacing the traditional time fixed effects with a time-varying intercept. Finally, we decompose this intercept into its macroeconomic drivers, including the average aggregate labor market tightness principal component. This decomposition allows us to compare the relative influence of sectoral versus aggregate labor market tightness on wage growth.

5.1 Estimates of a fixed coefficient on sectoral labor market tightness

As a benchmark for our model with time-varying parameters, we investigate the relationship between sectoral variation in labor market tightness and CBA wage growth in a fixed parameters TWFE model. We do so by estimating equation (1) using standard panel OLS techniques. By construction, the relationship between sectoral labor market tightness and CBA wage growth is constant over time. The results are shown in Table 4.

Our results indicate that sectoral differences in labor market tightness cannot explain much of the variation in CBA wage growth once time fixed effects are added. In the first column of Table 4, we show that for the full sample β_1 is positive and significantly different from zero when we do not include fixed effects. The addition of CBA fixed effects (second column) does not alter this result. However, once we include both CBA *and* time fixed effects, the relationship between sectoral labor market tightness and CBA wage growth reduces in size and becomes insignificantly different from zero. This is also the case for all the subsamples. The within R^2 statistics in Table 5 confirm this: differences in the sectoral vacancy rate explain almost none of the variation in CBA wage growth after accounting for CBA and time fixed effects.

Table 4: Estimates of β_1 from equation (1)

	No FE	Only CBA FE	CBA and time FE
	(1)	(2)	(3)
Full sample	0.025***	0.023***	0.003
	(0.002)	(0.002)	(0.005)
2005-2008Q3	0.003	−0.009	0.011
	(0.008)	(0.019)	(0.019)
2008Q4-2019	0.019***	0.021***	0.005
	(0.001)	(0.001)	(0.004)
2020-2025Q2	0.001	0.002	0.002
	(0.004)	(0.004)	(0.012)
CBA fixed effects		✓	✓
Time fixed effects			✓

Notes: Standard errors are in parentheses, clustered by CBA, although assuming homoskedastic standard errors does not alter our conclusions regarding statistical significance, see Table 16 in Appendix D. FE refers to 'fixed effects'. Statistical significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 5: R^2 statistics pertaining to equation (1)

Period	R^2	Adjusted R^2	Within R^2
Full sample	0.6584	0.6019	0.0001
2005-2008Q3	0.7140	0.3764	0.0056
2008Q4-2019	0.5626	0.4536	0.0006
2020-2025Q2	0.6633	0.5072	0.0000

Notes: R^2 statistics pertain to the specification in column (3) of Table 4, which includes both CBA and time fixed effects. The within R^2 measures the explanatory power of the regressor after partialling out the fixed effects.

To better understand why labor market tightness no longer significantly affects CBA wage growth after the introduction of time fixed effects, we assess the relative importance of cross-sectional versus common aggregate variation. We do so by decomposing the total

variance of $\Delta w_{i,s,t}$ and $\Delta\theta_{s,t-2}$ into three components: between-CBA, common aggregate, and idiosyncratic variance.

First, we calculate the *between-CBA variance* by taking the variance of the time-averaged means across the 949 CBAs in our sample: $\sigma_B^2 = \text{Var}(\bar{x}_i) = \frac{1}{N-1}\sum_i(\bar{x}_i - \bar{x})^2$, $i = 1, ..., 949$. This captures structural, cross-sectional differences in the levels of the variables across different CBAs. In a standard regression framework, this variation is absorbed by CBA fixed effects.

Second, the *common aggregate variance* is defined as $\sigma_C^2 = \text{Var}(\bar{x}_t) = \frac{1}{T-1}\sum_t(\bar{x}_t - \bar{x})^2$, where $\bar{x}_t$ represents the quarterly cross-sectional means. This captures shared fluctuations affecting all units simultaneously—such as macro-shocks or the synchronized post-pandemic recovery—which are typically captured by time fixed effects.

Finally, we compute the *idiosyncratic variance* from the "double-demeaned" residuals: $\sigma_I^2 = \text{Var}(x_{it} - \bar{x}_i - \bar{x}_t)$. This represents the purely cross-sectional variation remaining after removing both CBA and time fixed effects. Crucially, this "within-within" variation is the *only* source of data used to identify the β coefficients in a standard TWFE model.

In Table 6, we report these variance components for $\Delta w_{i,s,t}$ and $\Delta\theta_{s,t-2}$ relative to their respective total variances. The results indicate that while the between-CBA variance is relatively small, the common aggregate variance is the dominant source of fluctuation for both variables. This is particularly evident for labor market tightness ($\Delta\theta_{s,t-2}$), where the aggregate signal is nearly four times larger than the idiosyncratic component. These findings suggest that a standard TWFE estimator effectively discards the most informative variation in the data, forcing the identification of the wage-tightness relationship to rely on a much smaller, and potentially noisier, idiosyncratic residual.

Our panel OLS results and the comparison of the variance components of both CBA wage growth and sectoral labor market tightness suggest that macroeconomic variation, captured by the time fixed effect, plays a relatively large role in determining CBA wage dynamics. In the next subsection, we investigate whether the effects of sectoral labor market tightness and macroeconomic conditions, as captured by the time fixed effects, on CBA wage growth are time-varying.

Table 6: Comparing between, common and idiosyncratic variance of $\Delta w_{i,s,t}$ and $\Delta\theta_{s,t-2}$

	$\Delta w_{i,s,t}$		$\Delta\theta_{s,t-2}$	
Variance	Absolute	Relative	Absolute	Relative
Total	2.20	1.00	74.12	1.00
Between	0.39	0.18	6.96	0.09
Common	1.18	0.54	61.10	0.82
Idiosyncratic	0.87	0.39	16.12	0.22

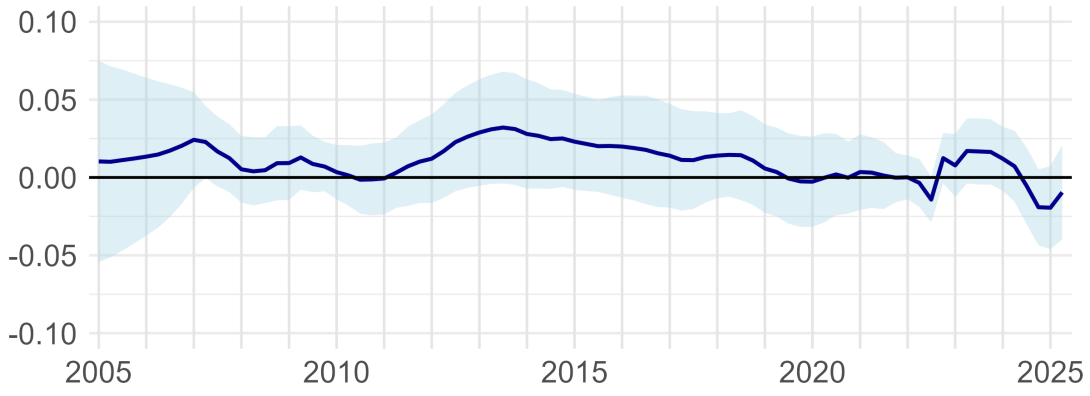
Note: Relative variance is calculated as the component variance divided by the total variance. The relative shares do not necessarily sum to unity as the covariance terms do not perfectly cancel out.

5.2 Estimates of a time-varying coefficient on sectoral labor market tightness

While our preliminary panel OLS estimates suggest that the relationship between sectoral labor market tightness and wage growth remains statistically insignificant even across specific temporal windows, this result may be sensitive to the manual partitioning of the data. To address this, we employ a state space framework to estimate a time-varying coefficient on sectoral labor market tightness, $\beta_{1,t}$. A state space model treats the parameter evolution as a stochastic process. By optimizing the likelihood across the full sample, the model endogenously identifies the optimal degree of time variation in $\beta_{1,t}$. This approach provides a more efficient and flexible test of whether the relationship has changed over time.[20]

By estimating equation (3) with time-varying parameters, we find that the smoothed estimate of $\beta_{1,t}$, shown in Figure 8, reinforces the view that cross-sectional differences in labor market tightness contribute little to explaining wage growth. Even though the parameter value does vary over time, it is almost never significantly different from zero.[21] Moreover, a model with a fixed coefficient on sectoral tightness yields very similar estimates, see column 1 of Table 17 in Appendix D. Hence, the post-2020 acceleration in CBA wage growth cannot be explained by differences in labor market tightness between sectors.

Figure 8: Kalman smoothed estimates of $\beta_{1,t}$ in model (3)



Note: The shaded area indicates the 95%-confidence interval.

By contrast, the estimated path of the time-varying intercept τ_t, depicted in Figure 9, sheds new light on the role of aggregate dynamics in explaining wage growth. We observe a sharp rise in τ_t just prior to the full onset of the Global Financial Crisis in the Netherlands,

[20] We also estimated a state space model with fixed parameters to confirm that our results are not sensitive to our modeling approach, see the first column in Table 17 in Appendix D.

[21] Note that the state space model assumes the variance of the irregular term to be homoskedastic. We chose to keep this part of the model parsimonious, since the significance of the parameters in the OLS estimates was not affected by our choice of standard errors: the clustered standard errors are very similar to homoskedastic standard errors, see Tables 4 and 16.

followed by a steep decline during the ensuing recession and the European sovereign debt crisis. Because the within-transformation centers the data around entity-specific means, the negative values of τ_t during this period reflect aggregate wage growth falling below its long-run average. Furthermore, as τ_t is estimated conditional on sectoral labor market tightness, it effectively captures the common macroeconomic shocks that sectoral variation fails to explain. From 2016 onward, τ_t climbs steadily until the COVID-19 pandemic, after which it rises sharply during the post-pandemic recovery and the inflationary surge following Russia's invasion of Ukraine in 2022. Therefore, factors common to *all* sectors predominantly drove higher CBA wage growth during and after the post-pandemic recovery.

Figure 9: Kalman smoothed estimates of τ_t in model (3)



Note: The shaded area indicates the 95%-confidence interval.

Since τ_t captures the variation explained by all common aggregate dynamics, this opens the possibility that *aggregate* labor market tightness contained within τ_t *does* impact wage growth. Therefore, in the next section, we decompose τ_t to examine to what extent aggregate dynamics in labor market tightness drive the dynamics in τ_t, compared to other macroeconomic variables.

5.3 Decomposition results and estimates of the effects of macro labor market tightness

We now seek to identify the macroeconomic determinants of τ_t, in particular the role of aggregate labor market tightness. Our approach addresses the missing intercept problem by explicitly modeling the aggregate fluctuations that were previously absorbed by the time-varying intercept. Specifically, we investigate whether *aggregate* labor market tightness, rather than differences in tightness between sectors, can explain observed changes in CBA wage growth.

We investigate this by first estimating the extended model (5), which includes both sectoral labor market tightness $\Delta\hat{\theta}_{s,t-2}$ and its aggregate counterpart $\Delta\theta_{t-2}$. We further

expand this in model (6) by introducing a vector of macroeconomic control variables $\boldsymbol{M}_{t-2}$, including the remaining principal components, the inflation shortfall variable (8), and the CDW. Table 7 compares the end-of-sample parameter estimates and model fit across these specifications. We benchmark these against the results of model (3) in column 1. Columns (2) and (3) show the results for equation (5) and (6), respectively. Column 4 presents a version of model (6) where $\boldsymbol{M}_{t-2}$ also includes interaction terms between the inflation variables and a time dummies to account for nonlinearities in the post-pandemic period.[22] A detailed version containing all parameter estimates is provided in Table 17 in Appendix D.

Table 7: End of sample parameter estimates and model diagnostics across specifications

	(1) Benchmark model (3)	(2) Extended model (5)	(3) Extended model (6)	(4) Full model (6)
Time-varying intercept at time T				
τ_T	0.832^{***}	1.144^{***}	-0.495^{**}	-0.869^{***}
	(0.097)	(0.161)	(0.233)	(0.175)
Sectoral labor market tightness coefficient at time T				
$\beta_{1,T}$	-0.010	-0.010	-0.008	0.003
	(0.017)	(0.014)	(0.015)	(0.004)
Aggregate labor market tightness coefficient at time T				
$\beta_{2,T}$		1.025^{**}	0.843^{**}	0.907^{**}
		(0.433)	(0.412)	(0.431)
Decomposition of τ_t				
$\Delta\theta_{t-2}$		✓	✓	✓
$\boldsymbol{M}_{t-2}$			✓	✓
Inflation interaction terms in $\boldsymbol{M}_{t-2}$				✓
Log-likelihood	-9815.871	-9784.075	-9713.702	-9684.626
AIC	0.252	0.252	0.250	0.249
$\sqrt{\text{Var}(\tau_t)}$	0.935	0.888	0.206	0.115
$\sqrt{\text{Var}(\tau_t)}$ (rel.)	1.000	0.950	0.220	0.123

Note: Standard errors in parentheses. ***, **, and * denote significance at 1%, 5%, and 10%.

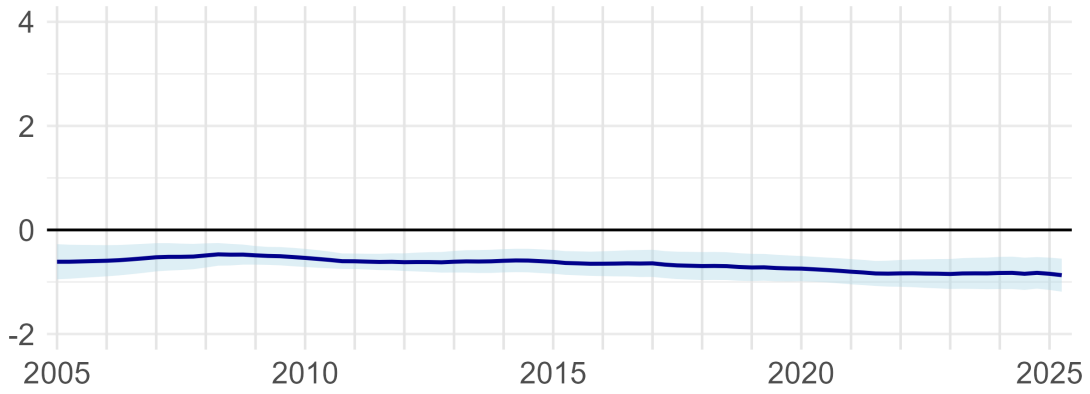
Traditional model selection criteria such as the Akaike Information Criterion (AIC) are less informative in this context. Due to the Kalman filter's smoothing mechanism, τ_t is estimated to minimize prediction errors, which can result in similar AIC values across specifications. We therefore propose an additional model selection criterion: the estimated

[22] We interact a time dummy that equals one when $t \geq$ 2021Q3 and zero otherwise with the principal components $\pi^{(1)}_{t-2}$ and $\pi^{(2)}_{t-2}$, as well as the inflation shortfall variable $\tilde{\pi}^{p-1}_{i,s,t}$. We found that the coefficient on $\pi^{(1)}_{t-2}\mathbb{1}_{t-2\geq 2021Q3}$ was negative, which is an indication that one dummy may not fully capture the complex wage-inflation dynamics in the post-COVID period. Therefore, we also interacted $\pi^{(1)}_{t-2}$ with a dummy that equals one for $t \geq$ 2022Q3. That coefficient has a positive sign, see Table 17 and section D.

volatility of the time-varying intercept $\sqrt{\text{Var}(\tau_t)}$ relative to the baseline model (3). Since τ_t captures residual aggregate dynamics not explained by the regressors, a lower variance implies that a greater share of systematic macroeconomic variation is accounted for by the model. We use this reduction as an empirical metric to compare the benchmark model to the extended model with aggregate labor market tightness and other macroeconomic variables.

Column 4 of Table 7 shows that the most comprehensive specification, which includes the inflation interaction terms in $\boldsymbol{M_{t-2}}$ in model (6), provides the best fit to the data. It yields the highest log-likelihood and the lowest AIC, while the volatility of τ_t is reduced to just 12.3% of its baseline value. This results in a nearly flat time-varying intercept, as shown in Figure 10. Notably, restricting τ_t to a constant in this specification results in very similar values for the log-likelihood and AIC, and produces nearly identical estimates for $\beta_{2,t}$ (see Appendix E.2.1). We therefore conclude that our decomposition is successful: the included macroeconomic variables account for virtually all common aggregate variation in CBA wage growth.

Figure 10: Kalman smoothed estimates of τ_t in full model (6)



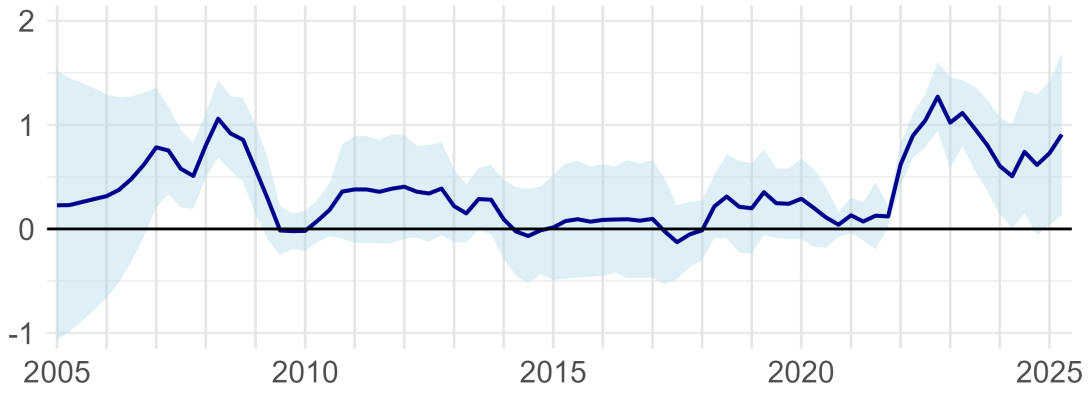
Note: The shaded area indicates the 95%-confidence interval.

The corresponding smoothed estimate of the time-varying coefficient on aggregate labor market tightness indicates that while the relationship between labor market tightness and CBA wage growth is predominantly positive over the sample period, it only becomes statistically significant during the 2006Q4-2009Q2 period and after 2021Q3, see Figure 11. The point estimate for $\beta_{2,t}$ peaks in both 2008 and 2022 at values that exceed 1. At the end of the sample the Kalman smoothed estimate of $\beta_{2,t}$ is 0.907, which implies that a one-standard-deviation increase in aggregate labor market tightness is associated with a 0.907 percentage point increase in negotiated wage growth.

The significant relationship between $\Delta\theta_{t-2}$ and wage growth can be rationalized as follows. The period before the Global Financial Crisis saw an increase in real economic activity, as seen during the 2006-2008 period in Figure 7. The growth rate of labor market

tightness $\Delta\theta_{t-2}$ decelerates in the first quarters of 2008 until it declines sharply in early 2009, while its estimated coefficient $\beta_{2,t}$ peaks in the second quarter of 2008. After that period, $\beta_{2,t}$ starts decreasing and it becomes statistically insignificant a year later. The significant relationship observed toward the end of the sample coincides with a sharp acceleration in the principal component that captures aggregate labor market tightness, shown in Figure 7. These patterns support the presence of nonlinearities in the wage Phillips curve, consistent with recent findings by Benigno and Eggertsson (2023) and Gitti (2024).[23]

Figure 11: Kalman smoothed estimates of $\beta_{2,t}$ in full model (6)

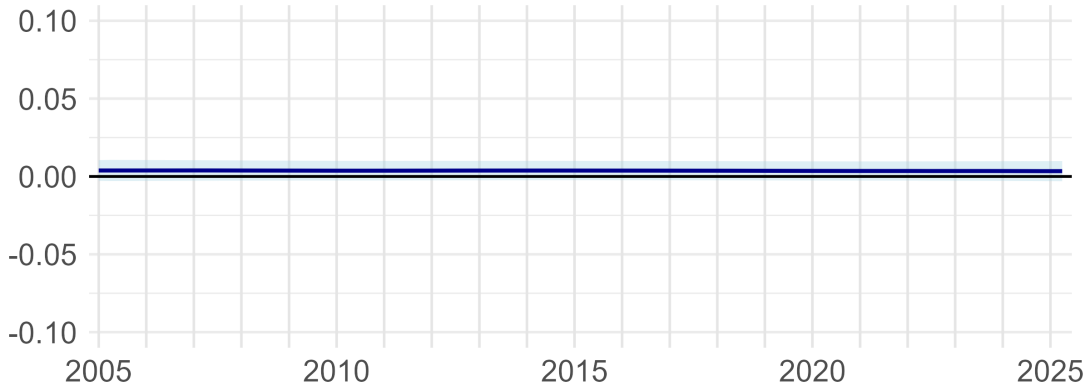


Note: The shaded area indicates the 95%-confidence interval.

For this specification, the smoothed estimate of the time-varying coefficient on sectoral labor market tightness, $\beta_{1,t}$, converges endogenously to an insignificant constant. The sensitivity of the sectoral tightness estimates to the inclusion of macroeconomic data stems from the fact that the inflation variables in our macroeconomic dataset exhibit idiosyncratic cross-sectional variation. As detailed in Section 4.1, this variation arises from the staggered nature of CBA contracts and varying forecast horizons. Importantly, while the precision of the estimate improves, our qualitative conclusions regarding the relative unimportance of sectoral labor market tightness remain unchanged by the inclusion of these macroeconomic variables.

[23]Note that in both the first and final quarters of the sample, the confidence interval around the point estimates of $\beta_{2,t}$ widens considerably. This loss of precision is correctly recognized by the Kalman smoother since fewer observations are available at both ends of the sample. On top of this, the widening of the confidence interval is also due to relatively sparse data at the beginning of the sample.

Figure 12: Kalman smoothed estimates of $\beta_{1,t}$ in full model (6)



Note: The shaded area indicates the 95%-confidence interval.

5.4 Robustness checks

We performed several robustness checks to determine whether our findings are sensitive to specific modeling choices.

First, we examined whether our estimates are influenced by the high concentration of goods-producing sectors in the dataset, particularly Sector C (manufacturing), which accounts for nearly half of the sample. To address this, we re-estimated the benchmark model using two distinct subsamples: one consisting solely of goods-producing sectors and another of services-sector CBAs. The resulting estimates for τ_t and $\beta_{1,t}$, presented in Appendix E.1.1, were nearly identical to our primary results. Consequently, we did not re-estimate our extended models using these subsamples.

Second, we investigated whether partitioning the sample into firm-specific and industry-wide CBAs affected our conclusions. The estimates for the benchmark model shown in Appendix E.1.2 indicate that the values for τ_t remain nearly identical across these groups. While the estimates for $\beta_{1,t}$ vary slightly, as $\beta_{1,t}$ for industry CBAs is reduced to a fixed parameter within the model, neither estimate is significantly different from zero. These results suggest that the type of CBA does not fundamentally alter our main findings, rendering further estimation of the extended models for these subsamples unnecessary.

Third, we checked whether the estimates of the coefficient on our measure of sectoral labor market tightness $\beta_{1,t}$ in full model (6) are robust to adding sectoral control variables: energy cost sensitivity $\pi^e_{s,t-4}$ (lagged by four quarters since the data are at an annual frequency), and sectoral productivity $\Delta\hat{z}_{s,t-2}$. As shown in Appendix E.2.2, the coefficients on these variables are not statistically significant and their inclusion does not qualitatively alter our primary findings.

Fourth, we tested the sensitivity of our full model (6) to different measures of sectoral labor market tightness, by including the sectoral vacancy rate in levels rather than growth rates. While the latter is our default specification, the former is more conventional in the

wage Phillips curve literature. The resulting estimates for τ_t, $\beta_{1,t}$, and $\beta_{2,t}$ are presented in Appendix E.2.3. The results remain largely consistent with our baseline growth-rate specification: although $\beta_{1,t}$ is statistically significant at the start of the sample period, its magnitude is negligible and diminishes over time. Similarly, we tested whether measuring sectoral labor market tightness as a four-quarter moving average of vacancy growth affected the results. The results shown in Appendix E.2.4 indicate that this is not the case.

We also tested whether using the sectoral equivalent of our aggregate measure of tightness, a principal component of sectoral economic data, has similar outcomes as the sectoral vacancy growth rate. As some of the sectoral variables are only available at the annual level, the principal component is less precise than the vacancy rate growth rate. The resulting estimated coefficient shown in Appendix E.2.5 is larger, but also more volatile as well as insignificant, which persuaded us not to alter our main conclusions.

Finally, because our aggregate tightness measure is the standardized average of $\Delta\theta^{(1)}$ and $\Delta\theta^{(2)}$, we check whether the results depend on this averaging by replacing the average with either component alone (Appendices E.2.6 and E.2.7). Using $\Delta\theta^{(1)}$ yields very similar estimates. Using $\Delta\theta^{(2)}$ yields an estimate of $\beta_{2,t}$ that is not significant in the 2007–2009 period, when $\Delta\theta^{(2)}$ displays considerably smaller fluctuations than $\Delta\theta^{(1)}$. To us this demonstrates the benefits of combining the information available in the two components, as averaging makes the aggregate measure less dependent on the dynamics of either one.

6 Discussion and conclusion

This paper investigates the nonlinear relationship between labor market tightness and collectively bargained wage growth in the Netherlands using a dynamic panel state space model. Our approach addresses the identification problem inherent in fixed-effects panel models, where time dummies absorb aggregate variation and obscure the macroeconomic channel through which labor market tightness affects wages. We propose a novel decomposition of the time-varying intercept using principal components derived from a broad set of macroeconomic indicators in addition to variables pertaining to the Dutch collective bargaining institutions. This enables us to recover both cross-sectional and time-series variation in tightness and assess their distinct contributions to wage growth dynamics.

We find that cross-sectional differences in sectoral vacancy rates explain little of the variation in CBA wage growth, even when allowing for time-varying slope coefficients. By contrast, aggregate measures of labor market tightness, particularly those summarized by an average of the two principal components which load heavily on vacancy-related indicators, are strongly associated with CBA wage growth, especially during the post-pandemic period. This suggests that aggregate macroeconomic dynamics, rather than sector-specific labor market conditions, are the dominant drivers of negotiated wage outcomes in the Netherlands. Furthermore, our results suggest that higher CBA wage

growth during the post-pandemic recovery cannot be explained by labor shortages that only certain sectors faced.

A central finding of our analysis is the presence of a *nonlinear* relationship between aggregate labor market tightness and CBA wage growth. While this relationship is generally positive across the sample, its magnitude increases significantly after 2021, coinciding with the post-COVID recovery and unprecedented vacancy growth. In our preferred specification, the coefficient on macro tightness becomes statistically significant only during periods where the growth of real activity and labor demand accelerate. This pattern is consistent with recent work (e.g. Benigno and Eggertsson (2023); Gitti (2024)), which suggests that the wage Phillips curve steepens under conditions of extreme tightness. It also aligns with our qualitative findings: union representatives report that macroeconomic indicators—including inflation and national labor shortages—become more salient bargaining inputs during tight labor market episodes.

Our empirical framework, which combines time-varying intercepts and coefficients within a state space system, allows us to detect these dynamic and potentially regime-dependent effects. That the macro tightness–wage link strengthens only after 2021 implies that conventional linear models may understate the wage responsiveness to tightness during booms. Accurately modeling these nonlinearities is therefore essential for understanding wage dynamics and wage-push inflation risks under persistent labor shortages.

These results have several implications. First, macroeconomic conditions — notably labor market tightness and inflation — play a critical role in shaping centrally bargained wage outcomes in a highly institutionalized labor market. Second, despite sector-level heterogeneity and staggered bargaining schedules, central wage setting remains responsive to common macroeconomic conditions, reflecting a high degree of coordination in the bargaining process. Third, empirical strategies that fail to account for time variation and nonlinearity may mischaracterize the true sensitivity of wages to labor market conditions, particularly in periods of structural change.

Other OECD countries, particularly those in continental Europe and Scandinavia, are characterized by comparable collective bargaining institutions with similar degrees of coordination. As a result, our finding that cross-sectional differences in labor market tightness cannot explain wage growth may not be specific to the Netherlands. In addition, the nonlinear interaction between wage growth and the macroeconomy that we find for the Netherlands is potentially at play in other countries as well. Furthermore, since our findings pertain specifically to collectively bargained wages, they do not rule out a stronger role for sector-specific tightness in determining individually negotiated wages or wage drift. Future work should explore these channels, potentially through linked microdata on individual wages.

References

Benigno, Pierpaolo and Gauti B Eggertsson, "It's baaack: The surge in inflation in the 2020s and the return of the non-linear Phillips curve," Technical Report, National Bureau of Economic Research 2023.

Beraja, Martin, Erik Hurst, and Juan Ospina, "The aggregate implications of regional business cycles," *Econometrica*, 2019, *87* (6), 1789–1833.

Bhuller, Manudeep, Karl Ove Moene, Magne Mogstad, and Ola L Vestad, "Facts and fantasies about wage setting and collective bargaining," *Journal of Economic Perspectives*, 2022, *36* (4), 29–52.

Bonam, Dennis, Jakob de Haan, and Duncan Van Limbergen, "Time-varying wage Phillips curves in the euro area with a new measure for labor market slack," *Economic Modelling*, 2021, *96*, 157–171.

Caloia, Francesco G, Jante Parlevliet, and Mauro Mastrogiacomo, "Staggered wages, unanticipated shocks and firms' adjustments," *Journal of Macroeconomics*, 2023, *76*, 103521.

Calvo, Guillermo A, "Staggered prices in a utility-maximizing framework," *Journal of Monetary Economics*, 1983, *12* (3), 383–398.

CPB Netherlands Bureau for Economic Policy Analysis, "Centraal Economisch Plan 2026," 2026.

Doornik, Jurgen A, *Object-Oriented Matrix Programming using Ox*, Timberlake Consultants Press, 2021.

Durbin, James and Siem Jan Koopman, *Time Series Analysis by State Space Methods*, 2nd ed., Oxford University Press, 2012.

Erceg, Christopher J, Dale W Henderson, and Andrew T Levin, "Optimal monetary policy with staggered wage and price contracts," *Journal of Monetary Economics*, 2000, *46* (2), 281–313.

Fitzgerald, Terry, Callum Jones, Mariano Kulish, and Juan Pablo Nicolini, "Is there a stable relationship between unemployment and future inflation?," *American Economic Journal: Macroeconomics*, 2024, *16* (4), 114–142.

Francke, Marc K., Siem Jan Koopman, and Aart F. De Vos, "Likelihood functions for state space models with diffuse initial conditions," *Journal of Time Series Analysis*, 2010, *31* (6), 407–414.

Galí, Jordi, "The return of the wage Phillips curve," *Journal of the European Economic Association*, 2011, *9* (3), 436–461.

Gertler, Mark and Antonella Trigari, "Unemployment fluctuations with staggered Nash wage bargaining," *Journal of Political Economy*, 2009, *117* (1), 38–86.

_ , Luca Sala, and Antonella Trigari, "An estimated monetary DSGE model with unemployment and staggered nominal wage bargaining," *Journal of Money, Credit and Banking*, 2008, *40* (8), 1713–1764.

Gitti, Giulia, "Nonlinearities in the Regional Phillips Curve with Labor Market Tightness," 2024.

Górnicka, Lucyna and Gerrit Koester, *A forward-looking tracker of negotiated wages in the euro area: Eurosystem wage tracker experts* number 338, ECB Occasional Paper, 2024.

Hazell, Jonathon, Juan Herreño, Emi Nakamura, and Jón Steinsson, "The slope of the Phillips Curve: evidence from US states," *The Quarterly Journal of Economics*, 2022, *137* (3), 1299–1344.

Hooper, Peter, Frederic S Mishkin, and Amir Sufi, "Prospects for inflation in a high pressure economy: Is the Phillips curve dead or is it just hibernating?," *Research in Economics*, 2020, *74* (1), 26–62.

Huizinga, Anna and Bram Hendriks, "Price passthrough via production networks in the small open economy," forthcoming. CPB Discussion Paper.

Jungbacker, Borus and Siem Jan Koopman, "Likelihood-based Dynamic Factor Analysis for Measurement and Forecasting," *Econometrics Journal*, 2015, *18*, C1–C21.

Koopman, Siem Jan, Neil Shephard, and Jurgen A. Doornik, "Statistical algorithms for models in state space using SsfPack 2.2," *The Econometrics Journal*, 1999, *2* (1), 107–160.

Matthes, Christian, Naoya Nagasaka, and Felipe F Schwartzman, "Estimating the missing intercept," 2025.

Mavroeidis, Sophocles, Mikkel Plagborg-Møller, and James H Stock, "Empirical evidence on inflation expectations in the New Keynesian Phillips Curve," *Journal of Economic Literature*, 2014, *52* (1), 124–188.

McLeay, Michael and Silvana Tenreyro, "Optimal inflation and the identification of the Phillips curve," *NBER Macroeconomics Annual*, 2020, *34* (1), 199–255.

Ministry of Social Affairs and Employment, "Cao-loonpatronen in tijden van corona en inflatie," 2025. Accessed on April 21, 2026.

Ministry of the Interior and Kingdom Relations, "Evaluatie van het referentiemodel: Onderzoek naar de loonontwikkeling in de sector overheid," Technical Report, Ministry of the Interior and Kingdom Relations 2017.

Mortensen, Dale T and Christopher A Pissarides, "Job creation and job destruction in the theory of unemployment," *The Review of Economic Studies*, 1994, *61* (3), 397–415.

Schmitt-Grohé, Stephanie and Martin Uribe, "Optimal fiscal and monetary policy in a medium-scale macroeconomic model," *NBER Macroeconomics Annual*, 2005, *20*, 383–425.

Statistics Netherlands, *Standard Industrial Classifications (SIC) 2008* Centraal Bureau voor de Statistiek 2024. Accessed on April 24, 2026.

_ , "Werknemers naar soort CAO en SBI naar achtergrondkenmerken, 2010-2024," CBS Maatwerk (Custom Table) 2025/41 2025. Accessed on April 24, 2026.

_ , "Vakbondsleden; omvang, geslacht en leeftijd," 2026. Accessed on April 24, 2026.

Stock, James H. and Mark W. Watson, "Macroeconomic Forecasting Using Diffusion Indexes," *Journal of Business and Economic Statistics*, 2002, *20* (2), 147–162.

van de Werve, Ilka, Siem Jan Koopman, Frank Weerman, and Arjan Blokland, "Exploring the crime drop in European Union homicide rates using econometric modeling," *European Journal of Criminology*, 2026, *23*, forthcoming.

van der Velden, Sjaak, "Loonstrijd en loonontwikkeling in Nederland," Technical Report, Wetenschappelijk Bureau voor de Vakbeweging De Burcht, Amsterdam 2016. Commissioned by FNV.

Visser, Jelle, "OECD/AIAS ICTWSS database: Detailed note on definitions, measurement, and sources," *Paris: Organisation for Economic Co-operation and Development*, 2021.

Wolf, Christian K, "The missing intercept: A demand equivalence approach," *American Economic Review*, 2023, *113* (8), 2232–2269.

Appendix

A Summaries of interviews

Below are summaries of the interviews we conducted with a CBA expert at the Ministry of Social Affairs, as well as (former) representatives of the largest union in the Netherlands, FNV, the employers' organization AWVN. We also include a summary of an e-mail conversation with FNV regarding the formulation of the centrally demanded wage growth percentage. The content below has been verified and approved by the interviewees for publication.

A.1 Laurens Harteveld (formerly AWVN, now Ministry of Social Affairs)

Date: August 24, 2022.
Participants: Rixt Swierstra, Marente Vlekke, Laurens Harteveld.

Q: *How does the polder model work for CBAs, and what is the role of the Stichting van de Arbeid and the SER?*

Laurens Harteveld: "The polder model is the Dutch consultation model in which employers, trade unions, and the government discuss and coordinate policy and practice regarding collective labor agreements (as well as socio-economic developments).

- **Stichting van de Arbeid (Foundation of Labor):**
  - Social partners (trade unions and employers' organizations) meet daily in thematic working groups.
  - **Composition:**
    - Trade union federations: FNV (largest), CNV, and VCP (smallest).
    - Employers' organizations: VNO-NCW (large companies), MKB (medium and small businesses), and LTO (agriculture).
  - The **Regiegroep** [Steering Group] discusses topics at a managerial level, exchanges information, and finalizes decisions. It consists of the Stichting van de Arbeid plus the Director-General of Employment from the Ministry of Social Affairs and Employment (SZW).
  - **Frequency of Regiegroep meetings:** Normally about once a month, but during crises (e.g., 2022), daily.
- **SER (Social and Economic Council):**

- Similar to the Stichting van de Arbeid, but also includes **Crown members** (independent experts, one-third of the members), appointed by the cabinet.
- The Stichting van de Arbeid consists solely of social partners."

Q: *What is the role of the FNV's centrally demanded wage growth percentage in negotiations after Prinsjesdag? How are changes in the economic outlook handled?*

Laurens Harteveld: "The centrally demanded wage growth percentage of FNV, published right before Prinsjesdag, forms the basis for all negotiations in the subsequent 12 months. This claim used to be translated into sector-specific demands. But in the past years, it has become more of a national wage claim that applies to all sectors. As the year progresses, new economic figures (e.g., inflation, productivity) are incorporated into the negotiations. If the economic situation changes significantly (e.g., higher inflation), unions will use this to demand higher wages, while employers will point to disappointing figures. FNV coordinates strategically: they encourage sectors where they are strong to reach successful agreements first, so these can serve as a reference for later negotiations for CBAs for which the willingness to go on strike is lower."

Q: *Is the FNV's centrally demanded wage growth percentage also adopted by CNV and VCP?*

Laurens Harteveld: "FNV is now the only trade union federation with a centrally demanded wage growth percentage. CNV used to write a labor terms memorandum until three years ago but no longer does so. They do formulate a 'wage request'. This is not a uniform percentage, but a range in which negotiators can determine their demands. They focus on sector-specific customization rather than a centrally demanded wage growth percentage. VCP does not publish a labor terms memorandum either."

Q: *Is the FNV's centrally demanded wage growth percentage always at least equal to the CPB August forecast of the minimum of CPI at constant tax rates inflation and the wage margin?*

Laurens Harteveld: "Yes, FNV aims for at least purchasing power preservation and uses the CPI inflation at constant tax rates forecast as a lower limit. The centrally demanded wage growth percentage is always equal to or higher than the minimum of the CPI inflation at constant tax rates forecast and the computed wage margin."

Q: *What is the exact formula for the wage margin according to the FNV?*

Laurens Harteveld: "Until 2013, the FNV used an objective formula: wage margin = CPB forecasts of productivity growth + GDP deflator – secondary labor conditions. They also took into account employment, economic growth and labor share dynamics. The FNV board then translated the wage margin into a wage claim. Since 2013, this formula has been abandoned. The FNV members' parliament now determines the wage claim. The formula is still reviewed but is no longer leading. This has led to higher centrally demanded wage growth percentages. This has led to higher wage claims that have been met with significantly less success."

Q: *How does the policy cycle around Prinsjesdag and CBA negotiations proceed?*

Laurens Harteveld: "Before Prinsjesdag: FNV publishes a draft version of the 'labor terms memorandum' ('Arbeidsvoorwaardenagenda') in which they present the centrally demanded wage growth percentage. In the autumn the agenda is definitively adopted. December marks the start of the CBA negotiations. From January 1st onwards, FNV starts coordinating the timing of negotiations strategically. Sectors with promising negotiations go first, so their agreements can serve as examples for other sectors. During the remainder of the year, new CPB forecasts are taken into account only to a certain extent, as the centrally demanded wage growth percentage remains the focus of the negotiations."

"FNV negotiates CBAs covering 5.4 million employees, CNV 4.6 million, and VCP/De Unie 3.1 million. In many CBA negotiations, multiple unions are involved. In larger CBAs, all three unions (FNV, CNV, VCP) are often involved."

Q: *How does the process of translating the labor conditions memo into concrete CBA agreements work?*

Laurens Harteveld: "In the past, the centrally demanded wage growth percentage was translated into sector-specific demands, depending on local conditions, such as labor market tightness and profitability. However, in recent years, this translation has disappeared, and the demanded wage growth in each sector is now equal to the centrally demanded wage growth percentage. The employers' organization AWVN is involved in negotiations for affiliated employers/sectors with the unions, by providing information, advice, and/or participating in consultations as a negotiator. The national CBA coordinator (FNV) checks whether the agreements in the new CBA align sufficiently with the labor conditions memo. If not, the FNV negotiator must return to the negotiating table. FNV coordinates its efforts to conclude CBAs with higher wage increases first, so they can serve as examples in other negotiations."

Q: *How are new CPB forecasts used? Are forecasts from other institutes also used?*

Laurens Harteveld: "CPB forecasts are the primary source for both employers and unions, as both parties use the same data. This ensures a common language. Since the CPB no longer publishes sectoral forecasts, data from commercial banks are used to substantiate sectoral differences."

Q: *What economic data would you use if you were to build a model to explain CBA wage developments?*

Laurens Harteveld: "Primarily labor productivity growth at the macro level and the GDP deflator. Of secondary importance are economic growth, labor market tightness, union strength, union membership rate, inflation, inflation net of taxes, the labor share."

Q: *How do AWVN and FNV differ in their views on wages and inflation?*

Laurens Harteveld: "AWVN emphasizes company-specific circumstances and labor productivity. Customization is the key focus. They believe wages do not need to follow inflation, as addressing price increases is the government's responsibility (e.g., through tax measures). They focus on investing in secondary employment conditions (e.g., training, pensions). FNV demands purchasing power preservation and refers to the CPI inflation net of taxes. They believe wages structurally lag behind and seek higher compensation for rising living costs."

Q: *Why are multi-year agreements often made? What estimates are they based on?*

Laurens Harteveld: "There are advantages for employers: predictable wage costs, avoiding annual negotiations, and spreading cost increases. Employees are given certainty about their wage development, especially in sectors with strong unions."

Q: *In multi-year agreements, we often see a high percentage in year 1 and lower percentages afterward. Do they negotiate a cumulative percentage or an average?*

Laurens Harteveld: "No cumulative percentage, but individual annual percentages. The first year often aligns with the centrally demanded wage growth percentage; later years are the result of negotiations. For year 1, the most recent CPB forecast is used. For later years, they look at long-term forecasts (e.g. inflation, productivity). Employers want gradual increases, employees want the highest increase upfront."

Q: *Do sectors with APC (automatic price compensation) have different characteristics?*

Laurens Harteveld: "Yes, sectors with APCs often have the following characteristics:

- Small to medium-sized companies (e.g., Rotterdam port, food industry in Zaanstreek).
- High organization rate and strong trade unions (e.g., FNV Havens).
- Capital-intensive sectors where the labor factor weighs less heavily in production costs.
- Two types of APC:
    1. **Structural:** Wages increase structurally and automatically with inflation, often as of January 1st based on October figures from the previous year (e.g., Boluda Towage, Cargill).
    2. **One-time:** Only applies to the current CBA, with no guarantee for the future (e.g., Painters, UMC). Often, the APC agreement only applies during the current term because it involves a long-term CBA (e.g., 4 years), and trade unions are only willing to agree to such a long term if there is at least a purchasing power guarantee."

Q: *Do you see differences between sectors that do or do not agree on nominal amounts?*

Laurens Harteveld: "Nominal amounts are mainly found in sectors with many low-skilled employees (e.g., logistics). Trade unions want to strengthen the lower end of the wage structure (e.g., an additional €100/month). Employers want to maintain the wage structure and are reluctant to agree to percentage increases that disrupt wage ratios. The more nominal amounts are used, the more the wage structure is compressed, and differences at the lower end of the wage structure disappear. This removes or reduces the incentive for employees to advance or seek promotions. If one sector agrees on a nominal amount, others often follow."

Q: *How is a nominal amount determined? What factors are taken into account?*

Laurens Harteveld: "Usually €100 per month as a starting point, aimed at the lowest pay scales. [Post-interview note by Laurens: This refers to the years 2021 and 2022.] Negotiation points:

- **Economic:** Reasonable wage increase vs. affordability for employers.
- **Tactical:** Aligning with previous agreements in other sectors ("competition").
- **Union strength:** Strong unions secure higher amounts.

The outcome is often a combination of a nominal amount + percentage increase (e.g., €100 + 2%)."

Q: *How do employers view nominal amounts?*

Laurens Harteveld: "Employers are not fans of nominal amounts, because they disrupt the wage structure. Percentage increases maintain the ratios between wages, while nominal amounts disproportionately strengthen the lower end. However, they still reach agreements because:

- Trade unions make demands (e.g., €100/month) and threaten with actions.
- Employers prefer agreeing on a nominal amount rather than an APC, which could be more expensive in the long term.
- Trend-following: If other sectors agree on nominal amounts, employers follow to maintain their competitive position.
- Employers often also recognize that employees at the lower end of the wage structure struggle to keep up during times of high inflation."

Q: *What is the difference between firm CBAs and industry CBAs, and how does this affect negotiations?*

Laurens Harteveld: "Firm CBAs are often influenced by foreign parent companies, especially in international concerns. These companies have less room to adapt to the Dutch context. Negotiations can then be more difficult because decisions are sometimes made elsewhere. For industry CBAs, large companies within a sector often want to offer good labor conditions. Small firms within the same sector count every euro and strive for minimal wage increases. This can lead to conflicting interests within the same CBA negotiations, especially if large and small firms fall under one CBA."

Q: *Do you think the 'polder model' has weakened in recent years?*

Laurens Harteveld: "Yes, the polder model has weakened over the past 10 years due to:

- **Trade unions:** Fewer members → less bargaining power and fewer opportunities for strikes.
- **Employers:** Increasingly part of foreign parent companies, which are less open to the Dutch polder model, the economic situation in the Netherlands (often using a global spreadsheet for compensation), and which want nothing to do with trade unions.

- **Government:** Weaker position, less strong leadership due to more frequent cabinet changes and long periods of caretaker status.

However, parties still cling to it because:

- All parties have become weaker, but together they are still strong.
- There is no alternative for a negotiating partner with perseverance (through the CBA, social partners can implement central agreements in practice on the work floor).
- The model offers predictability and stability, even as its influence declines. And it benefits all involved."

Q: *How is the organization rate (union members per sector) used in negotiations?*

Laurens Harteveld: "The organization rate (the percentage of employees who are union members) is crucial for the bargaining power of trade unions. Sectors with a high organization rate (e.g., ports, education) can make higher demands and have more success in strikes."

Q: *Why do we see higher wage increases in January than in December, and how does the FNV play a role in this?*

Laurens Harteveld: "CBA contracts concluded in December (effective January 1) often have lower wage increases because the negotiations have not yet been fully influenced by FNV coordination. Starting January 1, the FNV begins strategic coordination. They first focus on sectors where they are strong (high organization rate, significant dissatisfaction) to secure higher wage increases. These agreements serve as a reference for other sectors. If a CBA results in wage increases that are too low, the FNV labor conditions coordinator can reject it. Negotiators must then return to the negotiating table."

Q: *How do parties handle wage reductions that are 'hidden' in CBAs?*

Laurens Harteveld: "Wage reductions rarely occur openly but are often hidden in:

- Fewer vacation days [Post-interview note by Laurens: In recent years, there have been more free days rather than fewer].
- Lower allowances (e.g., overtime, irregular hours).

These measures are not reflected in official wage increase figures but still represent cuts for employees."

Q: *Why do high-profit sectors (e.g., the financial sector) prefer bonuses over structural wage increases?*

Laurens Harteveld: "Bonuses are more flexible: They do not need to be paid annually. Structural wage increases are fixed and can become a heavy burden in bad years. Capital-intensive sectors (e.g., banks, oil) have large buffers but want to prevent wages from rising too high in difficult times. Bonuses provide a safety net."

A.2 Interview with Jan de Kramer (AWVN)

Date: July 17, 2023.
Participants: Matthijs Katz, Marente Vlekke, Jan de Kramer.

Q: *What is your primary goal in the negotiations?*

Jan de Kramer: "There are two types of people at the table. The labor conditions advisor sits on the employers' side and is neutral—we emphasize the process. It's about achieving something between a gain-gain and a win-win outcome, as in the Harvard Negotiation Method. You have a sustainable relationship with the union; you never focus solely on content. Without that relationship, no sustainable outcome is possible. The outcome must provide continuity and certainty and prevent competition between sectors by setting a lower limit. Both parties win to some extent—you want to reach an agreement, and you often end up with an average. The question is: Am I satisfied with the outcome? Can we come to an agreement together?"

Q: *Do you formulate a maximum wage demand or a range beforehand?*

Jan de Kramer: "There is no range anymore. That used to be the case: AWVN secretaries in each sector would come up with a target figure per sector. Now, it's up to the individual negotiators. They try to mitigate and advise the representatives. They respond to the union's demands—the union's demands are the starting shot."

Q: *How do CBA negotiations generally proceed? Does the process vary greatly or minimally between different CBAs?*

Jan de Kramer: "Unions must first ask their members what they want to be addressed at the CBA table. Then, they come to the table with a letter of intent—that's statutorily determined. There's a wish list from the union, and typically, the employer formulates a response to that. Then, each point from the union is discussed: everything gets a yes, no, or maybe. That's a fairly confrontational way of negotiating. AWVN advocates for being

proactive: as an employer, formulate what you want instead of being reactive. Another approach is to present the themes that you, as an employer, consider important and discuss those. This generally leads to less unrest and confrontation."

Q: *Are there different patterns visible in different sectors?*

Jan de Kramer: "At the sectoral level, there has been a change since the yellow vests [demonstrations], think of schools, builders, and farmers. Sectors that were caught between the wheels during COVID now have higher wage demands and a more confrontational stance. In education, healthcare, and especially retail, members feel they've been strung along for years. Labor market tightness plays a role, but it's mainly in sectors that have recovered strongly post-COVID and where there's a long-standing feeling of being underpaid. In certain segments of the manufacturing industry, operators in the lower middle layer actually earn a lot, partly because production is easily measurable, and companies are more vulnerable to strikes."

Q: *What are often the bottlenecks in negotiations?*

Jan de Kramer: "Economic downturns, high inflation, increases in energy costs, tight housing markets, when people really feel it, that sharpens the dynamics. Companies with foreign owners also often just talk about money. With foreign ownership, you see more strikes and conflicts. The 'polder' factor is missing there."

Q: *Has the negotiation process changed significantly over the past twenty years?*

Jan de Kramer: "The representativeness of unions is increasingly being questioned. This also muddies the waters during negotiations. This is related to a shift in perspective on the employer's side. When the polder model was established, it was about which movement the unions represented—liberal, Christian, etc. That isn't the case anymore. But the idea that unions do not represent younger employees is increasingly being discussed."

Q: *What do you think are the most important factors that determine a negotiation outcome?*

Jan de Kramer: "Good working relationships are crucial. Large companies have professionals who deal with this. Smaller companies don't have the time. Additionally, the real threat from unions, measurable results — this is why the industry is more successful than services — the degree of organization, and keeping employees involved play a major role. But perhaps even more important than bargaining power is negotiation competence: in companies with a good, sustainable working relationship with unions and a cooperative

approach to CBA negotiations, you see faster outcomes and higher satisfaction among all parties."

Q: *How do you think changing figures and forecasts of economic trends, unemployment, and labor market tightness influence the negotiations? How do you think changing figures and forecasts of inflation influence the negotiations? Which is more important? Are these even important variables?*

Jan de Kramer: "Labor market tightness plays a very large role. It influences employers' willingness to increase wages more than they otherwise would. But it also plays a role that in many sectors, there are CBA agreements, but no company adheres to them—take the hospitality industry: every hospitality employer pays much more than the CBA prescribes. The most important series for your estimate are mainly labor market tightness and inflation. Negotiations are ultimately a stabilization mechanism. They are part of a broader system of economic stabilization."

Q: *Does bad economic news have more impact than good news?*

Jan de Kramer: "Fear of loss, i.e. the idea that people have loss aversion, always plays a role. During an economic expansion, employers say they want to invest profits and not use it for wage increases. During downturns, there's also no room for wage increases. So it's always difficult."

A.3 Interview with Pim Paulusma and Robbert Coenmans (FNV)

Date: August 16, 2023.
Participants: Matthijs Katz, Marente Vlekke, Pim Paulusma, Robbert Coenmans.

Q: *Is there a predetermined order to collective wage bargaining, for example, does the Metal CBA always take the lead?*

Pim Paulusma: "No, that's a misunderstanding. At the start of each season, after establishing our centrally demanded wage growth percentage, we look at which contracts are on the agenda. We identify the dossiers we believe will be influential or where something 'exciting' is expected to happen. The dynamics differ per sector."

Q: *What factors make a CBA trendsetting?*

Robbert Coenmans: "Size and visibility are the most important criteria. Nationally, we focus on major contracts such as Metal, Construction, Transport & Logistics, and Supermarkets. These are influential due to the large number of employees, their national recognition, and often a tradition of strong labor relations. Success in these sectors directly impacts the rest."

Pim Paulusma: "Sometimes, this leading role is sector-specific. For example, in healthcare, the Hospital CBA is important for the rest of the sector. It's more challenging in sectors like ICT, where there are many individual company CBAs instead of one large sectoral agreement. Additionally, timing is crucial: a CBA might be influential on paper, but if the contract is still valid for another year, there's little to negotiate at that moment."

Q: *Has this method of coordination always been the same?*

Pim Paulusma: "In the past, the structure was more rigid. Nowadays, we reassess each year which CBAs are most influential at that time. Be careful when analyzing data: the closing date you see in the statistics doesn't always reflect the entire process. There can be a significant gap between the expiration date and the moment a signature is actually put on the agreement."

Q: *If you were in the CPB's position, which macroeconomic variables would you use to predict the outcomes of CBA negotiations?*

Pim Paulusma: "Profitability, the labor income share, productivity, and inflation are naturally used for the official formula [of FNV's centrally demanded wage growth percentage], as well as labor market tightness. Tightness and inflation in recent years have made it possible to negotiate higher wages."

Robbert Coenmans: "It's often about the *change* in variables. Think about the change in purchasing power."

Q: *Are there factors that limit bargaining power?*

Pim Paulusma: "Rising unemployment has a direct dampening effect. The fear among employees that they could be 'next' to lose their job tempers their willingness to negotiate. Additionally, there's always room for customization. If a specific company is demonstrably struggling, we take that into account. However, I should add: for industry CBAs, there's always an employer with a hard-luck story. It's up to us to determine what the actual economic capacity is."

A.4 FNV's explanation of the wage margin and centrally demanded wage growth percentage

The statement below is a summary of multiple email exchanges in 2026 between the researchers and Robbert Coenmans and Pim Paulusma. Compiled by the authors, it has been approved for inclusion in this document by Robbert Coenmans and Pim Paulusma on April 21, 2026.

"The wage margin is the capacity employers have to increase wages without reducing the labor share (the ratio between labor and capital in total earned income). The wage margin serves as a guideline for the final centrally demanded wage growth percentage, which is the percentage wage increase that the FNV (or other unions) demands during CBA negotiations. The centrally demanded wage growth percentage is always at least equal to inflation to ensure employees do not lose purchasing power. Additionally, the FNV considers the interests of employees, and economic conditions play a role: if the economy is performing well, profits are high, and the labor market is tight, the wage demand may exceed the wage margin.

Until 2022, the wage margin for the upcoming year was calculated using the August forecast from the CPB for labor productivity and the price increase of value added. The wage margin was calculated by multiplying these two percentages. For example: if labor productivity increased by 1.2% and the price of value added increased by 2.0%, the wage margin would be $1.012 \times 1.020 = 1.0324$, or 3.2%.

In response to the unprecedented high inflation in 2022, the centrally demanded wage growth percentage for 2023 was set equal to the year-on-year CPI inflation of August 2022 (14.3%) to ensure employees were compensated for the extreme inflation caused by rising energy prices.

In 2023, we published a range for 2024 for the first time: between 5% and 14%. This was primarily due to the observation that in many CBAs, the high inflation had not yet been addressed at all. This was also largely because many CBAs were not renegotiated in 2023. If we had not accounted for this, we would have accepted a significant loss of purchasing power. Therefore, we looked at inflation in both 2022 and 2023 and the wage increases that had already occurred in that period within the respective CBAs. This range was essentially a formula that resulted in different wage demands at different negotiation tables, depending on how much salary increase people had already received. Here, we still used regular inflation, not derived inflation, inflation cleaned for changes in tax rates, to avoid confusion.

Starting from 2024 (the centrally demanded wage growth percentage for 2025), we have returned to using a formula. This formula also focuses on reducing wage inequality. The main components are: 1) the derived inflation from the previous year; 2) the development

of labor productivity; 3) an additional adjustment to bring the labor share back to a healthier level over a number of years.

That said, we are closely monitoring the situation in Iran and how prices and the economy are being managed there. It is possible that, as a result, we may adjust the centrally demanded wage growth percentage for next year slightly differently.

An important note regarding all these calculations is that the centrally demanded wage growth percentage can always be adjusted by our members. It must be approved by the Bondsraad ['Federation Council'], previously by the Ledenparlement ['Members' Parliament']. In theory, the centrally demanded wage growth percentage can be adjusted both upward and downward."

B Data details

B.1 CBA wage data example

Table 8 shows the panel data structure in long format. Table 9 shows a snippet of the same data in state space format. In this representation, each CBA (i) is assigned a dedicated row in the observation vector. Quarters in which no new agreement was reached for a specific CBA are treated as missing observations (indicated by "–").

Table 8: Selected CBA data in long format

CBA ID	CBA Name	CBA type	SBI	Agreement date	Start date	Expiration date	$\Delta w^p_{i,s,t}$
405	Hairdressing	Industry	S	2008-02-14	2008-01-01	2009-06-30	3.3
405	Hairdressing	Industry	S	2009-11-17	2009-07-01	2010-06-30	0.5
405	Hairdressing	Industry	S	2010-10-21	2010-07-01	2011-06-30	0.6
405	Hairdressing	Industry	S	2014-06-19	2014-07-01	2016-06-30	1.3
405	Hairdressing	Industry	S	2018-01-24	2017-07-01	2019-06-30	1.0
405	Hairdressing	Industry	S	2019-05-16	2019-07-01	2020-01-31	0.5
405	Hairdressing	Industry	S	2022-05-09	2022-07-01	2023-06-30	4.5
405	Hairdressing	Industry	S	2023-04-24	2023-07-01	2024-06-30	7.3
405	Hairdressing	Industry	S	2024-04-22	2024-07-01	2025-06-30	7.0
662	Heineken	Firm	C	2008-11-17	2008-07-01	2009-09-30	2.7
662	Heineken	Firm	C	2010-03-17	2009-10-01	2010-12-31	0.6
662	Heineken	Firm	C	2011-06-09	2011-01-01	2012-12-31	1.5
662	Heineken	Firm	C	2013-05-21	2013-01-01	2013-12-31	2.0
662	Heineken	Firm	C	2014-12-05	2014-01-01	2015-09-30	1.4
662	Heineken	Firm	C	2015-12-10	2015-10-01	2017-12-31	0.9
662	Heineken	Firm	C	2017-12-12	2018-01-01	2019-12-31	1.7
662	Heineken	Firm	C	2020-02-03	2020-01-01	2021-12-31	2.8
662	Heineken	Firm	C	2022-01-17	2022-01-01	2023-12-31	1.3
662	Heineken	Firm	C	2024-02-08	2024-01-01	2024-12-31	7.0
662	Heineken	Firm	C	2024-12-11	2025-01-01	2026-12-31	2.8

Table 9: Snippet of matrix of sparse CBA wage data

	Quarterly Observation Vector ($\Delta w^p_{i,s,t}$)								
CBA ID	2005:Q1	...	2008:Q1	2008:Q2	2008:Q3	2008:Q4	2009:Q1	...	2010:Q1
405	–	...	3.3	–	–	–	–	...	–
662	–	...	–	–	–	2.7	–	...	0.6

Note: Values represent the weighted average annualized wage increase (%). The symbol '–' denotes a missing observation (no new agreement), which the state space model treats as filtered latent information.

B.2 Retroactive CBA wage increases

As shown in Table 10, a significant portion of the contracts in our sample are partially or even fully retroactive, meaning that the start date of at least one agreed-upon wage increase precedes the actual date of the agreement. This raises the question of whether our primary independent variable—labor market tightness lagged by two quarters—remains a relevant predictor for wage increases intended for the past. We maintain the two-quarter lag for two reasons. First, we are primarily interested in how recent labor market conditions influence negotiated wages. Even for retroactive periods, the final settlement reflects the bargaining power of the parties at the time of signing. Second, as Table 11 illustrates, the majority of retroactive increases have a start date within one year of the agreement date. Given the typical duration of the negotiation process, a two-quarter lag effectively captures the information set available to negotiators at the time the wage claims were formulated.[24]

Table 10: Distribution of CBA contracts by share of retroactive wage increases

Share of Retroactive Increases	n	%
No retroactive wage increases	2,767	38.5
Up to 25% retroactive wage increases	160	2.2
26–50% retroactive wage increases	1,939	27.0
51–75% retroactive wage increases	109	1.5
76–100% retroactive wage increases	2,207	30.7
Total	**7,182**	**100.0**

Note: Percentages may not sum exactly to 100% due to rounding.

[24]To ensure our results are not sensitive to this timing, we perform a robustness check using a four-quarter moving average of tightness, which yields consistent results, see Appendix E.2.4.

Table 11: Timing of first wage increase relative to CBA agreement date

Timing relative to agreement date	n	%
Retroactive, 1st increase more than one year before	368	5.1
Retroactive, 1st increase within a year	4,059	56.5
Immediate (Agreement date = date 1st increase)	175	2.4
Future, 1st increase within a year	2,535	35.3
Future, 1st increase more than one year after	45	0.6
Total	**7,182**	**100.0**

Note: Percentages may not sum exactly to 100% due to rounding.

B.3 Constructing time-weighted average CBA wage growth

As explained in the main text, there is considerable heterogeneity in contract length and the number of wage increases across CBAs. To make these heterogeneous growth rates comparable across different agreements, our measure accounts for both the compounding effect of staggered raises and their specific temporal distribution within the contract. We proceed as follows. First, we construct a cumulative wage index (I_k) for each scheduled increase, acknowledging that subsequent raises apply to a previously increased base salary. For a given contract period p consisting of n discrete wage increases, let v_k be the k-th percentage wage increase occurring at time t_k. Each increase remains in effect for a duration of d_k days, such that the total duration of the contract is $D_p = \sum_{k=1}^{n} d_k$. Second, we apply a time-weighting factor (d_k/D_p) to each index level. This ensures that the timing of a raise is accurately reflected; for instance, a wage increase implemented at the start of a contract carries significantly more weight in our calculation than an identical increase implemented near the expiration date. Finally, because contract lengths vary across industries and firms, we linearly annualize the resulting cumulative growth by the factor $365/D_p$. This normalization allows for a consistent comparison of wage growth rates across agreements with heterogeneous durations. The final calculation is expressed as

$$\Delta w_{i,s,t}^{p} = \left[\left(\sum_{k=1}^{n} \frac{d_k}{D_p} \cdot I_k \right) - 100 \right] \times \frac{365}{D_p}$$

where I_k represents the cumulative wage index at step k, accounting for compound growth:

$$I_k = 100 \times \prod_{j=1}^{k} \left(1 + \frac{v_j}{100} \right)$$

B.4 Sectoral control variables

As a robustness check, we include sectoral productivity growth $\Delta\hat{z}_{s,t}$ in addition to its aggregate counterpart Δz_{t-2} in $\boldsymbol{M}_{t-2}$. These are computed by taking the growth rate of the value added per hour worked in each sector at the SBI first digit level. The resulting data are shown in Figure 13. In addition, we include $\pi^{e}_{s,t-4}$, which measures the percentage point change in energy costs as a share of total costs per sector, again at the SBI first digit level. This variable, shown in Figure 14, is computed using annual input-output tables (for a description of the method, see Huizinga and Hendriks (forthcoming)). As Table 12 reveals, while both variables display a common trend, significant cross-sectional dispersion exists, providing the necessary variation to test the sensitivity of CBA outcomes to sectoral productivity and energy cost share shocks.

Figure 13: Sectoral productivity growth Δz_{t-2}, 2005-2025Q2

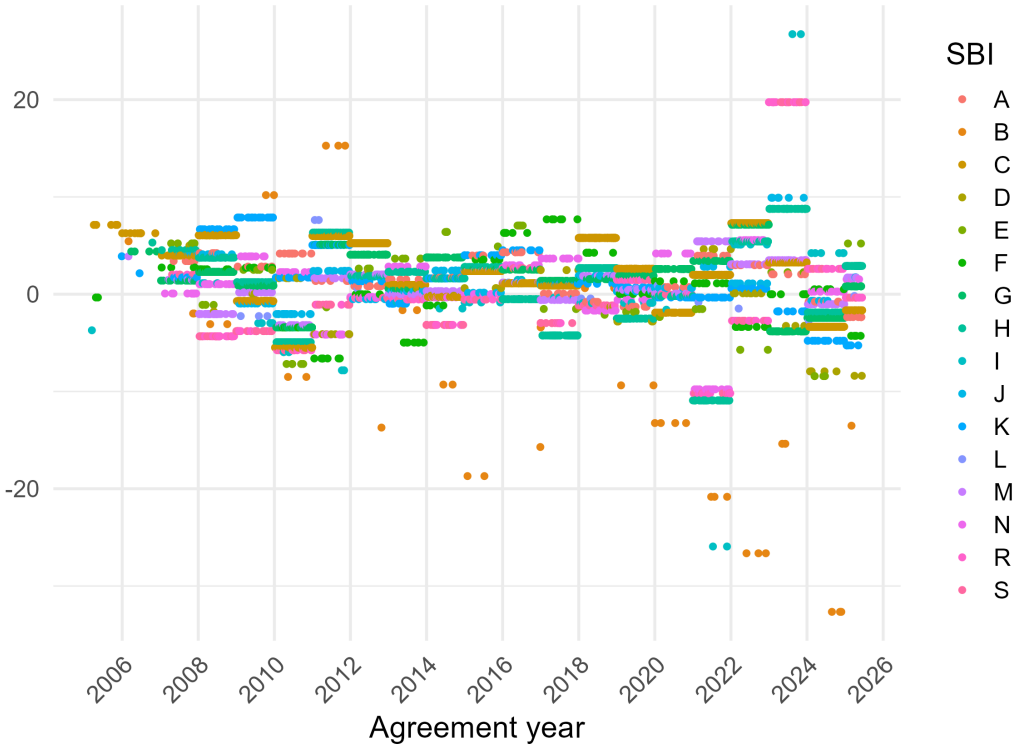


Figure 14: Percentage point change in energy costs as a share of total costs $\pi^e_{s,t-4}$, 2005-2025Q2

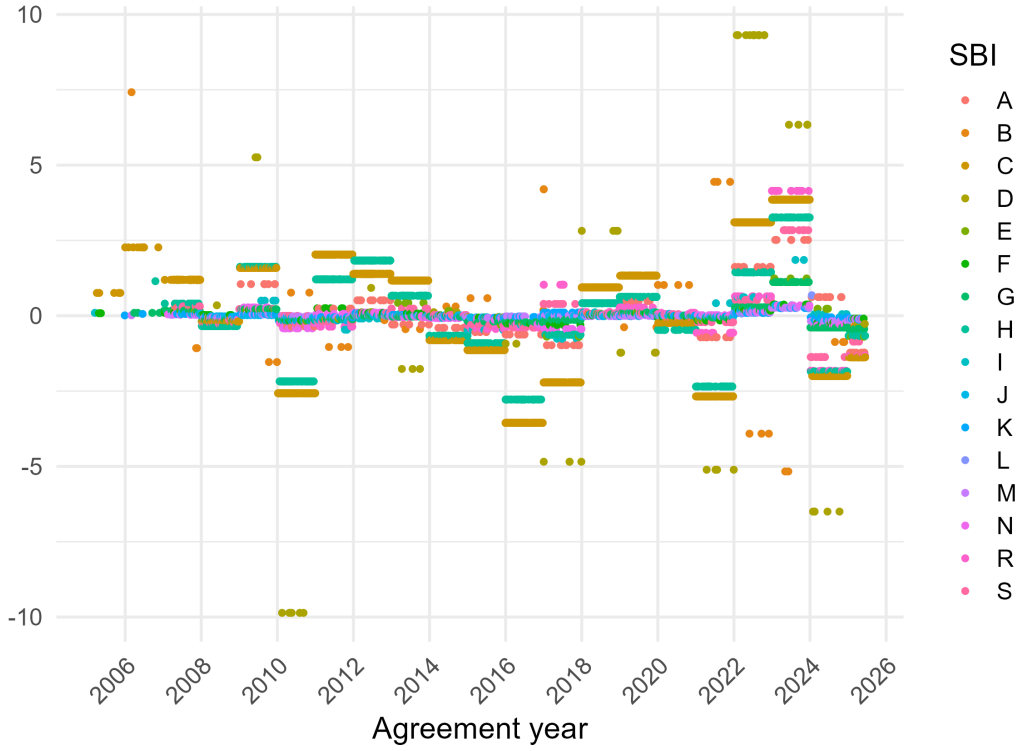


Table 12: Comparing between, common and idiosyncratic variance of $\Delta\hat{z}_{s,t-2}$ and $\pi^e_{s,t-4}$

	$\Delta\hat{z}_{s,t-2}$		$\pi^e_{s,t-4}$	
Variance	Absolute	Relative	Absolute	Relative
Total	15.84	1.00	2.81	1.00
Between	2.54	0.16	0.20	0.07
Common	6.48	0.41	1.80	0.64
Idiosyncratic	7.73	0.49	1.03	0.37

Note: Relative variance is calculated as the component variance divided by the total variance. The relative shares do not necessarily sum to unity as the covariance terms do not perfectly cancel out.

B.5 Details of the principal components analysis

Table 13: Description of macroeconomic variables used in the principal component analysis

Variable	Unit	Freq.	SA
Consumer sentiment	Y-o-Y Δ	Q (av.)	Yes
Consumers' perception of current economic conditions	Y-o-Y Δ	Q (av.)	Yes
Consumers' perception of economic conditions next 12 months	Y-o-Y Δ	Q (av.)	Yes
Consumers' perception of economic conditions previous 12 months	Y-o-Y Δ	Q (av.)	Yes
Consumption, final expenditures by households	Y-o-Y %Δ	Q	Yes
CPI, energy	Y-o-Y %Δ	Q (av.)	No
CPI, energy, net of taxes	Y-o-Y %Δ	Q (av.)	No
CPI, excluding food and energy	Y-o-Y %Δ	Q (av.)	No
CPI, food	Y-o-Y %Δ	Q (av.)	No
CPI, goods	Y-o-Y %Δ	Q (av.)	No
CPI, headline	Y-o-Y %Δ	Q (av.)	No
CPI, headline, CPB forecast of annual growth rate	Level	Q/B-A	No
CPI, net of taxes, CPB August forecast of annual growth rate	Level	Q/B-A	No
CPI, net of taxes, CPB forecast of annual growth rate	Level	Q/B-A	No
CPI, rent	Y-o-Y %Δ	Q (av.)	No
CPI, services	Y-o-Y %Δ	Q (av.)	No
Employed labor force	Y-o-Y %Δ	Q (av.)	No
Employed persons	Y-o-Y %Δ	Q	Yes
Exchange rate index	Y-o-Y %Δ	Q (av.)	No
GDP (expenditures approach)	Y-o-Y %Δ	Q	Yes
GDP (production approach)	Y-o-Y %Δ	Q	Yes
GDP deflator, CPB August forecast of annual growth rate	Level	Q/B-A	No
HICP, headline	Y-o-Y %Δ	Q (av.)	No
HICP, headline, CPB forecast of annual growth rate	Level	Q/B-A	No
Hours worked by employees	Y-o-Y %Δ	Q	No
Hours worked, total	Y-o-Y %Δ	Q	Yes
House price index	Y-o-Y %Δ	Q (av.)	No
Import index, prices	Y-o-Y %Δ	Q (av.)	No
Import index, volumes	Y-o-Y %Δ	Q (av.)	Yes
Inactive population	Y-o-Y %Δ	Q (av.)	No
Jobs of employees	Y-o-Y %Δ	Q	No
Jobs, total	Y-o-Y %Δ	Q	Yes

Continued on next page

Table 13 – continued from previous page

Variable	Unit	Freq.	SA
Labor force participation, gross	Y-o-Y Δ	Q (av.)	No
Labor force participation, net	Y-o-Y Δ	Q (av.)	No
Labor force, total	Y-o-Y %Δ	Q (av.)	No
Labor share	Y-o-Y Δ	Y	No
Productivity growth, CBS measure	Y-o-Y %Δ	A	No
Productivity growth, CPB August forecast	Level	Q/B-A	No
Productivity growth, own computation	Y-o-Y %Δ	A	No
Productivity growth market sector, own computation	Y-o-Y %Δ	A	No
Social security premiums	Y-o-Y %Δ	Q	No
Strikes: average number of employees involved	Y-o-Y %Δ	A	No
Strikes: labor days lost	Y-o-Y %Δ	A	No
Strikes: labor disputes	Y-o-Y Δ	A	No
Unemployed labor force	Y-o-Y %Δ	Q (av.)	No
Unemployment rate	Y-o-Y Δ	Q (av.)	Yes
Union members, total	Y-o-Y %Δ	A	No
Vacancies, expected growth	Level	Q (av.)	No
Vacancies, filled	Y-o-Y %Δ	Q	No
Vacancies, new	Y-o-Y %Δ	Q	No
Vacancies, open	Y-o-Y %Δ	Q	No
Vacancy rate	Y-o-Y Δ	Q	Yes

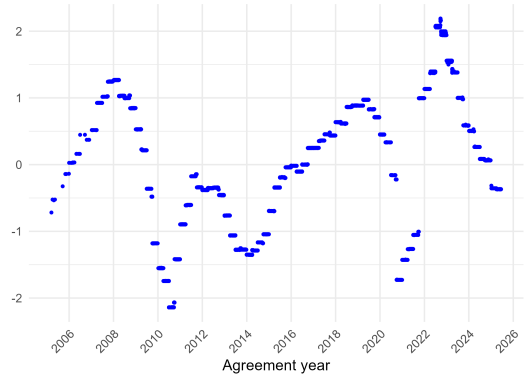
Notes: Y-o-Y %Δ denotes the year-on-year percentage change of the (average) index. Q (av.) indicates quarterly data derived from monthly averages. Q/B-A refers to a mixed frequency starting as quarterly and transitioning to bi-annual (twice-yearly) reporting. SA indicates whether the series has been seasonally adjusted by the source.

Table 14: Eigenvalues and explained variance of unrotated principal components

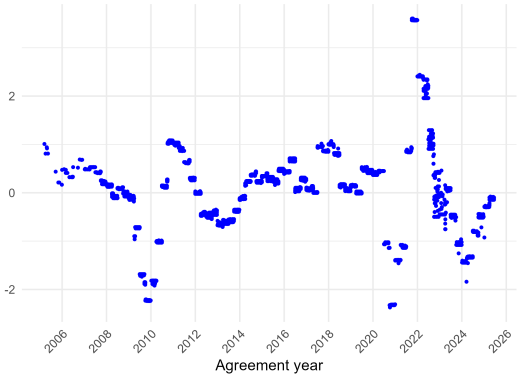
Component	Eigenvalue	Proportion	Cumulative
$\Delta\theta^{(1)}$	20.414	0.393	0.393
$\pi^{(1)}$	9.141	0.176	0.568
$\pi^{(2)}$	4.766	0.092	0.660
Δs	4.411	0.085	0.745
$\Delta\theta^{(2)}$	2.051	0.039	0.784
Δz	1.801	0.035	0.819
$\Delta u^{(1)}$	1.465	0.028	0.847
$\Delta u^{(2)}$	1.131	0.022	0.869

Notes: 'Proportion' is the share of total variance explained by each component individually, and is the running sum of the proportion column.

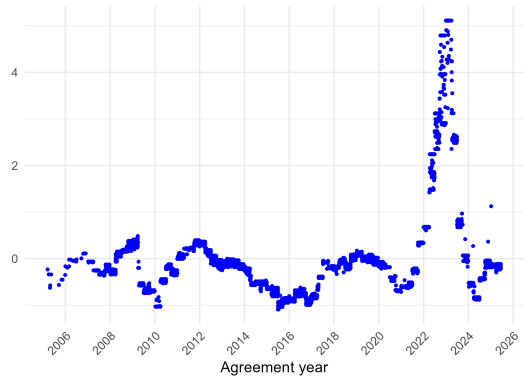
Figure 15: Principal components (1 of 2)



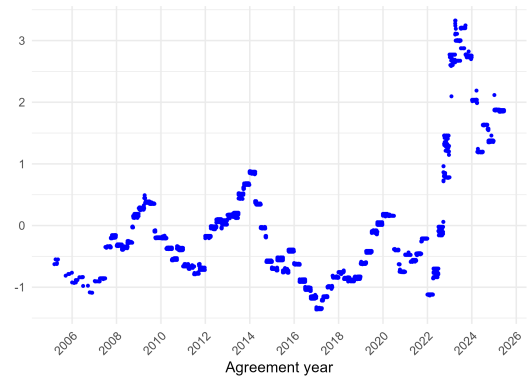
(a) $\Delta\theta^{(1)}$



(b) $\Delta\theta^{(2)}$

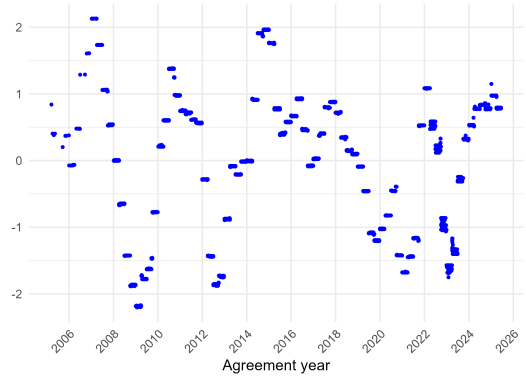


(c) $\pi^{(1)}$

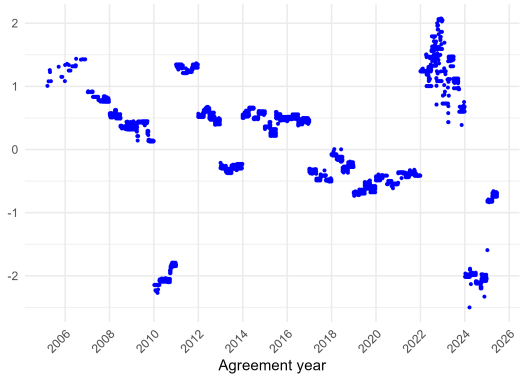


(d) $\pi^{(2)}$

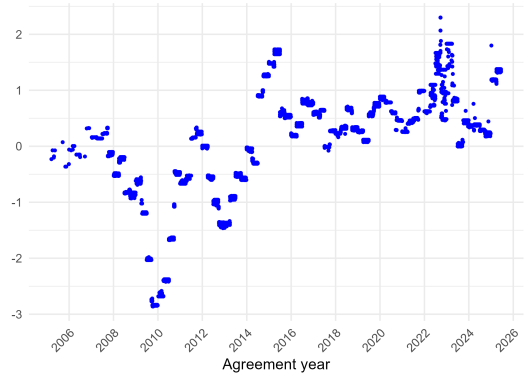
Figure 16: Principal components (2 of 2)



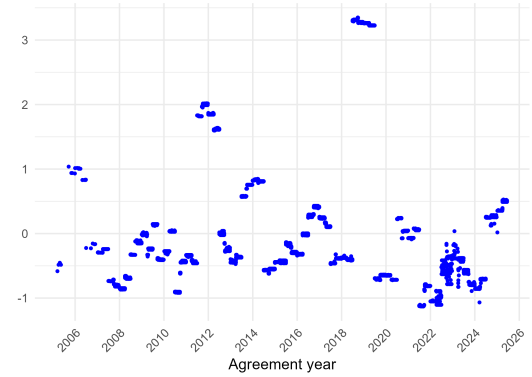
(a) Δs



(b) Δz



(c) $\Delta u^{(1)}$



(d) $\Delta u^{(2)}$

Table 15: R^2 of equations (7)

Variable	$\Delta\theta^{(1)}$	$\Delta\theta^{(2)}$	$\pi^{(1)}$	$\pi^{(2)}$	Δs	Δz	$\Delta u^{(1)}$	$\Delta u^{(2)}$
Consumer sentiment	0.038	0.077	0.194	0.055	0.917	0.063	0.012	0.016
Consumers' perception of current economic conditions	0.051	0.104	0.082	0.037	0.939	0.032	0.012	0.018
Consumers' perception of economic conditions next 12 months	0.147	0.040	0.135	0.018	0.614	0.118	0.009	0.028
Consumers' perception of economic conditions previous 12 months	0.000	0.138	0.019	0.043	0.887	0.001	0.085	0.004
Consumption, final expenditures by households	0.597	0.356	0.518	0.136	0.011	0.155	0.240	0.014
CPI, energy	0.261	0.112	0.849	0.036	0.062	0.239	0.077	0.000
CPI, energy, net of taxes	0.227	0.074	0.850	0.051	0.077	0.231	0.076	0.001
CPI, excluding food and energy	0.036	0.076	0.205	0.879	0.026	0.001	0.045	0.028
CPI, food	0.083	0.135	0.207	0.752	0.148	0.012	0.020	0.017
CPI, goods	0.273	0.020	0.898	0.269	0.123	0.218	0.086	0.004
CPI, headline	0.215	0.000	0.795	0.517	0.111	0.122	0.085	0.009
CPI, headline, CPB forecast of annual growth rate	0.162	0.003	0.648	0.397	0.050	0.006	0.082	0.006
CPI, net of taxes, CPB August forecast of annual growth rate	0.134	0.015	0.502	0.594	0.087	0.036	0.055	0.009
CPI, net of taxes, CPB forecast of annual growth rate	0.157	0.000	0.658	0.273	0.025	0.010	0.101	0.016
CPI, rent	0.209	0.012	0.086	0.055	0.093	0.024	0.042	0.000
CPI, services	0.010	0.197	0.008	0.804	0.005	0.097	0.004	0.018
Employed labor force	0.911	0.049	0.211	0.036	0.072	0.101	0.084	0.010
Employed persons	0.955	0.184	0.225	0.010	0.002	0.119	0.220	0.000
Exchange rate index	0.048	0.150	0.479	0.008	0.010	0.216	0.000	0.000

Continued on next page

Table 15 – continued from previous page

Variable	$\Delta\theta^{(1)}$	$\Delta\theta^{(2)}$	$\pi^{(1)}$	$\pi^{(2)}$	Δs	Δz	$\Delta u^{(1)}$	$\Delta u^{(2)}$
GDP (expenditures approach)	0.699	0.418	0.381	0.101	0.014	0.108	0.360	0.024
GDP (production approach)	0.716	0.380	0.415	0.128	0.009	0.120	0.319	0.028
GDP deflator, CPB August forecast of annual growth rate	0.184	0.014	0.345	0.693	0.031	0.000	0.037	0.002
HICP, headline	0.227	0.000	0.806	0.484	0.116	0.132	0.077	0.009
HICP, headline, CPB forecast of annual growth rate	0.173	0.001	0.661	0.346	0.054	0.010	0.086	0.004
Hours worked by employees	0.521	0.309	0.175	0.002	0.001	0.119	0.193	0.001
Hours worked, total	0.707	0.335	0.222	0.026	0.001	0.145	0.218	0.002
House price index	0.260	0.189	0.149	0.161	0.033	0.187	0.108	0.009
Import index, prices	0.272	0.284	0.805	0.020	0.022	0.281	0.083	0.003
Import index, volumes	0.082	0.772	0.020	0.127	0.082	0.141	0.161	0.003
Inactive population	0.605	0.000	0.229	0.219	0.167	0.072	0.003	0.033
Jobs of employees	0.935	0.250	0.235	0.001	0.000	0.124	0.180	0.000
Jobs, total	0.942	0.237	0.230	0.006	0.000	0.136	0.183	0.000
Labor force participation, gross	0.598	0.000	0.228	0.223	0.188	0.068	0.001	0.028
Labor force participation, net	0.937	0.076	0.197	0.019	0.046	0.120	0.097	0.011
Labor force, total	0.518	0.008	0.236	0.290	0.248	0.041	0.000	0.026
Labor share	0.299	0.081	0.077	0.115	0.068	0.054	0.069	0.014
Productivity growth, CBS measure	0.219	0.133	0.207	0.004	0.017	0.955	0.076	0.005
Productivity growth, CPB August forecast	0.365	0.104	0.073	0.077	0.003	0.110	0.491	0.014
Productivity growth, own computation	0.118	0.132	0.154	0.006	0.010	0.972	0.055	0.001
Productivity growth, own computation market sector	0.104	0.128	0.150	0.007	0.020	0.919	0.051	0.004
Social security premiums	0.284	0.005	0.169	0.144	0.175	0.000	0.106	0.006

Continued on next page

Table 15 – continued from previous page

Variable	$\Delta\theta^{(1)}$	$\Delta\theta^{(2)}$	$\pi^{(1)}$	$\pi^{(2)}$	Δs	Δz	$\Delta u^{(1)}$	$\Delta u^{(2)}$
Strikes: average number of employees involved	0.012	0.092	0.029	0.072	0.083	0.023	0.002	0.604
Strikes: labor days lost	0.012	0.019	0.006	0.006	0.004	0.004	0.016	0.870
Strikes: labor disputes	0.006	0.117	0.015	0.257	0.093	0.031	0.048	0.019
Unemployed labor force	0.758	0.269	0.115	0.044	0.006	0.153	0.275	0.000
Unemployment rate	0.718	0.208	0.075	0.059	0.005	0.105	0.229	0.000
Union members, total	0.103	0.099	0.095	0.034	0.025	0.033	0.737	0.055
Vacancies, expected growth	0.587	0.406	0.082	0.006	0.087	0.078	0.455	0.001
Vacancies, filled	0.568	0.508	0.188	0.025	0.076	0.312	0.473	0.001
Vacancies, new	0.253	0.828	0.060	0.088	0.230	0.093	0.391	0.003
Vacancies, open	0.452	0.728	0.190	0.062	0.111	0.206	0.448	0.003
Vacancy rate	0.478	0.692	0.316	0.027	0.053	0.243	0.413	0.003

Notes: This table presents the R^2-values pertaining to bivariate regressions (7) for each extracted principal component and the macroeconomic variables. Indices (1) and (2) refer to the leading and secondary components for each factor group.

C State space estimation technique

The most extended model with time-varying parameters is given by equation (6) which can be written into a linear Gaussian panel state space model

$$y_{i,t} = z_{0,i,t}\tau_t + z_{1,i,t}\beta_{1,t} + z_{2,i,t}\beta_{2,t} + x_t'\gamma + \varepsilon_{i,t}, \tag{9}$$

where

$$y_{i,t} := \Delta\hat{w}^p_{i,s_i,t}, \qquad z_{0,i,t} := 1, \qquad z_{1,i,t} := \Delta\theta_{s_i,t-2}, \qquad z_{2,i,t} := \Delta\theta_{t-2}, \qquad x_t := M_{t-2},$$

where i refers to CBA i, t refers to the quarter t in the CBA pair (i,t), index $s_i := s(i)$ with categorical function $s(i)$ that returns the sector index associated with CBA i, for the quarterly time index $t = 1,\ldots,T$ where T is the length of the time series in quarters. The definitions of $z_{j,i,t}$, for $j = 0,1,2$, show that different levels of pooling take place in the model: τ_t and $\beta_{2,t}$ are pooled over all CBAs, and $\beta_{1,t}$ is pooled over the sectoral level of the CBAs, at time t. Finally, we have $i = 1,\ldots,N_t$ where N_t is the number of CBAs in quarter t which can clearly differ substantially for different quarters.

To represent the model in a multivariate linear Gaussian state space form, we define

$$y_t = (y_{1,t},\ldots,y_{N_t,t})', \qquad Z_t = (\bar{z}_{1,t},\ldots,\bar{z}_{N_t,t})', \qquad \varepsilon_t = (\varepsilon_{1,t},\ldots,\varepsilon_{N_t,t})',$$

where $\bar{z}_{i,t} = (z_{0,i,t}, z_{1,i,t}, z_{2,i,t})'$ for $i = 1,\ldots,N_t$, and obtain the measurement or observation equation

$$y_t = Z_t\alpha_t + x_t'\gamma + \varepsilon_t,$$

with state vector $\alpha_t = (\tau_t, \beta_{1,t}, \beta_{2,t})'$, for $t = 1,\ldots,T$ and irregular vector $\varepsilon_t \sim NID(0,\Sigma_\varepsilon)$ where we assume $\Sigma_\varepsilon = \sigma_\varepsilon^2 \cdot I_{N_t}$ and $\sigma_\varepsilon > 0$. The dynamics of all state elements are determined by a random walk process, and therefore we obtain the transition or state equation

$$\alpha_{t+1} = \alpha_t + \eta_t,$$

which is simply a multivariate random walk with Gaussian innovations $\eta_t \sim NID(0,\Sigma_\eta)$ where we assume a diagonal variance matrix Σ_η. Hence, all elements in the disturbance vectors ε_t and η_t are serially and mutually independent of each other. The initial state vector α_1 is treated as a random vector with a diffuse prior. The econometric treatment of this model can be carried out using the state space methods discussed in Durbin and Koopman (2012, DK). A non-standard feature of our model is the time-varying dimension of the observation vector y_t. Since we adopt the collapsed Kalman filter of Jungbacker and Koopman (2015) and the univariate treatment of multivariate series (DK, Section 6.4), the methods can be applied straightforwardly. The treatment of the fixed regression

coefficient vector γ and the initial state vector α_1 is based on the augmentation method (DK, Sections 5.7 and 6.2.3). A synopsis of this treatment is given below.

The simultaneous treatment of the initialization and the regression effect is based on a rewritten (but equivalent) observation equation given by

$$y_t = X_t\beta + u_t, \qquad u_t = Z_t\alpha_t^* + \varepsilon_t, \qquad \alpha_{t+1}^* = \alpha_t^* + \eta_t,$$

where $\alpha_1^* = 0$ and $X_t = (Z_t, x_t)'$ and $\beta' = (\alpha_1', \gamma')$. We notice that $\alpha_t^* \equiv \alpha_t - \alpha_1$. For a given β, we observe $u_t = y_t - X_t\beta$ and can apply the Kalman filter to u_t with initialization $\alpha_1^* = 0$. Hence, the prediction error decomposition is applied to the time series $u_1, \ldots, u_T$, using the Kalman filter. We denote the subsequent prediction errors with their variance matrices by v_t^u and F_t^u, respectively, for $t = 1, \ldots, T$. We notice that, with correct model specification, $v_t^u \sim NID(0, F_t^u)$. The log-likelihood function can be expressed in terms of the prediction errors by

$$\ell_T(u;\beta) = -\frac{NT}{2}\log 2\pi - \frac{1}{2}\sum_{t=1}^{T}\log|F_t^u| - \frac{1}{2}\sum_{t=1}^{T} v_t^{u\prime} F_t^{u\,-1} v_t^u.$$

Since the Kalman filter is a linear projection, we can apply the same Kalman filter used for the vector u_t series, also to the vector y_t series (to obtain v_t^y) and to each series represented by the columns of X_t (to obtain v_t^X). Since $u_t = y_t - X_t\beta$, it follows that

$$v_t^u \;=\; v_t^y \;-\; v_t^X\beta\,, \qquad t = 1, \ldots, T.$$

By substituting these expressions in the log-likelihood function, we obtain

$$\ell_T(u;\beta) = -\frac{NT}{2}\log 2\pi - \frac{1}{2}\sum_{t=1}^{T}\log|F_t^u| - \frac{1}{2}\sum_{t=1}^{T}(v_t^y \;-\; v_t^X\beta)' F_t^{u\,-1}(v_t^y \;-\; v_t^X\beta),$$

from which it follows from classical inference results that the maximum likelihood estimate of β equals the generalized least squares estimate $\widehat{\beta} = S_T^{-1}s_T$ where

$$S_T = \sum_{t=1}^{T} v_t^{X\prime} F_t^{u\,-1} v_t^X, \qquad s_T = \sum_{t=1}^{T} V_t^{X\prime} F_t^{u\,-1} v_t^y.$$

The profile log-likelihood function is obtained when β is replaced by $\widehat{\beta}$ in the log-likelihood function, we have

$$\ell_T^*(y) = -\frac{NT}{2}\log 2\pi - \frac{1}{2}\sum_{t=1}^{T}\log|F_t^u| - \frac{1}{2}(q_T - s_T S_T^{-1} s_T),$$

where

$$q_T = \sum_{t=1}^{T} v_t^{y\prime} F_t^{u\,-1} v_t^y.$$

The diffuse log-likelihood is the same as this profile log-likelihood function, but it accounts for the diffuse prior for the initial state vector by adding the term $-\frac{1}{2}\log|S_T^{\alpha_1}|$ where $S_T^{\alpha_1}$ is the upper 3×3 block of S_T associated with the initial state α_1 (DK, Section 7.2.3).

We can adopt this profile (or diffuse) log-likelihood for the estimation of the remaining unknown parameters in the model such as the unknown variances σ_ε^2 and on the diagonal Σ_η. Hence, these unknown parameters are estimated by the method of maximum likelihood which requires the numerical maximization of the profile log-likelihood function with respect to these model variances. We typically do not estimate a variance directly, but we make it subject to an appropriate transformation, for example $\sigma_\varepsilon^2 = \exp(2\,h_\varepsilon)$, and estimate h_ε instead. Our panel state space model only contains 4 unknown variances, this numerical optimization task is therefore light and fast on any modern computer.

Once parameter estimation is completed, the estimation of the time-varying parameters in the state vector α_t can be carried out using Kalman filtering and smoothing methods (DK, Sections 4.2 and 4.3). The confidence intervals for the time-varying paths can also be constructed using the output of these methods. Finally, to verify whether the model specification can be justified for the datasets at hand, diagnostic residual statistics and graphs can be computed and studied. The standardized prediction errors from the Kalman filter are typically used for these tasks.

This completes the discussion of a full maximum likelihood treatment of a panel model with time-varying parameters. Despite the large-dimensional panel, we have shown that a solid statistical treatment can be provided using partly standard regression methods and partly maximum likelihood methods. Given the reliance on Kalman filter methods, all necessary computations can be carried out computationally efficiently using the suite of `C` routines from Koopman et al. (1999).

D Detailed main empirical results

Table 16: Estimates of β_1 from equation (1), standard errors not clustered

	No FE	Only CBA FE	CBA and time FE
	(1)	(2)	(3)
Full sample	0.025***	0.023***	0.003
	(0.002)	(0.002)	(0.004)
2005-2008Q3	0.003	−0.009	0.011
	(0.006)	(0.012)	(0.013)
2008Q4-2019	0.019***	0.021***	0.005
	(0.002)	(0.002)	(0.004)
2020-2025Q2	0.001	0.002	0.002
	(0.004)	(0.005)	(0.011)
CBA fixed effects		✓	✓
Time fixed effects			✓

Notes: Standard errors are in parentheses. In contrast to Table 4, they are conventional (homoskedastic) and are not clustered by CBA. FE refers to 'fixed effects'. Statistical significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 17 is the more extended version of Table 7 and also contains the results for the benchmark state space model with fixed parameters in column 1 to demonstrate that our estimates are not sensitive to the use of state space methods[25].

Our final model, the extended model with inflation interaction terms, as shown in column 5 contains an extra interaction term $\pi^{(1)}_{t-2}\mathbb{1}_{t-2\geq 2022Q3}$. We included this interaction term, as $\pi^{(1)}_{t-2}\mathbb{1}_{t-2\geq 2021Q3}$ had a negative coefficient. Since the coefficient on $\pi^{(1)}_{t-2}\mathbb{1}_{t-2\geq 2022Q3}$ is positive again, we conclude that the period between the end of the last COVID-lockdown, when inflation first started increasing, and the peak of the energy price shock following Russia's invasion of Ukraine in 2022, represents a unique structural break in the relationship between inflation and wage growth. This suggests that while initial post-pandemic inflation did not immediately translate into higher centrally negotiated wages, the surge in inflation

[25]The slightly lower maximized log-likelihood for model (2) relative to (1) is due to bounded parameter spaces and how diffuse initializations in the Kalman filter are handled. Model (1) with a fixed intercept is a limiting case of model (2) with a time-varying intercept, and hence these two log-likelihoods are not directly comparable, see the discussion in Francke, Koopman and De Vos (2010).

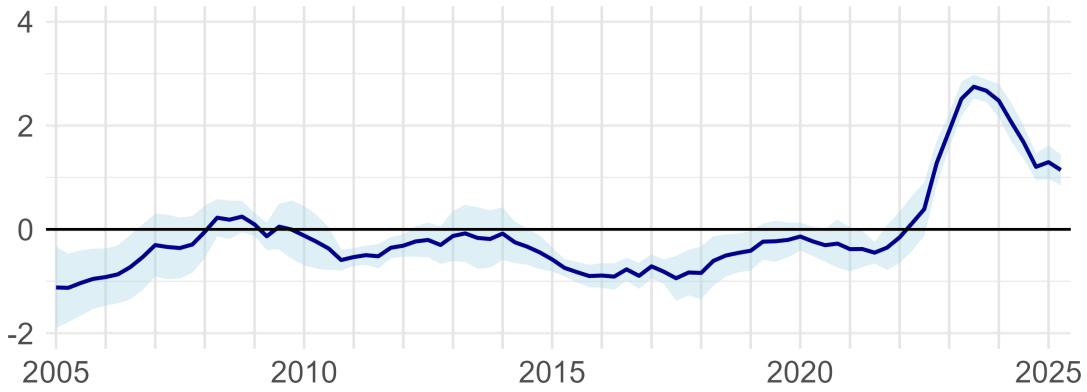
Table 17: End of sample parameter estimates and model diagnostics across specifications

	(1) Fixed par. model (1)	(2) Benchmark model (3)	(3) Extended model (5)	(4) Extended model (6)	(5) Full model (6)
Time-varying intercept at time T					
τ_T	0.806^{***}	0.832^{***}	1.144^{***}	-0.495^{**}	-0.869^{***}
	(0.096)	(0.097)	(0.161)	(0.233)	(0.175)
Sectoral labor market tightness coefficient at time T					
$\beta_{1,T}$	0.005^{*}	-0.010	-0.010	-0.008	0.003
	(0.003)	(0.017)	(0.014)	(0.015)	(0.004)
Aggregate labor market tightness coefficient at time T					
$\beta_{2,T}$			1.025^{**}	0.843^{**}	0.907^{**}
			(0.433)	(0.412)	(0.431)
Coefficients on aggregate control variables					
CDW_t				0.124^{***}	0.107^{***}
				(0.019)	(0.017)
$\tilde{\pi}_{i,s,t}^{p-1}$				0.079^{***}	0.074^{**}
				(0.011)	(0.030)
$\tilde{\pi}_{i,s,t}^{p-1}\mathbb{1}_{t-2\geq 2021Q3}$					0.005
					(0.032)
$\pi_{t-2}^{(1)}$				-0.146^{***}	0.352^{***}
				(0.053)	(0.106)
$\pi_{t-2}^{(1)}\mathbb{1}_{t-2\geq 2021Q3}$					-0.822^{***}
					(0.123)
$\pi_{t-2}^{(1)}\mathbb{1}_{t-2\geq 2022Q3}$					0.238^{***}
					(0.049)
$\pi_{t-2}^{(2)}$				0.423^{***}	0.058
				(0.082)	(0.102)
$\pi_{t-2}^{(2)}\mathbb{1}_{t-2\geq 2021Q3}$					0.638^{***}
					(0.129)
Δs_{t-2}				-0.011	-0.072
				(0.054)	(0.044)
Δz_{t-2}				0.057	-0.028
				(0.046)	(0.039)
$\Delta u_{t-2}^{(1)}$				0.029	0.042
				(0.070)	(0.054)
$\Delta u_{t-2}^{(2)}$				-0.000	-0.020
				(0.034)	(0.027)
Log-likelihood	-9812.232	-9815.871	-9784.075	-9713.702	-9684.626
AIC	0.252	0.252	0.252	0.250	0.249
$\sqrt{\text{Var}(\tau_t)}$	0.933	0.935	0.888	0.206	0.115
$\sqrt{\text{Var}(\tau_t)}$ (rel.)	0.998	1.000	0.950	0.220	0.123

Note: Standard errors in parentheses. ***, **, and * denote significance at 1%, 5%, and 10%.

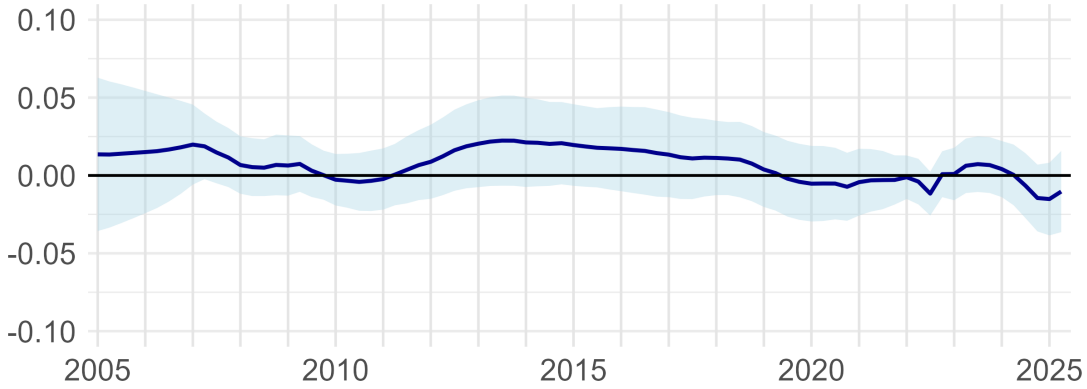
after 2021 eventually forced a stronger upward adjustment in CBA wage outcomes. Below we show the remaining graphs associated with the estimated models in Table 17 not shown in the main text.

Figure 17: Kalman smoothed estimates of τ_t in extended model (5)



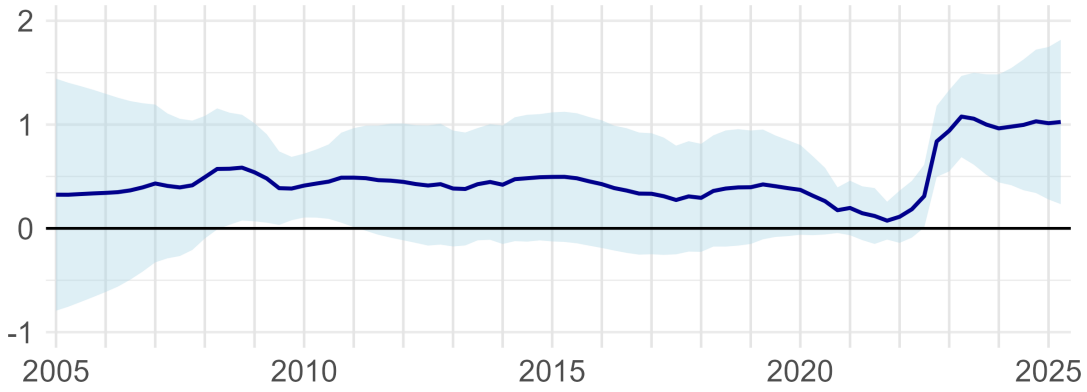
Note: The shaded area indicates the 95%-confidence interval.

Figure 18: Kalman smoothed estimates of $\beta_{1,t}$ in extended model (5)



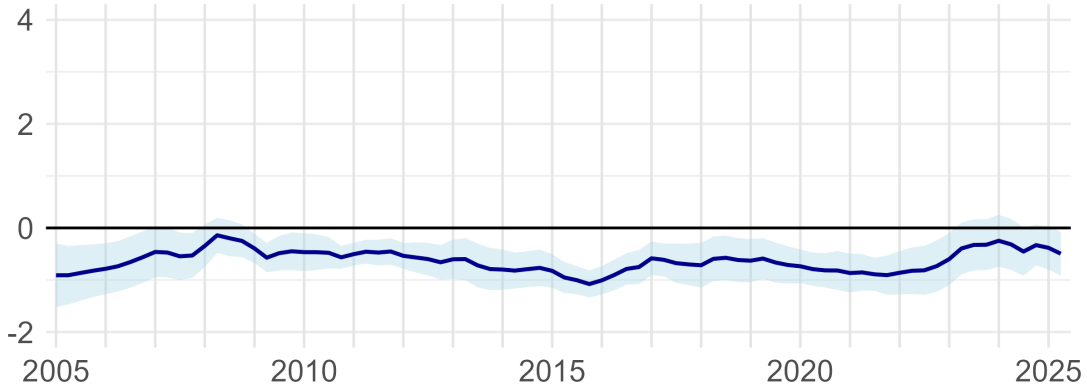
Note: The shaded area indicates the 95%-confidence interval.

Figure 19: Kalman smoothed estimates of $\beta_{2,t}$ in extended model (5)



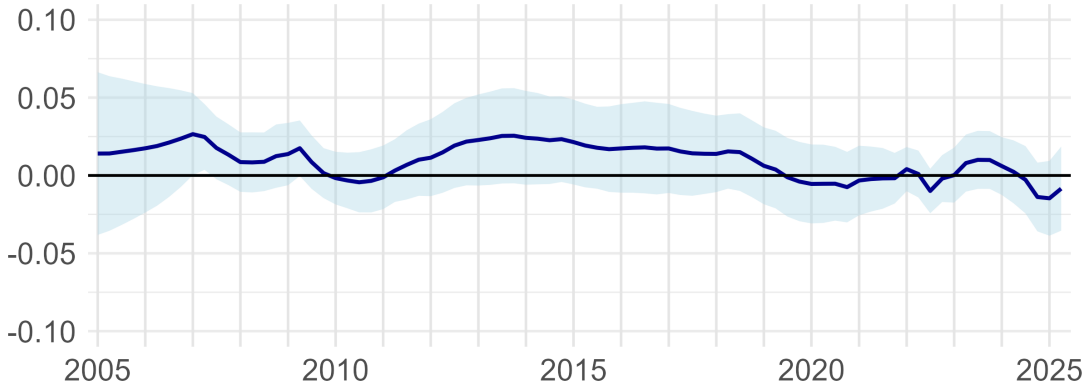
Note: The shaded area indicates the 95%-confidence interval.

Figure 20: Kalman smoothed estimates of τ_t in extended model (6)



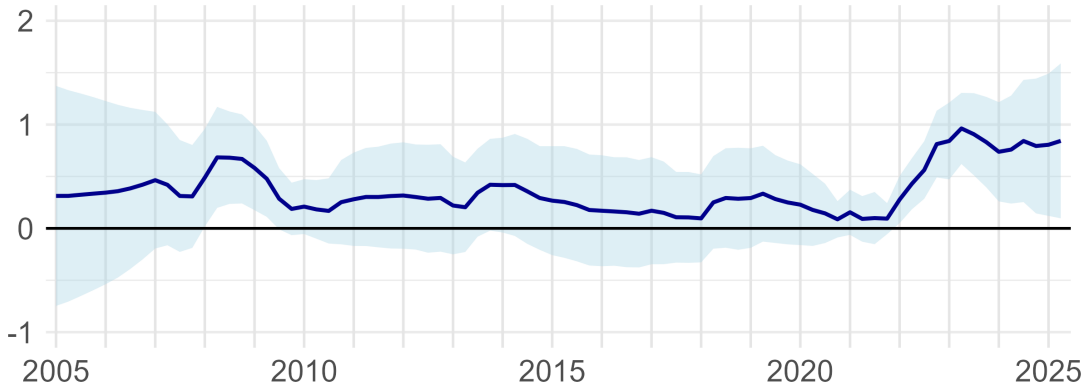
Note: The shaded area indicates the 95%-confidence interval.

Figure 21: Kalman smoothed estimates of $\beta_{1,t}$ in extended model (6)



Note: The shaded area indicates the 95%-confidence interval.

Figure 22: Kalman smoothed estimates of $\beta_{2,t}$ in extended model (6)



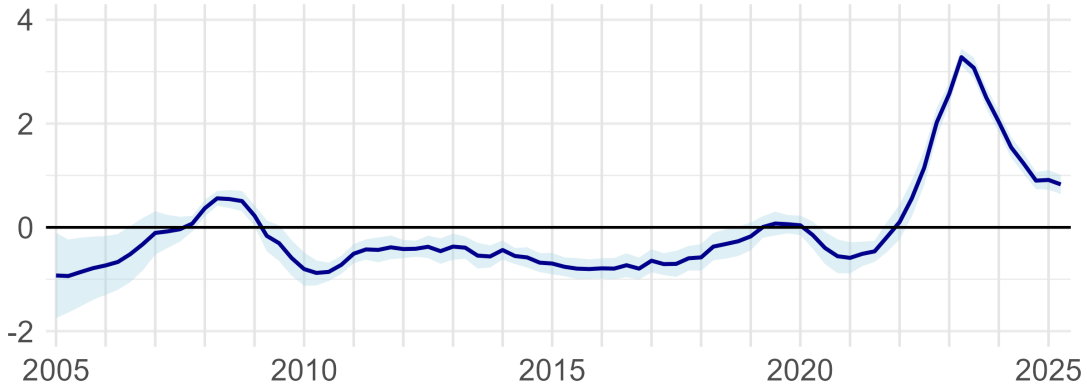
Note: The shaded area indicates the 95%-confidence interval.

E Robustness checks

E.1 Robustness checks of benchmark model (3)

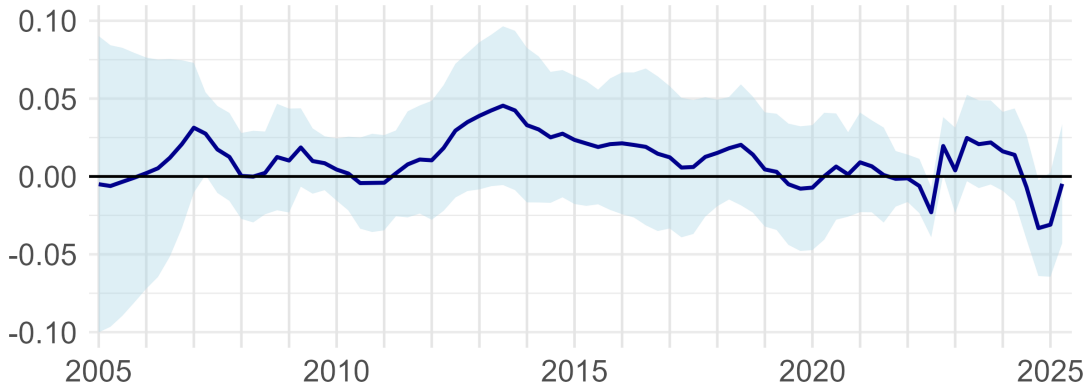
E.1.1 Estimates of benchmark model (3) for the goods and services sector

Figure 23: Kalman smoothed estimates of τ_t in benchmark model (3) for services CBAs



Note: The shaded area indicates the 95%-confidence interval.

Figure 24: Kalman smoothed estimates of $\beta_{1,t}$ in benchmark model (3) for services CBAs



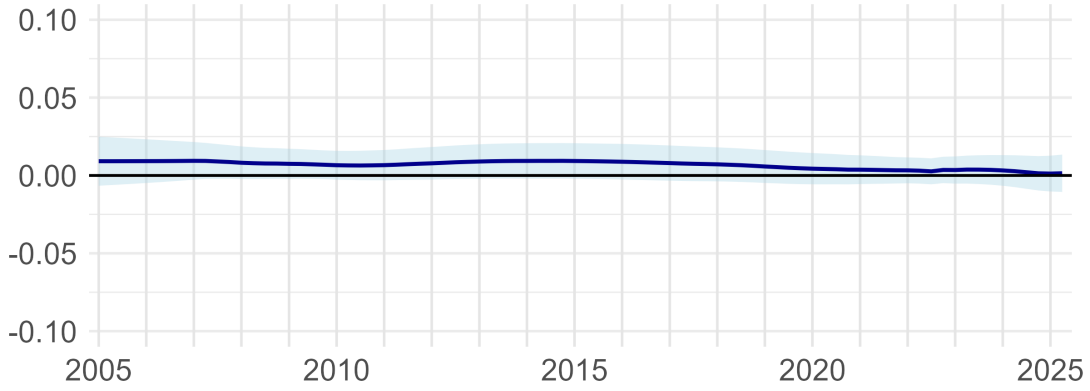
Note: The shaded area indicates the 95%-confidence interval.

Figure 25: Kalman smoothed estimates of τ_t in benchmark model (3) for goods CBAs



Note: The shaded area indicates the 95%-confidence interval.

Figure 26: Kalman smoothed estimates of $\beta_{1,t}$ in benchmark model (3) for goods CBAs



Note: The shaded area indicates the 95%-confidence interval.

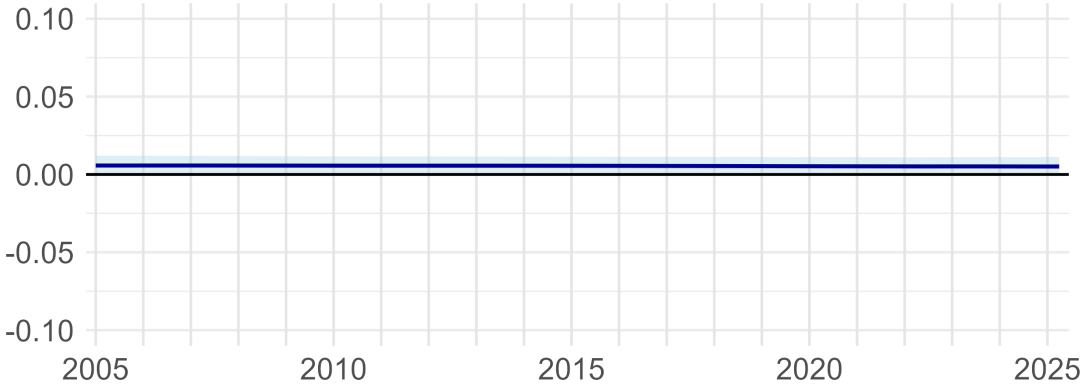
E.1.2 Estimates of benchmark model (3) for industry and firm CBAs

Figure 27: Kalman smoothed estimates of τ_t in benchmark model (3) for industry CBAs



Note: The shaded area indicates the 95%-confidence interval.

Figure 28: Kalman smoothed estimates of $\beta_{1,t}$ in benchmark model (3) for industry CBAs



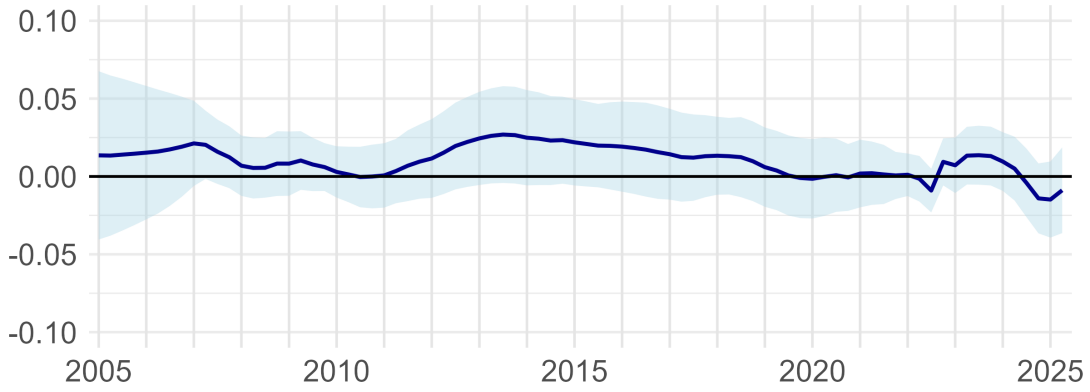
Note: The shaded area indicates the 95%-confidence interval.

Figure 29: Kalman smoothed estimates of τ_t in benchmark model (3) for firm CBAs



Note: The shaded area indicates the 95%-confidence interval.

Figure 30: Kalman smoothed estimates of $\beta_{1,t}$ in benchmark model (3) for firm CBAs



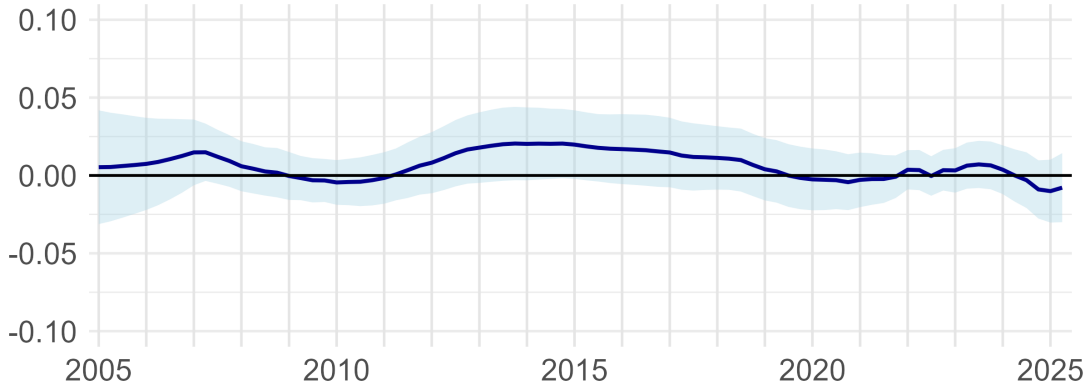
Note: The shaded area indicates the 95%-confidence interval.

E.2 Robustness checks of full model (6)

E.2.1 Estimates of full model (6) with τ_t restricted to a constant

Figure 32 shows the smoothed estimate of $\beta_{2,t}$ for a version of the full model (6) in which τ_t is restricted to a constant. The trajectory of $\beta_{2,t}$ is almost identical to the estimates of $\beta_{2,t}$ in Figure 11 in the main model where τ_t is allowed to vary over time. The estimates of $\beta_{1,t}$ also display more dynamics, but are never significantly different from 0, as in the main model. Column 2 in Table 18 shows that the log-likelihood and AIC are only marginally worse than those of the full model.

Figure 31: Kalman smoothed estimates of $\beta_{1,t}$ in full model (6) with τ_t restricted to a constant



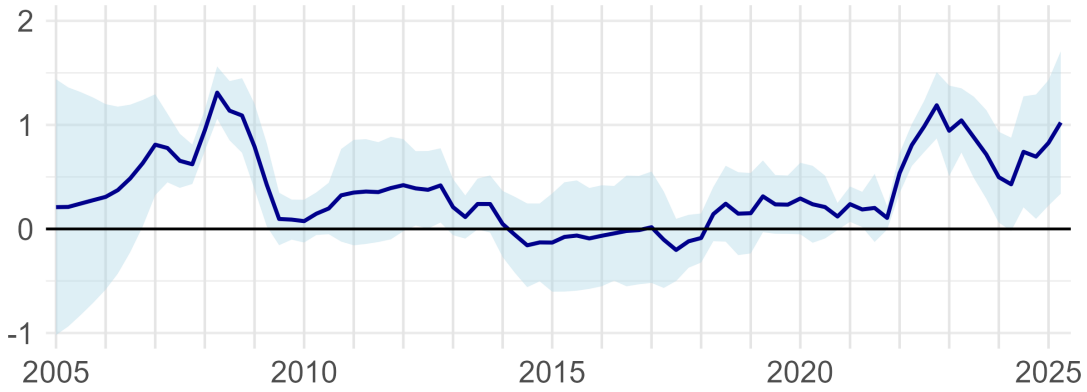
Note: The shaded area indicates the 95%-confidence interval.

Table 18: Robustness checks: comparison with the full model

	(1) Full model (6)	(2) Full model (6) with constant τ	(3) Full model (6) with sectoral controls
Time-varying intercept at time T			
τ_T	-0.869^{***}	-0.607^{***}	-0.828^{***}
	(0.175)	(0.065)	(0.178)
Sectoral labor market tightness coefficient at time T			
$\beta_{1,T}$	0.003	-0.008	0.004
	(0.004)	(0.012)	(0.003)
Aggregate labor market tightness coefficient at time T			
$\beta_{2,T}$	0.907^{**}	1.023^{***}	0.785^{*}
	(0.431)	(0.378)	(0.441)
Coefficients on aggregate control variables			
CDW_t	0.107^{***}	0.093^{***}	0.107^{***}
	(0.017)	(0.015)	(0.017)
$\tilde{\pi}_{i,s,t}^{p-1}$	0.074^{**}	0.076^{**}	0.074^{**}
	(0.030)	(0.030)	(0.030)
$\tilde{\pi}_{i,s,t}^{p-1}\mathbb{1}_{t-2\geq 2021Q3}$	0.005	0.003	0.004
	(0.032)	(0.032)	(0.032)
$\pi_{t-2}^{(1)}$	0.352^{***}	0.320^{***}	0.358^{***}
	(0.106)	(0.081)	(0.107)
$\pi_{t-2}^{(1)}\mathbb{1}_{t-2\geq 2021Q3}$	-0.822^{***}	-0.761^{***}	-0.835^{***}
	(0.123)	(0.105)	(0.125)
$\pi_{t-2}^{(1)}\mathbb{1}_{t-2\geq 2022Q3}$	0.238^{***}	0.232^{***}	0.243^{***}
	(0.049)	(0.049)	(0.049)
$\pi_{t-2}^{(2)}$	0.058	0.067	0.053
	(0.102)	(0.090)	(0.103)
$\pi_{t-2}^{(2)}\mathbb{1}_{t-2\geq 2021Q3}$	0.638^{***}	0.621^{***}	0.615^{***}
	(0.129)	(0.114)	(0.131)
Δs_{t-2}	-0.072	-0.074^{**}	$-0.075*$
	(0.044)	(0.037)	(0.045)
Δz_{t-2}	-0.028	-0.006	-0.029
	(0.039)	(0.033)	(0.041)
$\Delta u_{t-2}^{(1)}$	0.042	-0.016	0.043
	(0.054)	(0.041)	(0.054)
$\Delta u_{t-2}^{(2)}$	-0.020	-0.017	-0.020
	(0.027)	(0.025)	(0.027)
Coefficients on sectoral control variables			
$\pi_{s,t-4}^{e}$			0.007
			(0.015)
$\pi_{s,t-4}^{e}\mathbb{1}_{t-2\geq 2021Q3}$			0.030
			(0.022)
$\Delta\hat{z}_{s,t-2}$			-0.007
			(0.004)
Log-likelihood	-9684.626	-9686.855	-9692.101
AIC	0.249	0.249	0.250
$\sqrt{\text{Var}(\tau_t)}$	0.115	0.000	0.106
$\sqrt{\text{Var}(\tau_t)}$ (rel.)	0.123	0.000	0.113

Note: Standard errors in parentheses. ***, **, and * denote significance at 1%, 5%, and 10%. column 1 reproduces the full model from Table 17. column 2 restricts the time-varying intercept τ_t to a constant. column 3 adds sectoral control variables.

Figure 32: Kalman smoothed estimates of $\beta_{2,t}$ in full model (6) with τ_t restricted to a constant

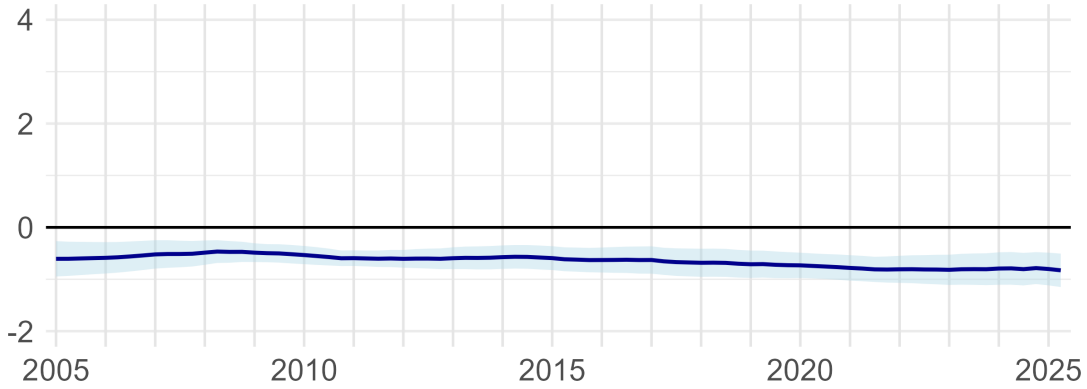


Note: The shaded area indicates the 95%-confidence interval.

E.2.2 Estimates of full model (6) with sectoral control variables

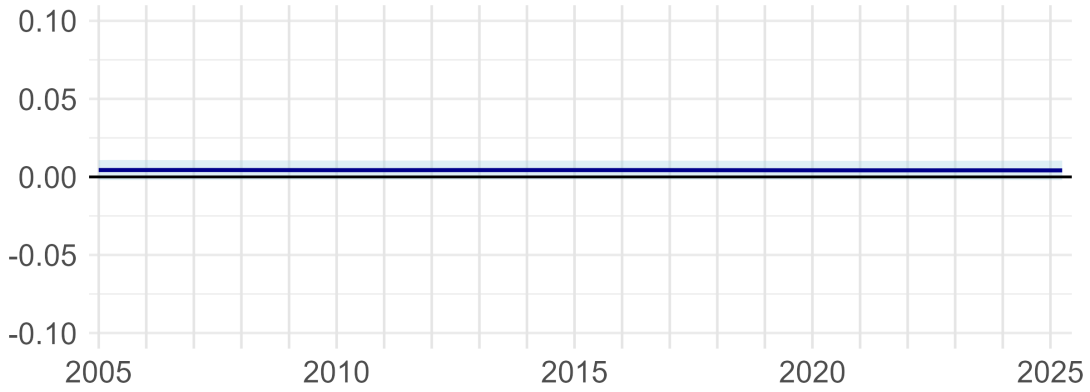
Column 3 in Table 18 and Figures 33-35 show that adding $\pi^e_{s,t-4}$ and $\Delta\hat{z}_{s,t-2}$ has no substantial impact on either the model fit or the Kalman smoothed estimates of the time-varying parameters in the full model. The coefficients on the sectoral control variables are not significantly different from 0.

Figure 33: Kalman smoothed estimates of τ_t in full model (6) with sectoral control variables



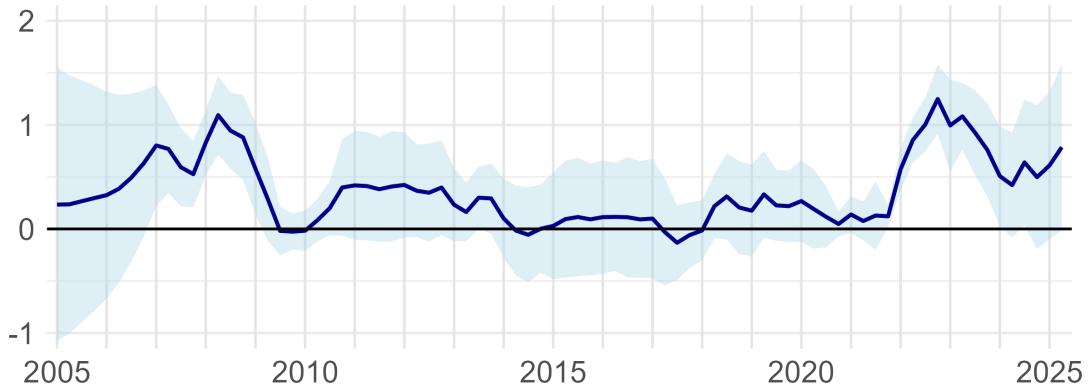
Note: The shaded area indicates the 95%-confidence interval.

Figure 34: Kalman smoothed estimates of $\beta_{1,t}$ in full model (6) with sectoral control variables



Note: The shaded area indicates the 95%-confidence interval.

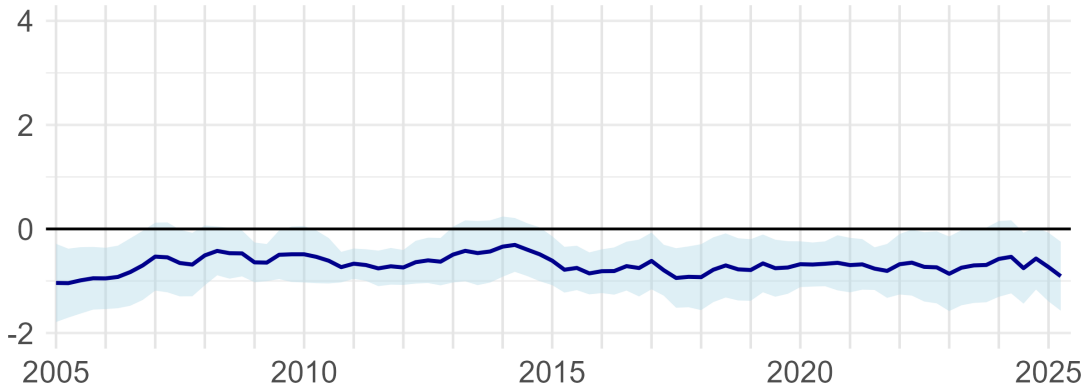
Figure 35: Kalman smoothed estimates of $\beta_{2,t}$ in full model (6) with sectoral control variables



Note: The shaded area indicates the 95%-confidence interval.

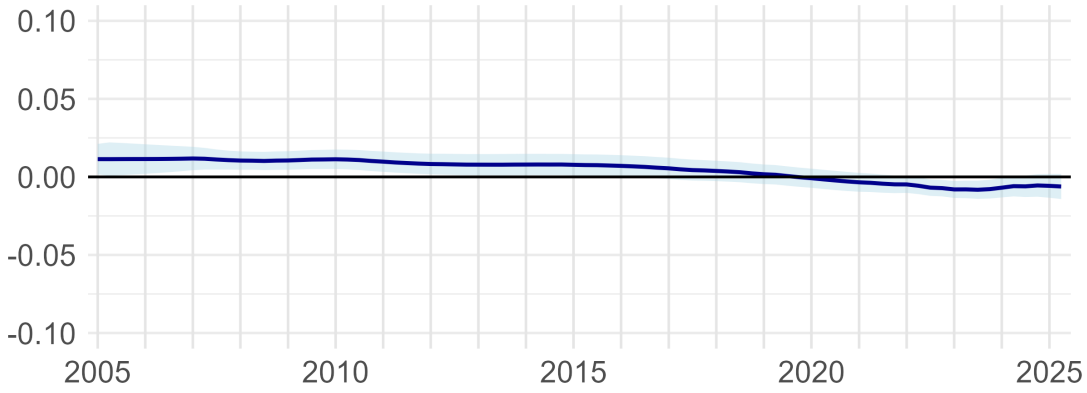
E.2.3 Estimates of full model (6) with sectoral labor market tightness measured in levels

Figure 36: Kalman smoothed estimates of τ_t in full model (6), with sectoral labor market tightness measured in levels



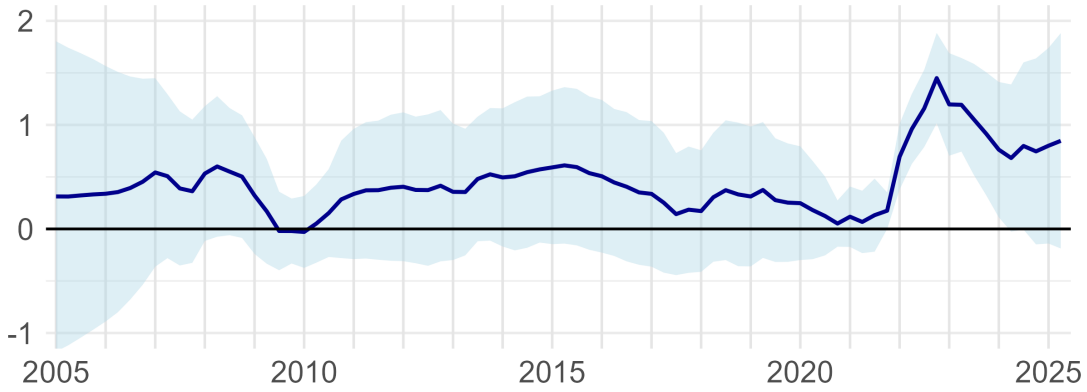
Note: The shaded area indicates the 95%-confidence interval.

Figure 37: Kalman smoothed estimates of $\beta_{1,t}$ in full model (6), with sectoral labor market tightness measured in levels



Note: The shaded area indicates the 95%-confidence interval.

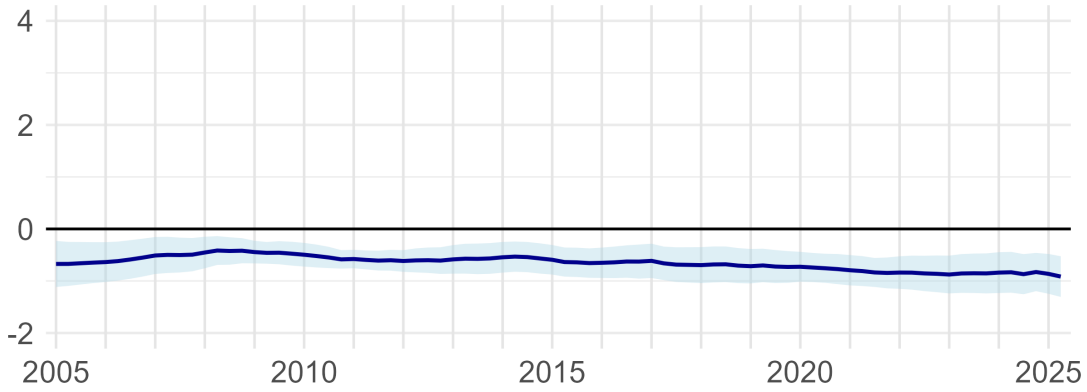
Figure 38: Kalman smoothed estimates of $\beta_{2,t}$ in full model (6), with sectoral labor market tightness measured in levels



Note: The shaded area indicates the 95%-confidence interval.

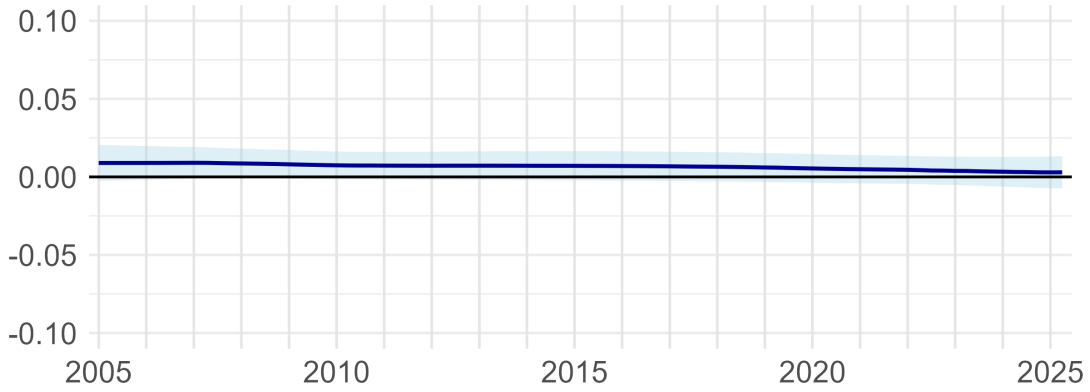
E.2.4 Estimates of full model (6), with sectoral labor market tightness measured as a four-quarter moving average of the vacancy rate growth rate

Figure 39: Kalman smoothed estimates of τ_t in full model (6), with sectoral labor market tightness measured as a four-quarter moving average of the vacancy rate growth rate



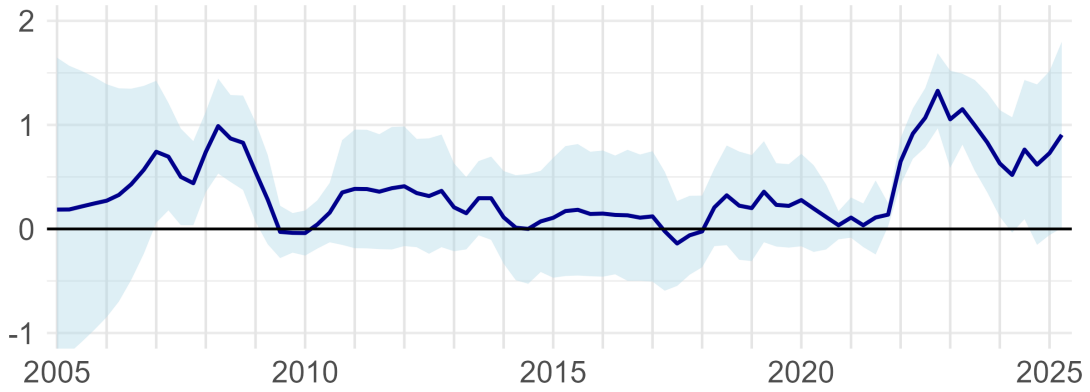
Note: The shaded area indicates the 95%-confidence interval.

Figure 40: Kalman smoothed estimates of $\beta_{1,t}$ in full model (6), with sectoral labor market tightness measured as a four-quarter moving average of the vacancy rate growth rate



Note: The shaded area indicates the 95%-confidence interval.

Figure 41: Kalman smoothed estimates of $\beta_{2,t}$ in full model (6), with sectoral labor market tightness measured as a four-quarter moving average of the vacancy rate growth rate



Note: The shaded area indicates the 95%-confidence interval.

E.2.5 Estimates of full model (6), with sectoral labor market tightness measured with a principal component at the sectoral level

All the variables related to labor market tightness that we could use for a sectoral principal component analysis are listed in Table 19 below. As the table shows, many of the sectoral equivalents of the macroeconomic data used to measure labor market tightness are either not available, or contain missing values for several sectors. For the variables measuring employed persons, hours worked and jobs growth sectoral data was only available at the annual level.

Table 19: Description of variables used in the sectoral principal component analysis

Variable	Unit	Freq.	SA
Employed persons	Y-o-Y %Δ	Y	No
GDP (production approach)	Y-o-Y %Δ	Q	No
Hours worked, total	Y-o-Y %Δ	Y	No
Jobs, total	Y-o-Y %Δ	Y	No
Labor share	Y-o-Y Δ	Y	No
Vacancies, open	Y-o-Y %Δ	Q	No
Vacancy rate	Y-o-Y Δ	Q	Yes

Notes: Y-o-Y %Δ denotes the year-on-year percentage change of the (average) index. SA indicates whether the series has been seasonally adjusted by the source.

From the sectoral PCA we obtained two sectoral labor market tightness principal components with an eigenvalue larger than 1, which we label $\Delta\theta^{(1)}_{s,t-2}$ and $\Delta\theta^{(2)}_{s,t-2}$, see table 20. We also ran regressions (7) for the sectoral principal components, the results of which are shown in table 21. The table shows that, similar to the results for the PCA on aggregate data, one principal component loads strongly on the vacancies data, while the other loads more on other labor market indicators. Analogously to the aggregate measure of tightness, we again took the average of these two components and standardized it to obtain our final measure of sectoral labor market tightness, shown in Figure 42. The resulting estimates of our full model with our sectoral principal component of labor market tightness are shown in the figures 43-45 below.

Table 20: Eigenvalues and explained variance of unrotated sectoral principal components of

Component	Eigenvalue	Proportion	Cumulative
$\Delta\theta^{(1)}_{s,t-2}$	3.076	0.439	0.439
$\Delta\theta^{(2)}_{s,t-2}$	2.143	0.306	0.745

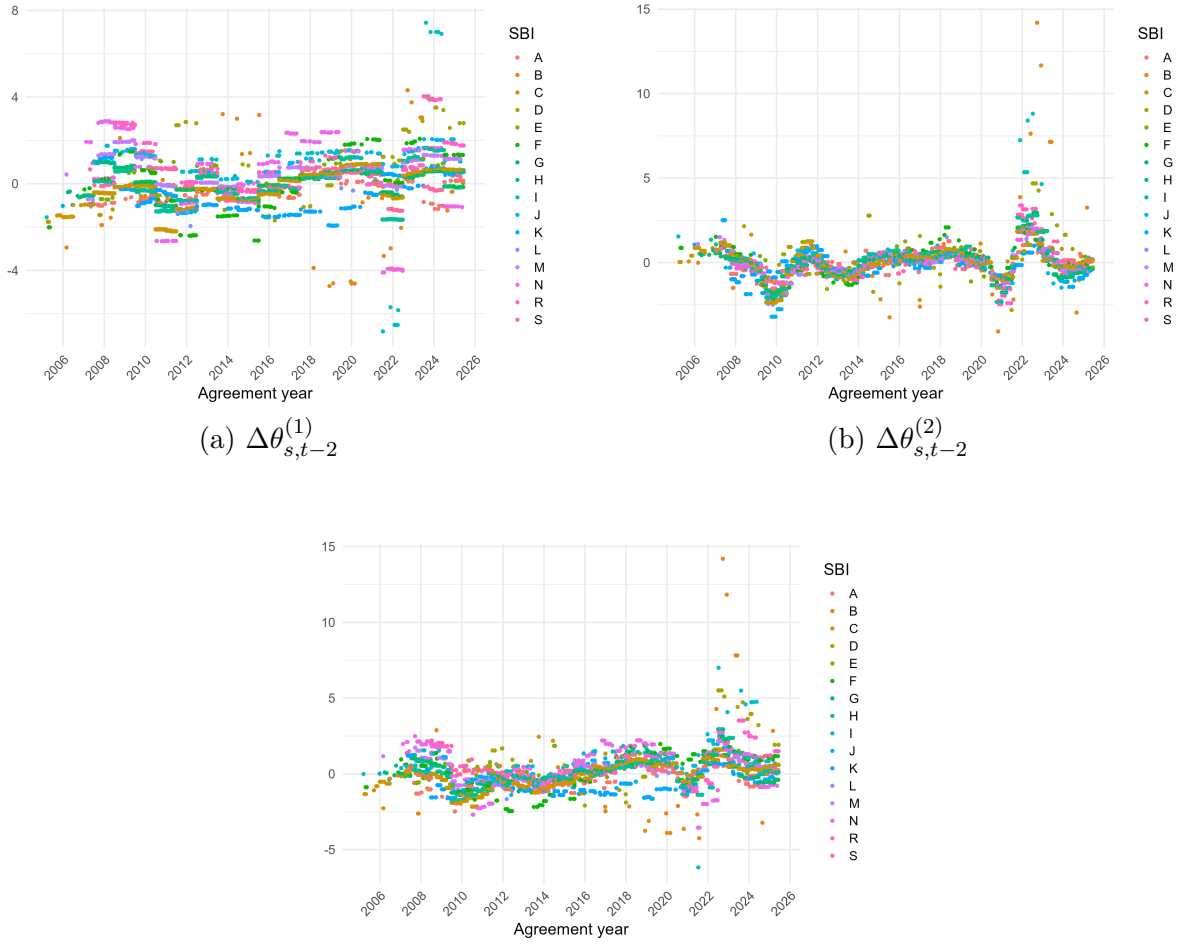
Notes: 'Proportion' is the share of total variance explained by each component individually, and *Cumulative* is the running sum of the proportion column.

Table 21: R^2 of equations (7) for sectoral data

Variable	$\Delta\theta^{(1)}_{s,t-2}$	$\Delta\theta^{(2)}_{s,t-2}$
Employed persons	0.888	0.023
GDP (production approach)	0.017	0.554
Hours worked, total	0.832	0.048
Jobs, total	0.909	0.026
Labor share	0.147	0.046
Vacancies, open	0.073	0.860
Vacancy rate	0.054	0.860

Notes: This table presents the R^2-values pertaining to bivariate regressions (7) for each extracted sectoral principal component and the macroeconomic variables.

Figure 42: Sectoral principal components and their standardized average

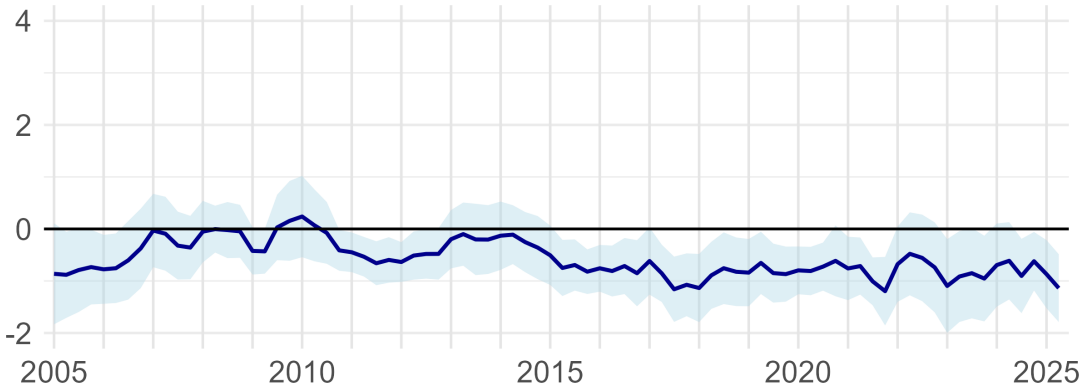


(a) $\Delta\theta^{(1)}_{s,t-2}$

(b) $\Delta\theta^{(2)}_{s,t-2}$

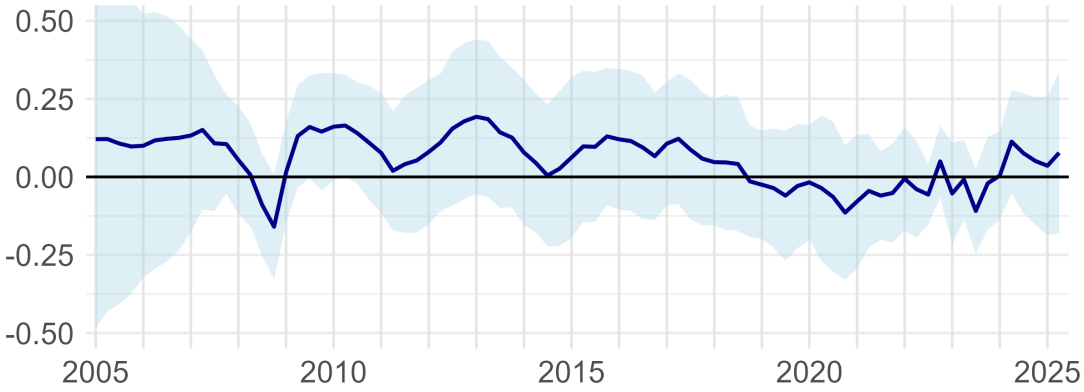
(c) Standardized average $\Delta\theta_{s,t-2}$

Figure 43: Kalman smoothed estimates of τ_t in full model (6), with sectoral labor market tightness measured with a principal component at the sectoral level



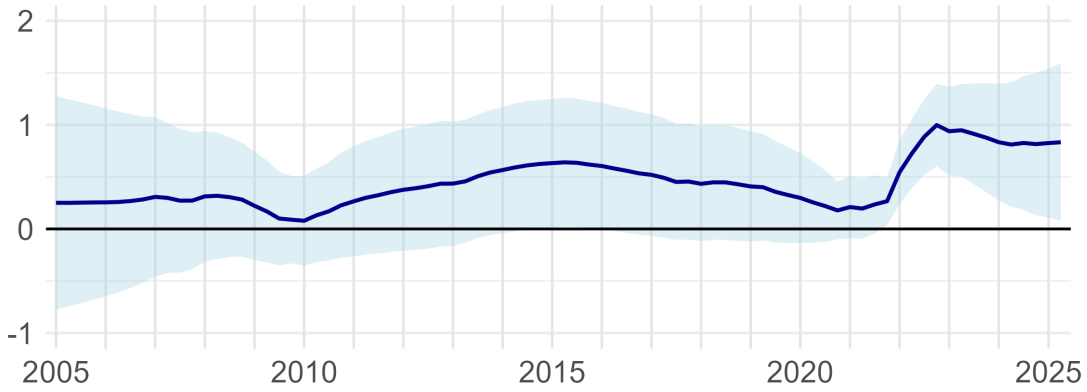
Note: The shaded area indicates the 95%-confidence interval.

Figure 44: Kalman smoothed estimates of $\beta_{1,t}$ in full model (6), with sectoral labor market tightness measured with a principal component at the sectoral level



Note: The shaded area indicates the 95%-confidence interval.

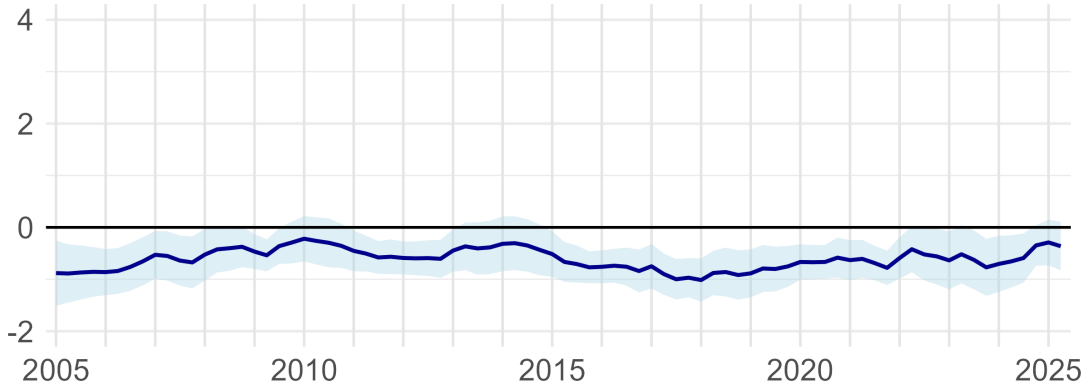
Figure 45: Kalman smoothed estimates of $\beta_{2,t}$ in full model (6), with sectoral labor market tightness measured with a principal component at the sectoral level



Note: The shaded area indicates the 95%-confidence interval.

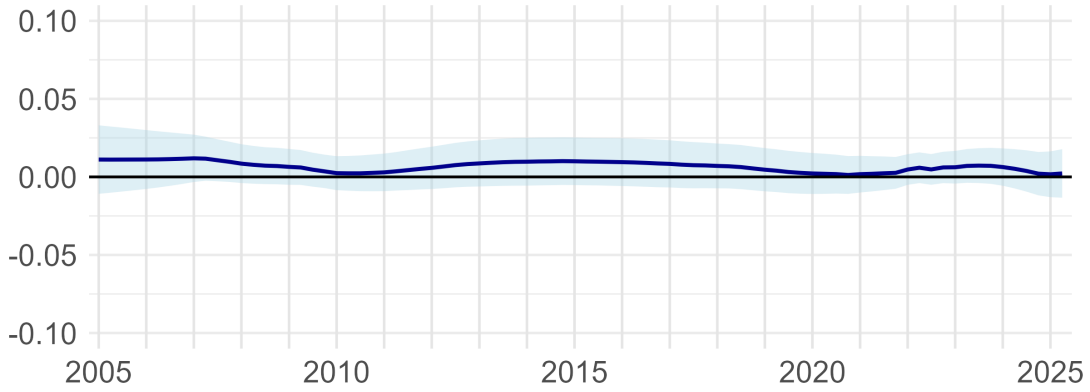
E.2.6 Estimates of full model (6), with aggregate tightness measured by principal component $\Delta\theta^{(1)}$

Figure 46: Kalman smoothed estimates of τ_t in full model (6), with aggregate tightness measured by principal component $\Delta\theta^{(1)}$



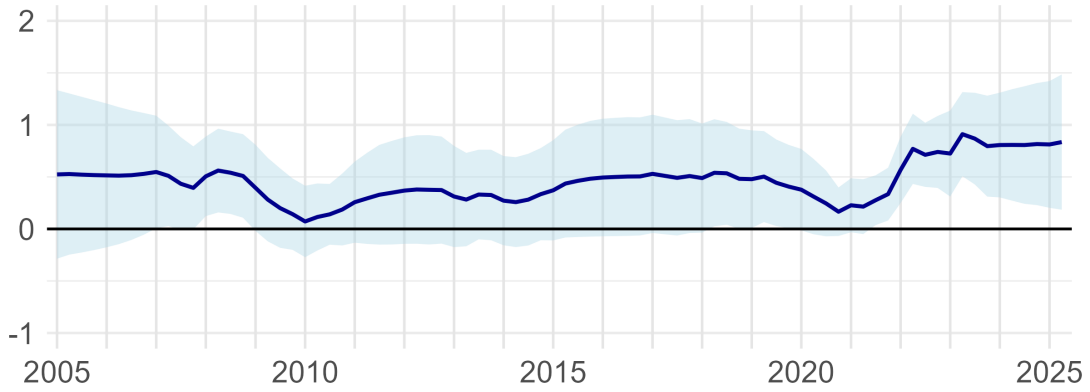
Note: The shaded area indicates the 95%-confidence interval.

Figure 47: Kalman smoothed estimates of $\beta_{1,t}$ in full model (6), with aggregate tightness measured by principal component $\Delta\theta^{(1)}$



Note: The shaded area indicates the 95%-confidence interval.

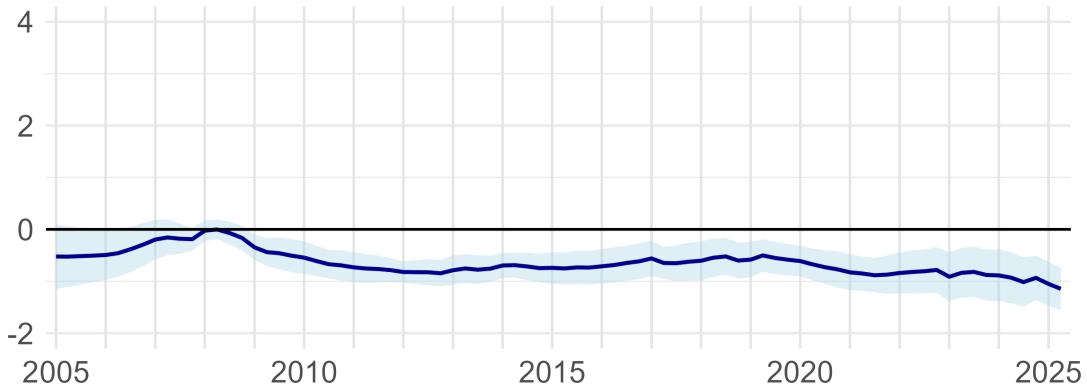
Figure 48: Kalman smoothed estimates of $\beta_{2,t}$ in full model (6), with aggregate tightness measured by principal component $\Delta\theta^{(1)}$



Note: The shaded area indicates the 95%-confidence interval.

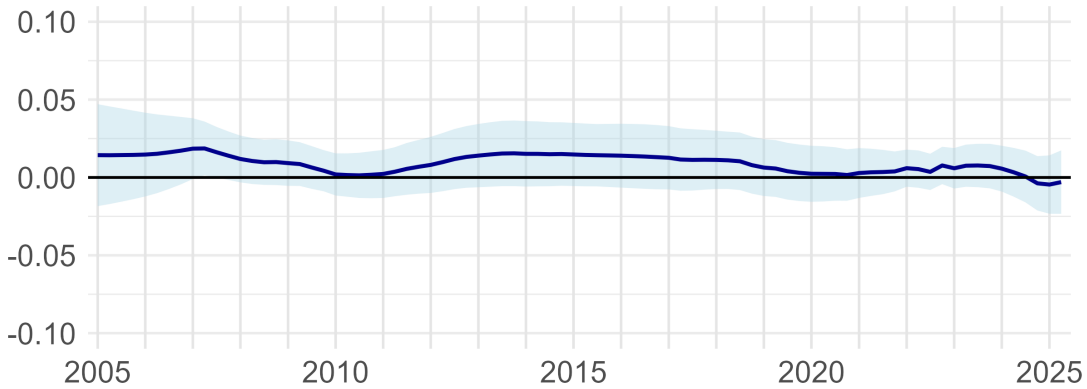
E.2.7 Estimates of full model (6), with aggregate tightness measured by principal component $\Delta\theta^{(2)}$

Figure 49: Kalman smoothed estimates of τ_t in full model (6), with aggregate tightness measured by principal component $\Delta\theta^{(2)}$



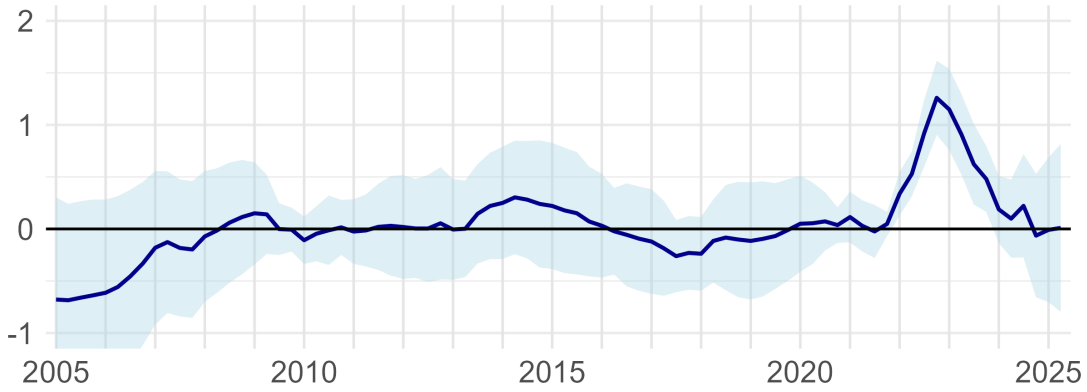
Note: The shaded area indicates the 95%-confidence interval.

Figure 50: Kalman smoothed estimates of $\beta_{1,t}$ in full model (6), with aggregate tightness measured by principal component $\Delta\theta^{(2)}$



Note: The shaded area indicates the 95%-confidence interval.

Figure 51: Kalman smoothed estimates of $\beta_{2,t}$ in full model (6), with aggregate tightness measured by principal component $\Delta\theta^{(2)}$



Note: The shaded area indicates the 95%-confidence interval.