Firm Sorting and Gender Pay Gaps in Local Labour Markets

Ruobing Li* Balázs Lengyel Virág Ilyés

Corvinus University of Budapest

August 23, 2026

Abstract

This study examines how persistent firm wage components relate to the cross-district dispersion in adjusted gender pay gaps and which local labour-market characteristics are associated with the remaining variation. Using Admin3 linked employer–employee records for 174 non-Budapest Hungarian residence districts during 2003–2017, it implements a three-step descriptive design. First, a pooled hourly-wage model recovers district-by-gender effects under common worker controls. Second, firm fixed effects are added and the dispersions of the recovered gaps are compared. Third, the baseline and firm-adjusted gaps are related to district characteristics. Firm fixed effects absorbed 68.8% of the baseline cross-district variance, although adjusted female penalties remained in every district. A larger public-sector employment share was associated with a smaller firm-adjusted female penalty, a relationship that persisted when gaps were re-estimated among private-sector workers. Greater industry concentration was associated with a larger penalty, although inference was weaker. The findings provide a fine-grained spatial accounting of firm-related compression and the correlates of the remaining district variation. They do not identify separate firm-side channels or causal effects.

JEL Classification: J16, J31, R23, C23

Keywords: Gender pay gap, firm fixed effects, local labour markets, linked employer–employee data, spatial wage dispersion, Hungary

National estimates of the gender pay gap (GPG) summarise average wage inequality but cannot reveal how adjusted wage differentials are distributed across local labour markets. The baseline estimates reported in this study show that the national mean conceals substantial cross-district variation even after observable worker characteristics have been taken into account. This variation matters because a common national estimate can coexist with markedly different local wage structures, making the spatial distribution of adjusted gaps an empirical object in its own right. Hungary offers a bounded setting for examining this distribution: public employment accounted for about one fifth of the study's eligible-wage sample, while foreign manufacturing

*Corresponding author: Corvinus University of Budapest. Email: `ruobing.li@stud.uni-corvinus.hu`. We thank Álmos Telegdy, Balázs Reizer, Rita Pető, István Boza, and participants in the doctoral seminar at Corvinus University of Budapest for helpful comments. We gratefully acknowledge the HUN-REN KRTK Databank for providing access to the administrative data used in this study. Any remaining errors are our own.

investment was geographically uneven across locations differing in market access, existing industrial activity, and inter-firm linkages (Békés, 2004; Kiss, 2007). The analysis therefore focuses on adjusted gender pay gaps across 174 non-Budapest residence districts during 2003–2017.

Observed district gender pay gaps need not represent pure place effects because they can also reflect the employers at which local women and men work. Evidence from linked employer–employee data shows that geographic moves traverse local employer pay hierarchies, so employer allocation can be reflected in measured geographic wage differences (Card et al., 2025). The same accounting issue arises here. Without firm identity, a district-by-gender effect combines a district wage component with the average premium paid by firms employing workers of that gender in the district. Cross-district differences may therefore combine location-related wage variation with differences in the firms employing women and men. Adding firm fixed effects provides an accounting comparison in terms of dispersion rather than levels. The comparison remains descriptive: it measures how much cross-district variation is absorbed by firm fixed effects, without attributing it to a specific firm-side mechanism.

A single district-level gender pay gap estimate cannot answer two distinct diagnostic questions: how much of its cross-district dispersion is absorbed by firm fixed effects, and which local characteristics are associated with the variation that remains. The first requires comparing adjusted district gaps estimated with and without firm fixed effects. This shows how much dispersion is absorbed by firm fixed effects, without identifying the underlying firm-related channels they reflect. The second requires treating the recovered gaps as district-level outcomes and examining their associations with local labour-market characteristics. Considering only the first would leave the remaining pattern undescribed; considering only the second would obscure whether those associations change once firm fixed effects are included. Both questions must therefore be addressed to distinguish firm-related compression of spatial dispersion from the correlates of the remaining district variation.

Existing studies have examined regional variation in gender pay gaps and the contribution of persistent firm wage components to aggregate gender wage inequality. Fuchs et al. (2021) decompose gaps across 401 German regions, while Nisic (2017) estimates conditional urban wage differentials from longitudinal data. Firm-decomposition studies use linked employer–employee records to estimate persistent employer components of aggregate gender wage gaps (Bruns, 2019; Masso et al., 2022). However, existing studies rarely connect a comparison of the cross-district dispersion in adjusted gender pay gaps before and after firm fixed effects with an analysis of the local characteristics associated with the remaining variation.

To address this gap, the study uses Admin3 linked employer–employee records covering 174 non-Budapest residence districts during 2003–2017 to quantify the dispersion absorbed by firm fixed effects and examine the local correlates of the remaining variation. First, it recovers adjusted district gender pay gaps from a pooled hourly-wage model containing district-by-gender effects and common worker controls. Next, it adds firm fixed effects and compares the dispersions of the two sets of recovered gaps. Finally, it treats the baseline and firm-adjusted gaps as district-level outcomes and relates them to local labour-market characteristics. Changes in dispersion are interpreted as variation absorbed by firm fixed effects, while the second-stage estimates are interpreted as associations rather than causal effects.

Firm fixed effects absorbed 68.8% of the baseline cross-district variance in adjusted gender pay gaps, although adjusted female penalties remained in every district. In the second stage, a larger public-sector employment share was associated with a smaller firm-adjusted female penalty. This relationship persisted when the gaps were re-estimated among private-sector workers. Greater industry concentration was associated with a larger penalty, although the inference was weaker. These descriptive associations do not identify their separate underlying mechanisms.

This study provides a fine-grained spatial accounting that connects persistent employer components to the cross-district dispersion in adjusted gender pay gaps and identifies which local structures are associated with the remaining variation. Because the outcome here is cross-district dispersion, that variance compression is not directly comparable with decomposition shares for average national gaps. For researchers, this distinction clarifies how firm-related evidence should be interpreted across spatial and aggregate designs. For regional-development and gender-equality analysts, the results may provide diagnostic evidence on employer and institutional composition for future evaluation. The analysis does not identify separate firm-side channels or the causal effects of potential interventions.

The remainder of the paper is organised as follows. Section I reviews regional gender pay-gap research and firm-decomposition studies, establishing the sequential diagnostic gap. Section II describes the Admin3 data, sample construction, wage measure, residence-based geography, and worker-level descriptive statistics. Section III explains the three-step empirical strategy. Section IV reports the dispersion comparison and district-level associations, while Section V assesses their robustness. Section VI interprets the findings, contributions, and limitations. Section VII summarises the main results and concludes.

I. Literature Review

A. From human capital to structural explanations

Human capital remains a necessary baseline for gender wage analysis, although it no longer provides a sufficient account of persistent gaps. While Mincer and Polachek (1974) link earnings to experience, interruptions, and family-related investment, Blau and Kahn (2017) show that educational convergence has sharply reduced the explanatory power of conventional endowments. Olivetti and Petrongolo (2016) show that employment selection and structural transformation condition observed gaps across industrialised countries; employee-only estimates therefore describe inequality among workers in employment rather than the full working-age population. In Central Europe, Mysíková (2012) finds that employed Hungarian women had stronger observed individual characteristics than employed men, yet remuneration differences and employment selection left a wage gap despite those advantages. Skills and experience therefore remain essential controls, while the residual Hungarian gap requires analysis of the jobs, firms, and wage structures through which those characteristics are rewarded.

Persistent gender gaps can arise when jobs attach different prices to availability and continuity, converting family-related constraints into unequal returns and job allocation. Goldin (2014) locates this mechanism in the price of temporal flexibility: where handovers are costly and workers are imperfect substitutes, pay rises nonlinearly with long or continuous availability. Evidence from elite professionals and changing overwork premiums shows how modest differences in hours or interruptions can cumulate into divergent careers

(Bertrand et al., 2010; Cha and Weeden, 2014). Kleven et al. (2019) provide a broader event-study anchor in Denmark, following labour-market outcomes around first childbirth. Their results show that parenthood persistently changes women's participation, hours, wages, occupational rank, sector, and firm allocation relative to men's, linking family constraints directly to sorting. Although these studies differ in population and design, they show that earlier constraints can generate observed occupation or employer placement, which should not be interpreted as a fixed preference. Temporal demands and parenthood therefore matter here because they shape the firms, sectors, and local jobs into which otherwise similarly qualified workers sort.

B. Firms, sorting, and within-firm wage setting

Persistent worker and employer wage components make employer allocation a distinct source of observed wage variation, as Abowd et al. (1999) demonstrate using French linked data. In West German workplace data, Card et al. (2013) further show that inequality can reflect both dispersion in persistent establishment premia and their covariance with worker components, although both remain proximate accounting relationships rather than causal mechanisms. Here, spatial sorting denotes the allocation of workers across places; gender sorting across firms denotes unequal allocation of women and men across high- and low-premium employers; occupation–firm sorting denotes gendered allocation across specific occupation–workplace cells; and within-firm wage setting denotes gender-differentiated returns inside the same employer. The first three concern allocation across units defined at different levels, whereas the fourth concerns pay determination within an employer environment. A common firm fixed effect may therefore absorb cross-district gender-gap variation produced by spatial sorting, gender sorting across firms, and occupation–firm sorting, since these operate through which firms employ women and men, but it cannot identify their separate contributions. It cannot, however, absorb within-firm wage setting: a firm fixed effect common to both genders applies equally to all workers at that employer, so any gender-differentiated pay within the same firm remains fully present in the recovered gender gap whether or not firm fixed effects are included.

Linked-data studies agree that firms matter for gender wage inequality, although reported sorting and within-firm shares differ with the estimand and counterfactual. Card et al. (2016) use Portuguese private-sector data to separate women's access to high-premium firms from gender differences in returns, finding that sorting dominates the measured firm contribution. Decomposing levels and changes separately, Bruns (2019) finds that sorting dominates West Germany's firm-related gap in levels, whereas its widening reflects changing gender-specific firm premia more than worker relocation. Masso et al. (2022) attribute 40.1% of Estonia's raw gap to firm premia, although their sorting–bargaining split changes markedly with the gender-specific firm-effect counterfactual. By contrast, Penner et al. (2023) use occupation–firm fixed effects rather than worker–firm mobility models and find that more than nine-tenths of Hungary's adjusted gap remains within occupation–firm cells. That residual is, however, neither an AKM bargaining component nor direct evidence of discrimination. Collectively, these studies show that measured firm contributions depend on sector coverage, sample, normalisation, and whether the object compared is the level of the gap or its dispersion across places.

Search and mobility frictions give employers wage-setting power, whose gender implications depend on how strongly women and men respond to firm-level wage offers. Modern monopsony does not require a dominant employer: imperfect information and mobility costs make firm-level labour supply upward-sloping, while concentration is only one observable source of employer power (Manning, 2003; Azar et al., 2022). Estimating firm-level labour-supply elasticities from separations and recruitment in German linked data, Hirsch et al. (2010) find women less responsive to wage offers than men, although the implied wage consequences remain model-dependent. In Hungary, Lovász (2008) relates firm-level adjusted gaps to three-digit industry competition measures, while Pető (2025) combines worker and firm fixed effects with acquisition and new-hire comparisons for 2003–2017. Lovász associates stronger product-market competition with smaller within-firm gaps; Pető finds larger gaps in foreign firms, with new-hire patterns consistent with greater discretion at hiring. Hungarian evidence thus identifies competition and ownership as dimensions of employer heterogeneity that need not align, so no single firm-level proxy summarises local employer structure, while leaving unresolved how persistent employer components map onto gender-gap dispersion across districts.

C. Local opportunity structures and the public sector as a contextual local institution

District gender pay gaps combine worker composition with the spatial distribution of employers, linking local opportunity structures to the firms into which workers sort. Dense local markets reorganise matching and learning (Glaeser et al., 2009; Wheeler, 2001; Glaeser and Maré, 2001; Andersson et al., 2007; de la Roca and Puga, 2017), yet French employment-area evidence shows that workforce sorting accounts for a substantial part of observed spatial wage disparities (Combes et al., 2008). Card et al. (2025) extend this logic by showing that geographic moves also traverse local employer hierarchies, making employer composition part of measured place differences. Petrongolo and Ronchi (2020) show that goods-producing industries are more clustered and female-intensive services more dispersed, implying that local sector composition generates unequal spatial variation in men's and women's job opportunities. Consistent with this synthesis, Fuchs et al. (2021) treat regional gender gaps as outcomes of both worker composition and establishment geography. A district is therefore a joint worker–firm opportunity structure, although its nominal boundaries need not define the same effective employer set for women and men.

Women and men in the same district may face different effective labour markets because commuting constraints and household location decisions shape reachable employers (Sorenson and Dahl, 2016; Madden, 1981; Rapino and Cooke, 2011; Phimister, 2005). Using French job-search records, Le Barbanchon et al. (2021) show that women accept shorter commutes and trade wages for proximity, consistent with a narrower feasible employer set, although commuting preferences account for only part of the residual gap. Hirsch et al. (2013) interpret metropolitan–rural gap differences through spatial monopsony, in which thicker markets constrain employer power over less mobile workers, a reading that Nisic (2017) extends to partnered women's mobility constraints. England (1993) finds that women with greater domestic responsibilities did not systematically work closer to home, so spatial entrapment cannot be assumed from household roles alone, although the evidence rests on a single metropolitan case study. Residence-based Hungarian LAU 1 districts are therefore imperfect proxies for gendered accessibility because commuting crosses boundaries and district employer counts need not represent either gender's feasible choices.

Urban scale does not predict a stable regional gender pay gap because raw, conditional, and unexplained gaps respond differently to employment structure and selection. Across 401 German NUTS 3 regions, Fuchs et al. (2021) find that high gaps often reflect well-paid male employment, whereas low gaps can reflect weak male opportunities rather than strong female wages. Across 16 Polish NUTS 2 regions, Majchrowska and Strawiński (2016) find larger raw gaps in urbanised regions with male-intensive high-wage industries, but larger unexplained shares in rural ones. For partnered German women, Nisic (2017) finds a narrower conditional gap in large cities, consistent with wider urban opportunities. Murillo Huertas et al. (2017) use matched worker–workplace decompositions and regional panel models in Spain, associating gaps with workplace allocation and wage institutions; density is negatively associated with raw gaps, while its unexplained association is specification-sensitive. NUTS 2 averages may mask finer heterogeneity, motivating Hungarian LAU 1 analysis without treating districts as closed functional labour markets. Hence, unstable scale associations redirect attention from market size to local institutional composition, especially public employment as a contextual local institution.

Public employment is treated here as a contextual local institution, on three grounds developed in the literature: the compression of the local wage distribution, the concentration of skilled female employment, and the alternatives available to women in the private sector. Evidence from Britain and Norway associates centralised public pay systems with narrower wage distributions, although estimated public–private differentials remain sensitive to establishment size, geography, and worker composition (Blackaby et al., 2018; Rattsø and Stokke, 2020). Public administration, education, and health concentrate both female and professional employment, and the German evidence shows that gender pay gaps in this sector are substantially smaller than in manufacturing (Fuchs et al., 2021); a district's public-sector share and its share of women in high-skilled occupations therefore partly index the same local employment structure. Using Hungary's 2002 public-pay reform, Telegdy (2018) finds that public exposure was followed by higher corporate wages, with responses extending beyond directly exposed workers and varying across worker groups, consistent with a gendered outside-option spillover. These channels imply the same association in a mixed-sector sample but diverge once public employees are excluded: a purely mechanical composition channel should then disappear, whereas cross-sector gender sorting and private-sector wage responses should not. As a result, a larger local public sector is expected to be associated with a smaller adjusted female penalty even among private-sector workers, which would be inconsistent with an explanation based solely on the direct inclusion of public-sector employees in the outcome sample without distinguishing sorting from spillovers.

A district's living and service environment forms a fourth dimension of the local opportunity structure. Childcare provision has the firmest empirical basis: early childhood generates persistent employment and earnings penalties for women that are substantially larger than those experienced by men (Kleven et al., 2019), while quasi-experimental evidence from Hungary shows that access to subsidised childcare significantly increases mothers' labour supply (Lovász and Szabó-Morvai, 2019). Local formal-care capacity may therefore condition the extent to which care responsibilities interrupt women's employment and accumulate into earnings losses. The remaining variables serve as broader indicators of local development and vitality and are not assigned separate causal mechanisms. Better-serviced districts may offer a wider range of employers and feasible job matches, expanding women's alternatives to local low-wage work. However, German evidence shows that stronger local economies need not narrow gender pay gaps when high-paying opportunities are concentrated in male-dominated firms and industries (Fuchs et al., 2021). Recorded crime and net out-migration enter as descriptive indicators of local social conditions and population loss through migration. Therefore, these features complete the set of local characteristics used to assess the remaining district-level variation.

D. *Empirical gap and contribution*

The two closest regional studies provide complementary design benchmarks, but their data structures identify different dimensions of spatial gender-pay inequality. Fuchs et al. (2021) use German social-security records for 2017 to estimate separate wage equations and Oaxaca–Blinder decompositions across 401 NUTS 3 regions, partitioning observed gaps into individual, establishment, and regional characteristics. This cross-sectional design provides geographically detailed composition accounting, although its firm information enters through observed establishment attributes and not persistent employer fixed effects. Nisic (2017) instead analyses SOEP panel data from 1992–2012 with hybrid within–between wage models and a dynamic employment-selection correction, using residential urban categories to compare partnered individuals. This longitudinal design addresses time-invariant individual heterogeneity and selection among movers; the estimated object is an urban wage differential, not the variance of district-specific gaps. The two designs answer different questions and cannot be treated as instances of the same cross-sectional decomposition: Fuchs identifies regional composition at one date, whereas Nisic estimates conditional within- and between-person urban differences over time. Neither study, however, measures how much cross-district GPG dispersion is absorbed by persistent firm fixed effects in linked employer–employee data.

Regional and firm studies leave an unresolved empirical link between persistent employer wage components and the dispersion of gender pay gaps across small local labour markets. Regional designs either decompose local gaps by worker and establishment characteristics or estimate conditional urban differentials, documenting spatial heterogeneity without persistent employer effects (Fuchs et al., 2021; Nisic, 2017). Firm-decomposition studies use worker mobility to estimate employer components of national or pooled gaps, yet do not ask how those components generate differences across small geographic units (Card et al., 2016; Bruns, 2019; Masso et al., 2022). Neither family preserves both fine geography and persistent employer components. The present study addresses this missing link using Hungarian linked employer–employee records for 2003–2017 covering 174 non-Budapest LAU 1 districts—a post-transition setting with a large public sector and spatially uneven foreign ownership (Békés, 2004; Kiss, 2007; Telegdy, 2018). It first recovers district-by-gender fixed effects from wage models with common worker controls, producing comparable adjusted wage structures across districts. It then compares the dispersion of these effects before and after firm fixed effects, treating the change as descriptive absorption without assigning it to a causal sorting mechanism. Finally, it relates the remaining district variation to workforce, employer, and institutional characteristics, including public-sector employment, while interpreting the second-stage estimates as associations. The study therefore provides a fine-grained spatial accounting that connects persistent employer components to local GPG dispersion and documents which measured district characteristics covary with the remaining variation.

II. Data

A. Institutional setting

Hungary entered the study period with a spatially uneven post-transition employer structure: foreign manufacturing investment clustered where market access, existing industry, and inter-firm linkages were strongest, while foreign-owned activity remained concentrated around Budapest and north-western Hungary (Békés, 2004; Kiss, 2007); public employment accounts for about one fifth of employees in the districts studied. That share is measured on the eligible-wage sample; the estimation panel, which additionally requires valid occupation and industry codes, is about 9% public (Appendix E). Public-sector wages were organised through a wage grid, and in September 2002 the government raised public employees' base wages by roughly 50% (Telegdy, 2018); that reform was followed by wage responses in exposed corporate labour-market segments. The statutory monthly minimum wage also rose from HUF 25,500 to HUF 40,000 in January 2001 and to HUF 50,000 in January 2002, nearly doubling within two years (Kertesi and Köllő, 2003). Taken together, these features characterised the employer structure and wage-setting environment prevailing during the study period, and they motivate treating firm ownership, public employment, and regulated wage-setting as characteristics of Hungary's local labour-market context rather than as causal factors in the district-level analysis.

B. Data source and sample construction

The empirical analysis uses the Admin3 linked employer–employee panel from the HUN-REN KRTK Databank. Admin3 follows a random half-sample of the Hungarian population at monthly frequency, integrating tax, pension and health-insurance registers into a single panel; this paper uses the 2003–2017 waves, and the unit of observation is the worker–firm–month. The linked structure supplies the identifiers and measures the design requires: a firm identifier for firm fixed effects, a residence-district identifier for the district-by-gender effects, and monthly earnings and working-time variables from which hourly wages are constructed. Because the geographic identifier records the worker's district of residence rather than the workplace, and because plant locations are not observed for multi-establishment firms, the design is residence-based. Section VI discusses what this implies for the interpretation of the district estimates.

The estimation sample is built through a sequence of restrictions. Workers are retained if they are aged 18–65, which focuses the analysis on the working-age population and limits the influence of schooling and retirement transitions. The analysis is restricted to workers in ordinary employment and statutory public-service relationships. Self-employed workers, cooperative members, business partners, and other non-employee categories are excluded because their earnings and firm attachment are not comparable with the wage-setting relationships the firm fixed-effects design requires. Participants in public work schemes (közfoglalkoztatás) are excluded on separate grounds: this is a fixed-term programme with separately regulated remuneration, and including it would blend programme participants with regular public-sector employees in the district public-sector share used in the second stage. Observations are dropped where monthly earnings are missing or non-positive, where days worked are missing, or where the firm or residence-district identifier is missing or invalid, since neither the hourly wage nor the fixed effects can otherwise be constructed. Invalid occupation codes are excluded, removing 99,462 worker–firm–month observations. Invalid industry codes are subsequently excluded for consistency, removing a further 2,265,850 observations. The estimates therefore describe wage differences among employees rather than earnings differences across all employed persons. Detailed cleaning steps are reported in Appendix A.

The main analysis excludes the 23 Budapest districts, whose scale and commuting structure differ markedly from those of the remaining districts and whose internal boundaries do not separate distinct local labour markets. The estimates reported below therefore describe employees residing in 174 non-Budapest districts. After these restrictions the cleaned analysis panel contains 221,278,959 worker–firm–month observations, covering 589,926 firms across 174 non-Budapest districts and fifteen calendar years. The regression estimation samples are smaller, since estimation excludes observations with missing values on any core control. The dominant restriction is the industry code, which is missing for 48,207,034 observations (21.8% of the cleaned panel); the baseline district-by-gender specification is therefore estimated on 172,645,146 observations. The baseline and firm-fixed-effects specifications are estimated from the same cleaned panel and cover the same 174 districts. Adding firm fixed effects leads to the additional exclusion of 11,869 singleton observations, approximately 0.007% of the baseline estimation sample, so the underlying worker–firm–month samples are nearly, although not exactly, identical. Descriptive statistics for the cleaned analysis panel are reported in Table 1.

C. Variables

The outcome in all specifications is the natural logarithm of the gross hourly wage. Because the analysis concerns the price of labour rather than total monthly pay, the hourly wage is constructed from monthly gross earnings, insurance-covered employment days, and reported weekly hours on the main job:

$$hourly\_wage_{ijmt} = \frac{monthly\_earnings_{ijmt}/days_{ijmt}}{weekly\_hours_{ijmt}/5}. \tag{II.1}$$

Throughout, i indexes workers, j firms, m calendar months, and t years, so that $ijmt$ identifies a single worker–firm–month observation. The numerator gives average earnings per insurance-covered employment day, and the denominator converts reported weekly hours on the main job into average daily hours under a five-day convention. The resulting measure should be read as an administrative average hourly wage rather than a contractual wage rate, since overtime, irregular schedules, and informal working-time arrangements are not fully captured by the administrative hours variable. Relative to monthly earnings, however, it removes much of the gender difference in days worked and in part-time arrangements from the wage comparison, which is the appropriate basis for a wage-rate analysis. This log hourly wage, $\log(hourly\_wage_{ijmt})$, is the outcome throughout.

The individual controls are age, age squared, education, occupation, industry, and year fixed effects. Education is a quasi-education proxy rather than a direct qualification record: it is constructed from occupation-based information and cumulative worker histories, and therefore measures attainment with error, particularly for workers whose occupational trajectory is atypical. Occupation is measured at the two-digit level, and industry is the harmonised one-digit TEÁOR section.

Public and private employment are distinguished by the employment legal-relationship code rather than by industry or ownership. The public category comprises civil servants, public employees, and the armed services; ordinary employment contracts are classified as private, so employees of state-owned enterprises are treated as private-sector workers throughout. This classification basis differs from those used in comparable regional studies. Fuchs et al. (2021), for example, use the German Employee History, which covers employees subject to social-security contributions but excludes civil servants, military personnel, and the self-employed; their public-sector analysis therefore omits groups that the Hungarian legal-status codes include. District public-sector shares should accordingly not be read as directly comparable in level to public-sector measures derived from other national datasets. Armed-services records account for 1.7% of the eligible-wage sample. Their recorded share averaged 0.86% in 2003–2009 and 2011, stood at 2.45% in 2010, and shifted to about 2.66% from 2012 onwards. Appendix F reports the second stage with the public-sector share defined over civil servants and public employees alone.

The geographic unit is the Hungarian district (járás, LAU 1), assigned by the worker's permanent home address, giving 174 non-Budapest districts. District characteristics come from two sources. Employment-based measures are constructed from Admin3 itself, counting unique workers rather than worker–month observations so that workers observed for more months in a year are not overweighted. Institutional and socio-economic indicators come from the T-STAR settlement-level database, aggregated to districts before ratios are formed and then averaged over 2003–2017. The T-STAR employment measure is unavailable for 2004–2006, so its district average uses the twelve available years; in logs it correlates at 0.994 with the Admin3 residence-based employment count. The construction and interpretation of the individual district covariates are set out in Section III.D.

D. Descriptive Statistics

Table 1 presents a descriptive puzzle. Women are more concentrated in higher-skill occupations than men (55.7 against 30.7%) and better represented in the upper education categories: 31.6% of female observations are classified as secondary and 24.2% as tertiary, against 13.4 and 22.3% of male observations. They are also slightly older on average, at 41.1 against 39.9 years. Their mean log hourly wage is nevertheless 6.356, against 6.448 for men, which is a raw female–male difference of -0.092 log points. These composition measures are not independent, since the quasi-education proxy is constructed partly from occupational information, but they point in the same direction. Despite this favourable observed composition, women have a lower mean log hourly wage; the conditional regressions examine how much of this contrast remains under common controls. Section III therefore recovers adjusted district-by-gender wage structures before examining how they change once firm fixed effects are added.

Table 1. Sample Composition by Gender

	Men	Women	Total
Worker–firm–month observations	114,512,444	106,766,515	221,278,959
Mean log hourly wage	6.448	6.356	6.404
SD log hourly wage	0.652	0.585	0.622
Mean age	39.9	41.1	40.5
Education (%)			
Elementary / vocational	64.29	44.25	54.62
Secondary	13.40	31.56	22.17
Tertiary	22.31	24.18	23.21
Occupation (%)			
Higher-skill occupations	30.74	55.73	42.80

Note: Cleaned Admin3 analysis panel, 2003–2017, Budapest districts excluded: 589,926 firms across 174 districts and fifteen calendar years, giving 348 district-by-gender cells. The raw female–male difference in mean log hourly wage is -0.092 log points. Higher-skill occupations are harmonised FEOR codes below 34, the same threshold used for the female higher-skill share in the second stage; lower-skill is the complement. Education is a quasi-education proxy constructed from occupation-based information and cumulative worker histories, not directly observed attainment. Percentages exclude observations with missing codes: 525 for occupation (217 male, 308 female) and about 0.25% for education. Full descriptive statistics, including wage percentiles, the age distribution, cell sizes, and industry composition, are reported in Appendix A.

III. Empirical Strategy

A. *Step 1: Measuring the Gender Pay Gap District by District*

The first step recovers, for each district, a gender pay differential that is adjusted for observed worker and job characteristics but not yet for the identity of the employing firm. Let $g(i) \in \{F, M\}$ denote the gender of worker i, and let d index the 174 non-Budapest districts, assigned by the worker's permanent home address rather than by workplace location. Month subscripts are suppressed throughout this section, since no parameter is indexed by month. The baseline specification is

$$y_{ijdt} = \mathbf{x}'_{ijdt}\beta + \mu_{d,g(i)} + \tau_t + \varepsilon_{ijdt}, \tag{III.1}$$

where y_{ijdt} is the log hourly wage constructed in equation (II.1), $\mu_{d,g(i)}$ denotes the 348 district-by-gender fixed effects, and τ_t denotes year fixed effects. The control vector contains a quadratic in age, categories of the quasi-education proxy, two-digit harmonised FEOR occupation codes, and one-digit TEÁOR industry sections. Education enters as a proxy rather than as a qualification record, since it is constructed from occupation-based information and cumulative worker histories (Section II.C), while the occupation and industry controls fix the broad job cells within which women and men are compared.

The adjusted gap in district d is the female-minus-male difference in these effects,

$$\widehat{GPG}_d^{\text{Base}} = \hat{\mu}_{d,F} - \hat{\mu}_{d,M}, \tag{III.2}$$

so that negative values indicate that women earn less than observationally comparable men residing in the same district. The levels of the fixed effects are identified only up to a normalisation, whereas the within-district difference in equation (III.2) is invariant to it, and that invariance is the property on which the following two steps rest. Because all effects are recovered from a single pooled regression, equation (III.2) should be read as a district-specific conditional differential estimated under common returns to observed characteristics, rather than as the gap a separately estimated district wage equation would produce; cross-district heterogeneity in those returns remains inside $\mu_{d,g(i)}$.

After observations with missing values on any core control are excluded, the Step 1 estimation sample contains 172,645,146 worker–firm–month observations, covering all 174 districts and fifteen calendar years. The dominant restriction is the industry code, which is missing for 21.8% of the cleaned analysis panel and whose incidence varies markedly across districts; the recovered gaps therefore describe workers with a valid industry code rather than the full district workforce, and Section V reports the design re-estimated with the industry control removed, so that observations with missing codes are retained.[1] Section III.B re-estimates equation (III.1) with firm fixed effects added and compares the dispersion of the two sets of recovered gaps.

B. Step 2: How Much of the Spatial Pattern Is Firm-Related

The district-by-gender effects in equation (III.1) are not pure place effects. Because firm identity is left out of that specification, $\mu_{d,g(i)}$ combines a district wage component with the average premium paid by the firms that happen to employ workers of gender g in district d; two districts whose employers set wages identically can therefore display different adjusted gaps simply because women and men there are attached to different firms. The second step provides an accounting adjustment for this by re-estimating equation (III.1) with firm fixed effects added, examining how the dispersion of the recovered district gaps changes:

$$y_{ijdt} = \mathbf{x}'_{ijdt}\beta + \tilde{\mu}_{d,g(i)} + \psi_j + \tau_t + \varepsilon_{ijdt}, \tag{III.3}$$

where ψ_j is a persistent wage component shared by all observations attached to firm j. The firm-adjusted gap is recovered exactly as before,

$$\widehat{GPG}_d^{\text{FirmFE}} = \hat{\tilde{\mu}}_{d,F} - \hat{\tilde{\mu}}_{d,M}. \tag{III.4}$$

The comparison of interest is between the dispersions of the two sets of recovered gaps, not between their levels. Section IV reports two objects: the cross-district standard deviation under each specification, and the joint distribution of the two, summarised by the Pearson correlation and the OLS slope. Firm fixed effects serve here as an accounting device rather than as a mechanism. As set out in Section I.B, ψ_j can absorb spatial sorting, gender sorting across firms, and occupation–firm sorting, but not within-firm wage setting, since ψ_j is common to both genders within a firm; the specification does not identify the separate

[1] Across the 174 districts the pooled missingness rate averages 23.2% with a standard deviation of 4.7 percentage points, ranging from 14.8 to 38.6%, and it is highest in the earliest and the final years of the panel. It is positively correlated with both recovered gaps—districts with less complete coverage display smaller adjusted female penalties—at $r = 0.37$ for the baseline gap and $r = 0.56$ after firm fixed effects. Its correlation with the firm-fixed-effects adjustment itself is not distinguishable from zero ($r = -0.08$, $p = 0.33$), so differential coverage across districts is not aligned with the compression documented in Section IV.

contributions of the three sorting-type mechanisms it does absorb. ψ_j is an average firm wage component, not a bargaining or rent-sharing parameter. Any change in dispersion is accordingly described throughout as variation absorbed by firm fixed effects rather than as variation explained by firms.

Firm fixed effects cannot be estimated for firms observed once, so Step 2 additionally drops 11,869 singleton observations and is estimated on 172,633,277 of the 172,645,146 observations used in Step 1, approximately 0.007% fewer. The two steps are therefore nearly, although not exactly, sample-matched, and the recovered effects are comparable across districts only within a connected set of firms and district-by-gender cells. Section V reports the dispersion comparison re-estimated with Step 1 restricted to the Step 2 sample.

C. Step 3: What the Remaining Spatial Pattern Is Associated With

The third step asks what the recovered gaps are associated with once they are treated as district-level outcomes in their own right. Each gap is regressed on a vector of district characteristics,

$$\widehat{GPG}_d = \alpha + \mathbf{z}_d'\gamma + u_d, \tag{III.5}$$

estimated across the 174 districts by weighted least squares. The vector $\mathbf{z}_d$ contains ten characteristics grouped into local workforce composition, job-opportunity and employer structure, and the institutional and socio-economic environment; their construction and sources are set out in Section III.D and Table 2. Equation (III.5) is estimated twice, once with $\widehat{GPG}_d^{\text{Base}}$ and once with $\widehat{GPG}_d^{\text{FirmFE}}$ as the dependent variable. Running both is the point of the exercise: attenuation after firm adjustment is consistent with overlap between a characteristic and persistent firm composition, without establishing that one is upstream of the other, whereas a characteristic that remains associated with the recovered gap after firm fixed effects are added is associated with local gender pay differences net of persistent firm wage components.

Two features of equation (III.5) require comment. First, the dependent variable is itself an estimate, which gives rise to a generated-regressor problem. Its sampling error enlarges u_d and therefore the residual variance, although it leaves γ unbiased as long as that error is unrelated to $\mathbf{z}_d$, and it is larger for small districts, whose gaps rest on fewer observations. Equation (III.5) is therefore weighted by district employment, which gives greater weight to the districts whose gaps are estimated on more observations. This is a partial rather than a complete remedy, since employment size proxies the precision of a recovered gap without being that precision.

Second, standard errors are heteroskedasticity-robust. Neighbouring districts share labour-market conditions, robust inference alone may be insufficient, and Section V re-estimates the second stage with standard errors clustered on the nineteen counties and with wild-cluster bootstrap p-values, which the small number of clusters makes necessary.

The sign convention inverts the intuitive reading of the second-stage coefficients, and is worth fixing before any results are reported. The dependent variable is the female-minus-male difference and is negative in every district, so it records a penalty by how far below zero it lies. A positive coefficient therefore means that districts scoring higher on a characteristic have gaps closer to zero, that is, a smaller adjusted female penalty; a negative coefficient means a larger penalty. Nor are the magnitudes elasticities: γ is the change in the log-point gap associated with a one-unit change in the characteristic, and because most characteristics are shares, a coefficient is read against a plausible movement in that share rather than against its full range. Coefficients are estimated in log points; for readability, magnitudes multiplied by 100 are described as approximate percentage-point changes in the log-wage gap. Throughout the paper, estimated wage differences for one gender are reported in per cent and differences between women and men in percentage points.

D. District Characteristics and Their Selection

The ten characteristics in $\mathbf{z}_d$ stand for the three ideas developed in Section I: who is locally available for well-paid work, what kind of employers are locally available to them, and what institutions and living conditions surround both. With only 174 districts the specification has to stay parsimonious, so candidate measures within each domain were screened rather than accumulated. A pairwise correlation above 0.7 flagged a pair for inspection but was not by itself disqualifying; the operative criterion was the variance inflation factor, with a value above five treated as a warning and resolved by retaining the measure with the stronger theoretical motivation or the more direct measurement. In the final specification the mean variance inflation factor is 2.74 and the maximum 4.94, on the female higher-skill share; Appendix B reports the full set alongside the pairwise correlations. Because the dependent variable is a single pooled estimate covering 2003 to 2017, each characteristic is averaged over the same window; Section V reports the second stage with the characteristics measured over 2003 to 2007 instead, which probes sensitivity to temporal ordering without resolving the joint determination of local characteristics and local pay gaps.

Four sets of candidate measures were tested and excluded under this rule. The total registered-jobseeker rate was dropped in favour of the female share of jobseekers, which bears more directly on women's outside options. Passenger cars per inhabitant, railway-station density and road-network density were each tested as measures of local mobility and each excluded. Primary-school teachers per 1,000 working-age inhabitants was excluded as overlapping the female higher-skill share. Firms per 1,000 working-age inhabitants was excluded as a composite of local scale and local development, correlating about 0.6 with each.

The retained measures also differ in the sample on which they are built. The public-sector share and employer concentration are constructed on the wider Admin3 eligible-wage sample, formed before occupation and industry codes restrict the analysis panel; the female higher-skill share uses the cleaned analysis panel. Industry concentration is computed only on observations carrying a valid industry code, rather than treating the missing code as an industry of its own: because the missing share ranges from 15 to 39% across districts, that alternative would make the index a mixture of industry structure and data coverage. Table 2 reports how the ten measures are distributed across the 174 districts together with the association each is expected to carry; Appendix B gives their construction, original variable codes and sample basis. The measures differ greatly in spread—the sewage-connected share runs from 0.07 to 0.94 across districts while the female share of jobseekers stays between 0.40 and 0.58—so a one-unit change means something very different from row to row.

Table 2. District-Level Covariates: Distribution and Expected Association

Variable	Mean	SD	Min	Max	Expected association
Female higher-skill occupation share	0.516	0.071	0.353	0.739	Smaller penalty
Female share of registered jobseekers	0.480	0.035	0.401	0.582	Larger penalty
Log district employment	9.60	0.69	8.14	11.51	Ambiguous
Employer concentration (HHI)	0.028	0.019	0.004	0.103	Larger penalty
Industry concentration (HHI)	0.178	0.042	0.113	0.320	Larger penalty
Public-sector employment share	0.201	0.041	0.126	0.348	Smaller penalty
Nursery places per 100 children aged 0–2	6.42	4.38	0.00	19.01	Smaller penalty
Sewage-connected dwelling share	0.589	0.197	0.068	0.942	Smaller penalty
Log crimes per 1,000 inhabitants	3.39	0.27	2.70	4.31	Larger penalty
Net migration per 1,000 inhabitants	−0.72	4.40	−7.71	16.65	Smaller penalty

Note: $N = 174$ non-Budapest districts; each measure is averaged over 2003–2017 before the district distribution is formed, and the statistics are unweighted across districts. "Expected association" states the direction anticipated in Section I, read against the sign convention of Section III.C: a smaller penalty corresponds to a positive second-stage coefficient and a larger penalty to a negative one. Nursery provision is recorded as zero throughout the period in ten districts, in which no settlement operated a crèche during the study period. Construction, original variable codes and sample basis are given in Appendix B.

IV. Results

A. Dispersion Before and After Firm Adjustment

Table 3 summarises the district-level gender pay gaps recovered from the two estimation steps. The baseline estimates measure the adjusted gap between observationally comparable women and men living in the same district. The firm-adjusted estimates additionally absorb persistent wage differences across firms. Their comparison therefore shows how much of the spatial dispersion in the baseline gaps is associated with the firms in which workers are employed.

Table 3. Distribution of Adjusted District Gender Pay Gaps, Main and Private-Sector Samples

Statistic	(1) Main sample Base	(2) Main sample Firm FE	(3) Private only Base	(4) Private only Firm FE
Mean	−0.1474	−0.1309	−0.1409	−0.1323
Standard deviation	0.0404	0.0226	0.0421	0.0242
Minimum	−0.3312	−0.2240	−0.3293	−0.2333
Maximum	−0.0637	−0.0631	−0.0445	−0.0565
Reduction in variance, Base → Firm FE	—	68.8%	—	67.0%
Correlation, Base with Firm FE	—	0.689	—	0.681
Correlation with main-sample estimate	—	—	0.994	0.991

Note. Adjusted district gaps in log points; $N = 174$ non-Budapest districts in every column, unweighted. Negative values indicate a female penalty. Columns (1) and (2) report the gaps from equation (III.2) and equation (III.4); columns (3) and (4) re-estimate both steps for ordinary employment relationships only (`fogvisz1 = 201`), as defined in Section II.C. The memo rows compare the two columns within each sample, except the last, which compares each private-sector estimate with its main-sample counterpart.

The baseline gap is negative in all 174 districts. Its mean is −0.147 log points, with a cross-district standard deviation of 0.040. The adjusted differential is therefore wider than the raw female–male difference of -0.092 log points recorded in Table 1: holding age, the quasi-education proxy, occupation and industry constant moves the gap away from zero rather than towards it, which is the pattern anticipated in Section II.D. The estimates range from -0.331 in the district with the largest adjusted penalty to -0.064 in the district with the smallest. The national mean therefore conceals substantial spatial variation even after differences in workers' observable characteristics have been taken into account.

Adding firm fixed effects makes the average gap less negative and sharply compresses its distribution. The mean rises from -0.147 to -0.131 log points, while the standard deviation falls from 0.040 to 0.023. This corresponds to a 44.1% reduction in the standard deviation and a 68.8% reduction in variance. Figure 1 shows this compression directly: the firm-adjusted distribution is narrower and more concentrated around its mean. The comparison concerns the dispersion of the district estimates rather than a decomposition of their average level. In particular, the reduction should be read as the share of spatial variation absorbed by persistent firm wage components, not as a causal effect of firms.

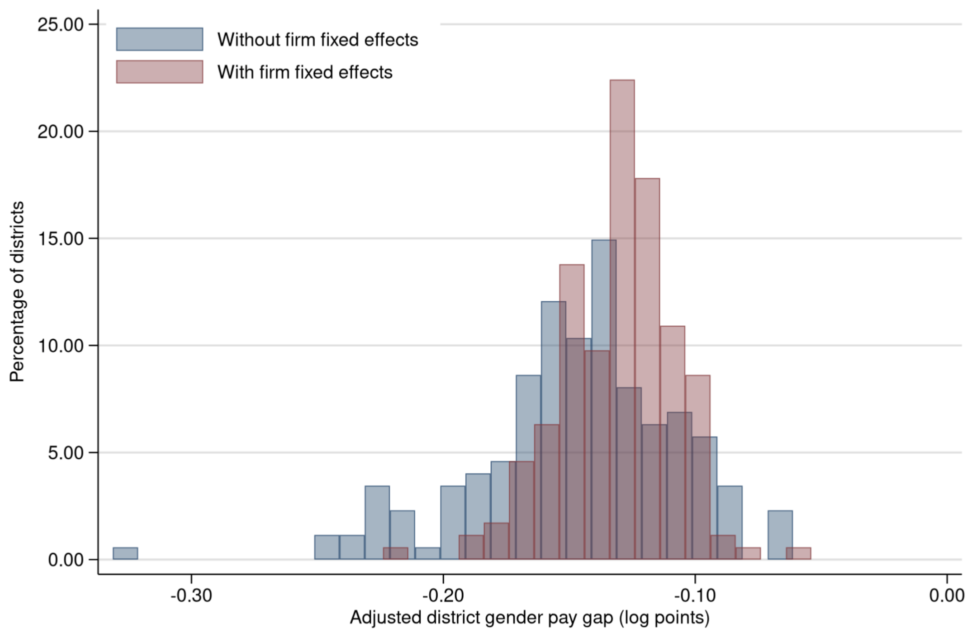

Figure 1. Distribution of Recovered District-Level Gaps, Before and After Firm Fixed Effects

Note. Unweighted distribution across 174 non-Budapest districts, in bins of 0.01 log points. Negative values indicate lower adjusted earnings for women.

The private-sector re-estimation produces almost the same result. After public-sector employment relationships are excluded, in columns (3) and (4) of Table 3, the standard deviation falls from 0.042 to 0.024, reducing the variance by 67.0%. The private-sector estimates are also closely aligned with their main-sample counterparts: the correlations are 0.994 for the baseline gap and 0.991 for the firm-adjusted gap, and Figure D.1 reports the corresponding distributions. The compression is therefore not generated solely by the inclusion of public-sector workers or by the more standardised wage structures found in parts of public employment.

Firm adjustment nevertheless does not affect every district in the same way. The correlation between the two main-sample estimates is 0.689, indicating that district rankings are only partly preserved, and the unweighted OLS slope of the firm-adjusted on the baseline estimate is 0.385. In Figure 2, 68.4% of districts lie above the 45-degree line, meaning that their estimated gap becomes less negative once firm fixed effects are introduced. A smaller group of districts lies below the line, showing that firm adjustment can also make the estimated local gap more negative.

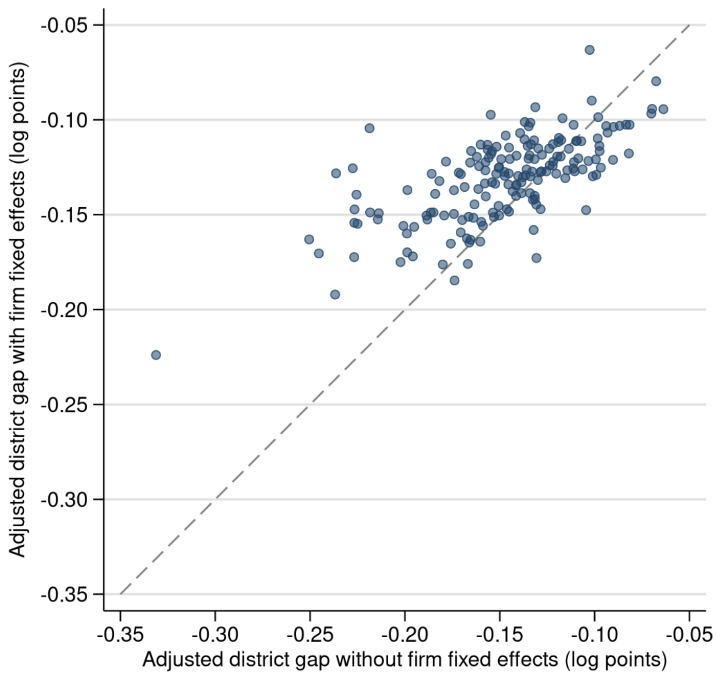

Figure 2. Baseline and Firm-Adjusted District Gender Pay Gaps

Note. Each point is one of the 174 districts. The dashed line is the 45-degree line; no fitted line is shown.

Figure 3 maps the two sets of estimates on a common scale. The baseline gaps in Panel A do not form a simple contiguous regional divide: large and small penalties are found in different parts of the country. The five largest penalties nevertheless fall in industrialised districts of the west and centre, and the five smallest in the southern and eastern periphery, which is the configuration Fuchs et al. (2021) report for Germany and Section I.C sets out, in which wide regional gaps often reflect well-paid male employment and narrow ones weak male opportunities. Panel B repeats the map for the firm-adjusted estimates, and the contrast between the two panels is the compression reported above: once persistent firm wage components are absorbed, the map becomes markedly more uniform.

The changes after firm adjustment are similarly dispersed across the country, as Figure 4 shows, and some of the largest improvements occur among districts with particularly negative baseline gaps. Paksi, Téti and Szentgotthárdi are all among the five districts with the largest baseline penalties and also among the five with the largest positive changes after firm adjustment. Because estimates in the tail are extreme partly on account of sampling error, some reversion of this kind would arise even without differential absorption, so the pattern is reported as a description of the two distributions rather than as evidence that firms matter more

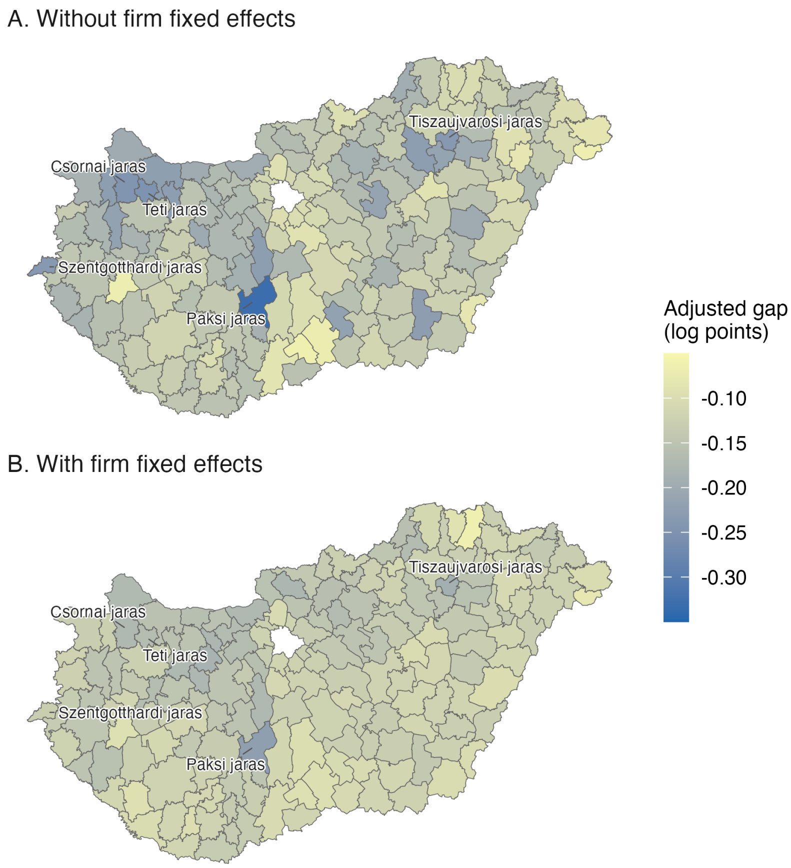

Figure 3. Adjusted District Gender Pay Gap Before and After Firm Fixed Effects

Note. Panel A reports the estimates from equation (III.2) and Panel B those from equation (III.4); more negative values indicate larger adjusted penalties for women. Both panels share one colour scale. Labels mark the five largest baseline penalties, shown in both panels. Budapest is excluded and appears unshaded.

where penalties are larger. Their paths subsequently differ: Paksi remains the district with the most negative firm-adjusted gap, whereas Szentgotthárdi moves to the centre of the distribution. These cases illustrate that firm wage components absorb a particularly large part of the initial gap in some districts, while a substantial local penalty remains in others.

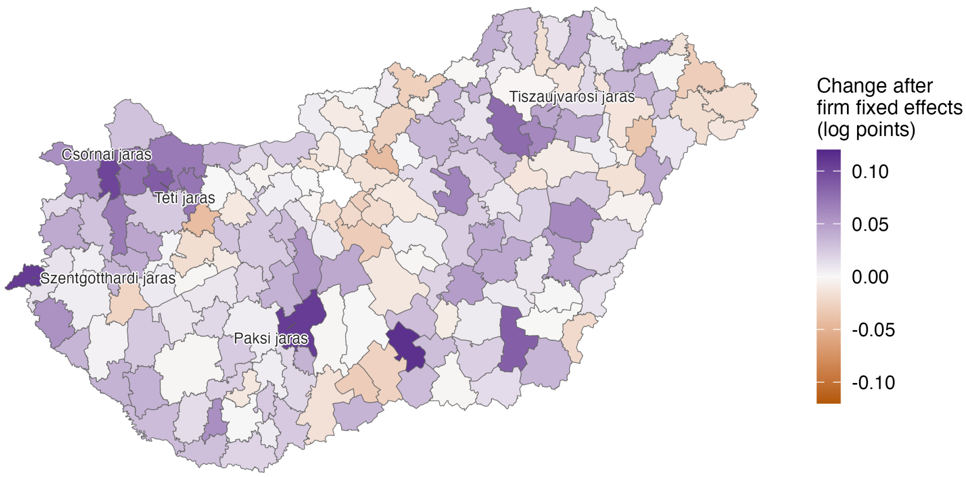

Figure 4. Change in the District Gender Pay Gap After Firm Adjustment

Note. The mapped change is the firm-adjusted minus the baseline estimate; positive values indicate that the gap becomes less negative once firm fixed effects are introduced. Labels mark the same five districts as Figure 3. See Section IV.A.

B. Local Characteristics and the Residual Gap

Table 4 relates the recovered district gaps to local workforce composition, employment structure, and institutional conditions. Models 1–3 build the firm-adjusted specification in stages. Model 4 applies the full set of covariates to the baseline gap, while Model 5 uses the private-sector firm-adjusted gap. All dependent variables are female-minus-male differences. A positive coefficient therefore indicates a smaller female penalty, whereas a negative coefficient indicates a larger one.

The workforce-composition variables initially follow the expected pattern. In Model 1, districts with a larger female higher-skill occupation share have a smaller firm-adjusted penalty, while districts in which women make up a larger share of registered jobseekers have a larger one. The respective coefficients are 0.110 and -0.177, and both are statistically significant at the 1% level. These associations do not persist once employment structure is included: in Model 2, both coefficients become small and statistically insignificant. The stepwise paths in Appendix Table C.1 locate the change, with the higher-skill share losing its association once industry concentration enters and the jobseeker share once employer concentration does. Entered on its own, the female higher-skill share yields a weighted R^2 of 0.156 against the baseline gap and 0.087 against the firm-adjusted gap. What the two measures index therefore overlaps substantially with differences in local employment opportunities and sectoral structure.

The public-sector employment share has the clearest positive association with the firm-adjusted gap. Its coefficient is 0.213 in Model 2 and rises to 0.267 in the full specification. A one-standard-deviation increase

Table 4. District-Level Regressions

	Model 1 Firm-adjusted	Model 2 Firm-adjusted	Model 3 Firm-adjusted	Model 4 Base gap	Model 5 Private-only firm-adjusted
Female higher-skill occupation share	0.110^{***}	-0.008	-0.028	0.016	-0.058
	(0.020)	(0.029)	(0.042)	(0.083)	(0.044)
Female share of registered jobseekers	-0.177^{***}	0.062	0.021	-0.061	0.044
	(0.058)	(0.063)	(0.071)	(0.137)	(0.078)
Log district employment		-0.007^{**}	-0.004	-0.005	-0.003
		(0.003)	(0.003)	(0.008)	(0.003)
Employer concentration (HHI)		0.057	0.049	-0.443	0.090
		(0.168)	(0.155)	(0.342)	(0.168)
Industry concentration (HHI)		-0.206^{***}	-0.160^{**}	-0.102	-0.173^{**}
		(0.058)	(0.067)	(0.129)	(0.071)
Public-sector employment share		0.213^{***}	0.267^{***}	0.564^{***}	0.324^{***}
		(0.069)	(0.090)	(0.173)	(0.095)
Nursery places per 100 children aged 0–2			0.000	-0.001	0.000
			(0.000)	(0.001)	(0.000)
Sewage-connected dwelling share			-0.023^{**}	-0.036^{*}	-0.019^{*}
			(0.009)	(0.021)	(0.010)
Log crimes per 1,000 inhabitants			0.002	0.013	0.002
			(0.006)	(0.014)	(0.007)
Net migration per 1,000 inhabitants			0.001	0.002^{*}	0.001
			(0.001)	(0.001)	(0.001)
Constant	-0.106^{***}	-0.100^{**}	-0.105^{**}	-0.176^{*}	-0.126^{**}
	(0.024)	(0.040)	(0.044)	(0.106)	(0.048)
R^2	0.150	0.434	0.474	0.412	0.474
Districts	174	174	174	174	174
F statistic	15.06	29.03	19.28	15.61	18.93

Note: District-level weighted least-squares regressions for 174 non-Budapest districts, weighted by district employment; heteroskedasticity-robust standard errors are reported in parentheses. Model 4 applies the full specification to the baseline gap. In Model 5, the dependent variable is recovered after Steps 1 and 2 are re-estimated for workers in an ordinary employment relationship (*fogvisz1* = 201). The mean VIF is 2.74 and the maximum is 4.94. Per-standard-deviation magnitudes quoted in the text use unrounded coefficients and standard deviations. Full stepwise entry paths for both dependent variables are reported in Appendix Table C.1 and Table C.2. $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

in the public-sector share, approximately 4.1 percentage points, is associated with a 1.10-percentage-point smaller female penalty. The estimate remains statistically significant after all ten covariates are included.

Comparing Models 3 and 4 helps locate this association. The public-sector coefficient is 0.564 for the baseline gap, equivalent to 2.32 percentage points for a one-standard-deviation increase, but 0.267 for the firm-adjusted gap. The latter is about 47% of the former. This comparison suggests that part of the baseline association is absorbed by persistent firm wage components, while a substantial association remains after those components are included. It is a descriptive comparison rather than a formal decomposition, and it does not establish that public employment itself changes local private wages.

Model 5 indicates that direct inclusion of public-sector employees in the outcome sample is unlikely to be the sole mechanical source of the association. When public-sector employees are excluded from the outcome estimation, the public-sector coefficient increases to 0.324 rather than falling towards zero. The main estimation panel is itself only about 9% public (Appendix D), so this is a smaller perturbation than the district shares in Table 2 might suggest; that the coefficient rises rather than falls is nevertheless inconsistent with the mechanical reading. This corresponds to a 1.33-percentage-point smaller private-sector penalty for a one-standard-deviation increase in the local public-sector share. The main-sample result therefore cannot be attributed solely to mixing workers on compressed public pay scales into the district gap.

The association is not the work of any one district, with a single exception. Across 174 leave-one-out regressions, the coefficient has a median of 0.267 and a standard deviation of 0.008, and omitting any district other than Paksi moves it by less than 7%. Omitting Paksi, which carries the largest adjusted penalty in the country, lowers it by about 30% to 0.186 while reducing the standard error by about two fifths, to 0.053, so the association becomes smaller and more precisely estimated rather than disappearing.

Two mechanisms remain consistent with this private-sector association. A larger public sector may provide women with stronger outside options and place pressure on private employers' wage offers. Alternatively, gendered selection between sectors may change the characteristics of the women who remain in private employment. Both mechanisms can produce the same district-level relationship. Because the analysis observes equilibrium employment and wages rather than transitions or outside offers, it cannot distinguish between them.

Industry concentration is the second stable correlate of the firm-adjusted gap. Its coefficient is -0.206 in Model 2 and -0.160 in the full specification. A one-standard-deviation increase in the industry HHI is associated with a 0.68-percentage-point larger female penalty. The estimate remains similar in the private-sector specification at -0.173. By contrast, its coefficient for the baseline gap is -0.102 and statistically insignificant. The stepwise paths in Appendix Table C.2 show the same distinction: the baseline-gap coefficient is -0.358 when industry concentration first enters but weakens after the public-sector share and institutional variables are added. The association is therefore more stable for the firm-adjusted gap than for the baseline gap.

Employer concentration does not show the same stability. Its firm-adjusted coefficient is 0.333 when it first enters, falls to 0.249 after industry concentration is added, and declines to 0.057 once the public-sector share enters. At the same step, R^2 rises from 0.380 to 0.434. In the full specification, employer concentration is unrelated to either the firm-adjusted gap or the baseline gap. The employer Herfindahl index therefore provides little independent information once broader industry and sectoral structure are taken into account: removing the public-sector share from the full specification returns it to 0.167 and strengthens industry concentration to -0.250 (Appendix Table G.2), so part of what the two concentration measures would otherwise register is shared with local public employment. Section I.B established that concentration is only one observable source of employer wage-setting power, so this is evidence about the measure rather than about employer power.

Log district employment follows a similar path. It is negative and significant in Model 2 but is not distinguishable from zero once the institutional variables are added, where a district 10% larger is associated with a difference of roughly 0.04 percentage points. Section I.C sets the possibility that thicker local markets improve women's matching and outside options against evidence that urban scale does not predict a stable regional gender pay gap. At the district level, with firm wage components absorbed, the full specification does not provide statistically precise evidence in either direction.

Among the institutional variables, the sewage-connected dwelling share is negatively associated with the firm-adjusted gap. Its coefficient of -0.023 implies that a one-standard-deviation increase is associated with a 0.46-percentage-point larger female penalty. This runs against the direction anticipated in Table 2, although not against the reading introduced in Section I.C, where Fuchs et al. (2021) find that stronger local economies need not narrow gender pay gaps when the high-paying opportunities are concentrated in male-dominated firms and industries; that is also the pattern Panel A of Figure 3 displays. A second interpretation consistent with the result is that wages in less-developed districts are more compressed towards statutory wage floors, leaving less room for gender differences to widen, while more-developed districts have a broader wage distribution. This second reading remains a conjecture: the second-stage regression does not measure the local incidence of minimum-wage binding or establish that wage compression generates the association.

The remaining institutional variables provide little evidence of independent relationships with the firm-adjusted gap. Nursery provision and recorded crime have coefficients close to zero. Net migration is positive, but its p-value is 0.128 in the full main-sample specification. Overall, the full model has a weighted R^2 of 0.474 against the firm-adjusted gap, with the most consistent associations concentrated in public-sector employment, industry structure, and infrastructure provision. The sewage-connected dwelling share and, against the baseline gap, nursery provision both run against the direction anticipated in Table 2: better-provided districts record wider adjusted gaps, not narrower ones. The nursery association does not survive the full specification and is absent from the firm-adjusted gap, so little weight can be placed on it, but the common sign of the two measures is consistent with the reading in Section I.C that stronger local economies need not narrow gender pay gaps where the high-paying opportunities are concentrated in male-dominated firms and industries. A further possibility specific to childcare is compositional: where subsidised provision raises mothers' employment, as Lovász and Szabó-Morvai (2019) find for Hungary, it also changes which women are observed in work, and Section I.A notes that employee-only estimates describe inequality among workers in employment rather than across the full working-age population.

All standard errors reported in Table 4 are heteroskedasticity-robust. Section V reports the alternative inference procedures and the earlier covariate window described in Section III.C.

V. Robustness Checks

Table 5. Summary of Robustness Checks

Check	Conclusion	Reported in
Step 1 re-estimated on the Step 2 estimation sample	Variance reduction is unchanged.	Table E.1
Individual-level industry control removed on the common valid-industry sample	Estimated compression increases to 74.7%, against 68.8% in the main specification.	Table E.1
Observations without a usable industry code restored	Relative to specification C, restoring records without a usable industry code lowers the compression from 74.7% to 72.8% and the public-sector coefficient from 0.267 to 0.195; industry concentration is unchanged (Appendix E).	Table E.1–Table E.2
Steps 1–3 re-estimated on ordinary employment relationships	Compression is almost unchanged and the public-sector association strengthens.	Table 3, columns (3) and (4); Table 4, Model 5; Figure D.1
Public sector defined using codes 202–203 rather than 202–204	Excluding armed services from the public-sector definition leaves the coefficient unchanged in sign and significance on all four outcomes; on the firm-adjusted gaps its per-standard-deviation value moves by 3.2% and 1.2% (Table F.4).	Table F.4
Influential districts removed singly, jointly, and by Cook's distance	The public-sector association attenuates and industry concentration strengthens, but neither changes sign.	Table G.2
County-clustered inference and wild-cluster bootstrap	Only the public-sector association clears the 5% level under the bootstrap.	Table G.1
District characteristics measured over 2003–2007	The public-sector association strengthens; industry concentration and infrastructure weaken but retain their signs.	Table G.1

Note: Each check re-estimates the design of Section III on the corrected panel. Coefficients, standard errors, and sample sizes are reported in the appendix tables listed.

Table 5 summarises the completed robustness exercises, which cover the estimation samples, industry-code coverage, outcome composition, the definition of public employment, influential districts, inference, and the timing of covariate measurement. The checks broadly preserve the principal point-estimate patterns, subject to the inference and industry-coverage qualifications below. One check strengthens the dispersion result. Removing the individual-level industry control on the common valid-industry sample raises the variance reduction from 68.8% to 74.7%, and the increase arises entirely in the baseline step: because each firm carries a single industry code, the industry indicators are collinear with firm identity once firm fixed effects are included, so the firm-adjusted district gaps are virtually unchanged and only the baseline gaps become more dispersed. What the main specification attributes to the control set is therefore compression that firm fixed effects would otherwise perform, and the estimate reported in Section IV is conservative. The remaining sample, measurement, and influence checks alter magnitudes without overturning the principal signs or conclusions.

Alternative inference provides the main qualification. County clustering leaves the point estimates unchanged but moves the public-sector association from the 1% to the 5% level. With only nineteen clusters, the cluster-robust asymptotics are unreliable, so wild-cluster bootstrap p-values are also reported: 0.005 for public-sector employment, 0.066 for industry concentration, and 0.054 for sewage-connected dwellings. The public-sector association therefore remains statistically significant under county clustering and wild-cluster bootstrap inference. The infrastructure and industry-concentration estimates retain their signs but are only marginally significant under the wild bootstrap. In Section IV.B, "stable" refers to the industry-concentration point estimate across specifications, not to its significance under every inference procedure.

These checks do not address the geographical definition of the labour market. The district is assigned by the worker's place of residence rather than by workplace, so commuting across district boundaries cannot be modelled with these data; Section VI takes up what this implies for the interpretation of the district estimates. Further exercises that the data allow and are left to future work include re-estimating the design on a coarser spatial aggregation or on functional labour-market areas and weighting the second stage by the estimated precision of the recovered district gaps.

VI. Discussion

This study examined how persistent firm wage components relate to the spatial dispersion of adjusted gender pay gaps and which local labour-market characteristics covary with the remaining variation. Using 2003–2017 Admin3 linked employer–employee records for 174 non-Budapest residence districts, the analysis compared recovered district-by-gender wage differentials before and after adding firm fixed effects. Adjusted female penalties were present in every district, while firm fixed effects absorbed 68.8% of their baseline cross-district variance without eliminating local differences. The private-sector re-estimation produced similar compression, showing that this pattern was not generated solely by public-sector wage structures. Public-sector employment had the clearest association with a smaller remaining female penalty, whereas greater industry concentration was associated with a larger penalty. These patterns remain descriptive: firm fixed effects do not identify the underlying mechanism, and the magnitude of the public-sector association is sensitive to industry-code coverage.

The large compression after firm adjustment indicates that persistent firm wage components are closely connected to the geography of adjusted gender pay gaps in Hungary. This result extends firm-level studies showing that women's access to high-premium firms and gender-specific firm premia account for material portions of aggregate gaps, although their relative importance varies with the estimand and counterfactual (Card et al., 2016; Bruns, 2019; Masso et al., 2022). It also complements broader evidence that spatial wage disparities reflect workforce sorting and establishment geography (Combes et al., 2008; Fuchs et al., 2021). The present analysis connects these literatures by asking how the distribution of district-specific gaps changes when persistent firm wage components are absorbed. Its 68.8% variance compression cannot be compared directly with decomposition shares reported in earlier studies, because the outcome here is cross-district dispersion rather than the average national wage gap. The specification also cannot determine which of the three sorting-type mechanisms (spatial sorting, gender sorting across firms, or occupation–firm sorting) accounts for the compression; within-firm wage setting is excluded by construction, since the common firm fixed effect does not vary by gender within a firm. Within these boundaries, the finding documents a strong accounting relationship between persistent firm wage components and cross-district dispersion, while preserving a distinct role for local labour-market context.

The public-sector result suggests that public employment forms part of the local wage-setting environment beyond the direct inclusion of workers on public pay scales. A one-standard-deviation increase in the local public-sector employment share was associated with a 1.10-percentage-point smaller firm-adjusted female penalty, while the coefficient rose from 0.267 to 0.324 when the wage gaps were re-estimated among private-sector workers. This pattern is consistent with evidence that centralised public wage-setting compresses wage distributions and that gender pay gaps in public-service sectors can be smaller than in manufacturing (Blackaby et al., 2018; Rattsø and Stokke, 2020; Fuchs et al., 2021). Its persistence in the private-sector outcome indicates that direct wage compression among public employees cannot account for the association alone. Stronger outside options for women could place pressure on private wage offers, an interpretation compatible with the corporate wage responses following Hungary's 2002 public-pay reform documented by Telegdy (2018). Alternatively, gendered selection between sectors could change the composition of women remaining in private employment. The present design observes neither transitions nor outside offers and cannot distinguish these mechanisms. Moreover, restoring observations without valid industry codes reduces the coefficient from 0.267 to 0.195. Therefore, public employment is best treated as a contextual local institution whose association is robust in direction, although neither its mechanism nor its magnitude is firmly identified.

Beyond public employment, the remaining spatial pattern is more closely associated with industry composition than with simple market size or employer concentration. A one-standard-deviation increase in industry concentration was associated with a 0.68-percentage-point larger firm-adjusted female penalty; the coefficient of -0.160 was similar in the private-sector specification at -0.173, although its wild-cluster bootstrap p-value was 0.066. The direction is partially consistent with Lovász (2008), who associates stronger product-market competition with smaller within-firm gaps, but the measures are not equivalent: the present industry HHI captures the concentration of district employment across broad industries, not product-market competition or employer wage-setting power. Evidence from Petrongolo and Ronchi (2020) that goods-producing industries are more spatially concentrated than female-intensive services suggests a complementary compositional interpretation, although the present analysis does not test that mechanism. Employer HHI and district employment size showed no independent association in the full specification, which limits conclusions about these particular proxies rather than about employer power generally. The negative sewage-coverage coefficient was also only marginal under wild-bootstrap inference. One interpretation, consistent with Fuchs et al. (2021), is that stronger local economies may contain male-dominated high-wage opportunities; a second, more conjectural reading is that statutory wage floors compress earnings more strongly in less developed districts. Neither explanation is directly observed. Therefore, the results indicate that these measures carry different conditional associations in the present specification while leaving its underlying channel unresolved.

The findings carry related methodological and policy implications. For researchers, the large compression after firm adjustment indicates the value of reporting spatial gaps before and after persistent firm wage components are absorbed; otherwise, firm-related variation may be attributed to geography alone. Analyses should also distinguish industry composition, employer concentration, and public employment, since these measures did not carry interchangeable associations. For regional-development and gender-equality policy analysts, the results suggest evaluating changes in local employment structure not only by the number of jobs created, but also by how employment opportunities are distributed across firms, sectors, women, and men. Public employment warrants attention as a contextual local institution because its association persists in the private-sector outcome, although the evidence does not show that expanding public employment would narrow private-sector gender pay gaps. Similarly, the industry-concentration result does not establish diversification as an effective intervention. Therefore, employer and institutional composition can serve as diagnostic dimensions in policy evaluation, while causal recommendations require exogenous changes, transition data, or other identification strategies.

Previous firm-level research has shown that employer allocation and gender-specific firm premia account for material portions of aggregate gender wage gaps (Card et al., 2016; Bruns, 2019; Masso et al., 2022), while regional studies have documented geographic heterogeneity without measuring how persistent firm wage components map onto small-area dispersion (Nisic, 2017; Fuchs et al., 2021). The present study connects these strands by showing that firm fixed effects absorb a large share of the cross-district variance in adjusted gender pay gaps. A second empirical addition concerns public employment: its association with a smaller adjusted female penalty remains when the outcome is estimated among private-sector workers. This finding is inconsistent with an account based solely on the direct inclusion of public-sector employees, although it cannot distinguish outside-option responses from gendered selection between sectors. The contribution is necessarily bounded. It provides an evidence-based spatial accounting of firm-related and institutional correlates across non-Budapest Hungarian districts, without offering a causal decomposition of their effects or a new theoretical model.

Several limitations define the evidential boundary of these findings. First, firm fixed effects operate as an accounting device that can absorb spatial sorting, gender sorting across firms, and occupation–firm sorting, without separating their respective roles; within-firm wage setting is not absorbed by construction, since the common firm effect does not vary by gender, and any such component remains in the recovered gender gap regardless of firm adjustment. Second, districts are assigned by residence rather than workplace. Cross-district commuting and the unobserved locations of establishments within multi-site firms may therefore blur the relevant employer geography. The relatively even distribution of armed-services employment across all 174 districts provides limited confirmation that the identifier behaves as a residence measure, but it does not validate districts as closed labour markets. and the resulting generated-regressor problem is only partly addressed by employment weighting, which proxies rather than measures their precision, while the 174-district sample limits model complexity and spatial inference. The second-stage relationships also remain jointly determined associations. Fourth, industry-code coverage is substantially lower among public-sector records, making district missingness and public-sector employment closely overlapping measures; the direction of the public-sector association is more stable than its magnitude. Moreover, public employment is defined by legal relationship rather than ownership, so employees of state-owned enterprises are classified as private and state-ownership presence may be understated, although the legal-relationship-based public-employment measure is consistently defined. Finally, the employee-only, non-Budapest sample does not correct for selection into employment. Future research could combine workplace or commuting-based geographies with employment transitions, exogenous institutional changes, and precision- or spatially informed second-stage methods.

VII. Conclusion

This study investigated spatial variation in adjusted gender pay gaps. It assessed how much of the cross-district dispersion was absorbed by firm fixed effects and examined local labour-market characteristics associated with the remaining district variation. The study drew on linked employer–employee administrative records covering 174 non-Budapest Hungarian residence districts between 2003 and 2017. A pooled hourly-wage specification with common worker controls was first used to estimate district-by-gender effects. Firm fixed effects were subsequently introduced, allowing the dispersion of the recovered gaps to be compared across specifications. The baseline and firm-adjusted gaps were then examined as district-level outcomes in relation to local labour-market characteristics. The resulting comparisons are descriptive and do not isolate individual firm-side mechanisms or establish causal effects.

The results show that firm fixed effects absorbed 68.8% of the baseline cross-district variance in adjusted gender pay gaps, substantially narrowing their spatial distribution. Nevertheless, the estimated female penalty remained negative in all 174 districts. A comparable 67.0% variance reduction in the private-sector re-estimation indicates that this compression was not generated solely by the inclusion of public-sector workers. Among the characteristics associated with the remaining variation, public-sector employment provided the clearest pattern. A higher district public-sector share was associated with a smaller firm-adjusted female penalty, including when the gap was estimated only among private-sector workers. Although this association was robust in direction, its magnitude was sensitive to industry-code coverage, and the underlying mechanism was not identified. Greater industry concentration was associated with a larger remaining penalty, although this evidence was weaker under alternative inference procedures. Overall, the findings link much of the initial spatial dispersion to persistent firm wage components while identifying public-sector employment and industry structure as the principal correlates of the residual variation.

The combined evidence indicates that spatial differences in adjusted gender pay gaps cannot be understood as purely geographic variation. Firm fixed effects absorbed most of the baseline cross-district variance, showing that persistent firm wage components are closely connected to the observed spatial distribution. Yet firm adjustment neither removed the female penalty in any district nor eliminated systematic associations with local employment structure. The geography of adjusted gender pay gaps therefore reflects two analytically distinct layers: substantial variation linked to the firms employing district residents, alongside residual variation associated with the institutional and industrial composition of districts. This distinction is central to interpreting the study's estimand. The reduction in dispersion describes a firm-related accounting relationship; the residual gaps remain associated with local labour-market conditions, without identifying the mechanisms within either layer.

The public-sector result provides a more specific interpretation of the residual district variation. Its persistence when wage gaps were re-estimated among private-sector workers is inconsistent with an explanation based solely on the direct inclusion of employees on public pay scales. Within this study, public employment is therefore best treated as a contextual local institution associated with the wider wage-setting environment. This interpretation does not imply that expanding public employment would reduce private-sector gender pay gaps. The observed relationship remains consistent with either stronger outside options for women or gendered selection between sectors. The present design cannot distinguish these alternatives because it contains no evidence on employment transitions or outside offers. The finding consequently identifies a robust contextual pattern while leaving its underlying process unresolved.

This study contributes to research on spatial gender pay inequality by clarifying how persistent firm wage components should enter the interpretation of local gaps. The substantial reduction in cross-district variance after firm adjustment shows that spatial estimates can contain firm-related wage variation, even under a common worker-control framework. Reporting district gaps before and after firm fixed effects are introduced can therefore show how much observed spatial dispersion is absorbed by persistent firm components. This comparison must remain distinct from decompositions of the average national gender wage gap because its estimand is cross-district dispersion. For researchers, the study provides a basis for interpreting spatial wage inequality without treating all observed local variation as a place-specific effect.

These findings should be interpreted within two main limitations. First, the common firm fixed effect is an accounting device that can absorb spatial sorting, gender sorting across firms and occupation–firm sorting without separating their contributions; it does not absorb within-firm wage setting. The observed variance compression therefore cannot be assigned to a specific firm-side mechanism. Future research could use employment transitions, worker mobility or exogenous institutional changes to distinguish the channels underlying this pattern. Second, districts are defined by workers' residence, so cross-district commuting and the unobserved establishment locations of multi-site firms may blur the relevant employer geography. Research using workplace- or commuting-based areas, functional labour markets or alternative spatial aggregations could test the sensitivity of the findings to geographic definition. Within these boundaries, the study provides evidence that persistent firm wage components are closely associated with spatial dispersion while identifying local institutional and industrial correlates of the variation that remains.

A. Data Construction and the Analysis Panel

This appendix documents how the analysis panel is built from the Admin3 records and reports the full descriptive statistics of which Table 1 presents an abridged version. Section II.B states the restrictions in substance; what follows gives the order in which they are applied, the variables and code values that implement them, and the observation counts at each stage.

A.1. Restrictions

The raw extract covers the 2003–2017 waves of the Admin3 linked employer–employee panel at monthly frequency, with the worker and recorded firm in the month as the unit of observation. Workers are retained if aged 18 to 65 (`kor`). The sample is then defined using `fogvisz1`, the legal-relationship classification recorded for the reference employment relationship on the fifteenth day of the month. Non-missing values below 205 are retained, which in this extract are exactly the four regular wage-setting relationships: ordinary private-law employment (201), civil servants and related officials (202), public employees (203), and armed services (204). The Admin3 H1 catalogue assigns public-work relationships (*közfoglalkoztatás*) the separate code 215; self-employment, cooperative membership, business partnerships, and the other non-standard relationships likewise carry codes above 204. Thus `fogvisz1 < 205` is a relationship-based operational definition of the regular wage-setting relationships studied in the paper: it excludes records classified in Admin3 as public work or another non-standard relationship, rather than attempting to classify individuals from every external programme-participation source. The public-sector share used in the second stage is defined consistently within this same classification, as records carrying codes 202–204 divided by all records carrying codes 201–204, and is computed on the eligible-wage sample defined in A.3 rather than on the cleaned analysis panel.

Observations are next dropped where monthly earnings (`w1`) are missing or non-positive, days worked (`nap1`) are missing or zero, weekly hours (`wh1`) are missing, or the residence-district identifier (`jaras_al`) is missing or takes invalid code 999. The three earnings and working-time variables are the components of the hourly-wage measure in equation (II.1). These restrictions remove 7,577,310 observations. The FEOR file is merged one-to-one on worker and month (`anon` and `t`); every remaining master observation is matched, so the merge removes no analysis record. Invalid occupation codes (`feor1_h2 < 0`) then remove 99,462 observations, and invalid industry codes (`teaor1_h1 < 0`) remove a further 2,265,850. Finally, the 23 Budapest districts are excluded. The resulting cleaned analysis panel contains 221,278,959 worker–firm–month observations.

A valid firm identifier is not imposed when the cleaned panel is saved because it is not required by Step 1. The cleaned panel contains 44,734 observations with a missing firm identifier; all also have a missing industry code and are therefore absent from the Step 1 complete-case sample before Step 2 is estimated.

Variable names differ between the source extract and the saved regression panel. After invalid values are removed, `feor1_h2` is cloned to `occ_h2`, with code 37 (artistic, cultural, sports and religious occupations) relabelled as 0 solely to serve as the omitted reference category. The variable `teaor1_h1` is cloned to `industry1` while retaining the labelled 1–17 TEÁOR sections. The wage regressions use `occ_h2` and `industry1`; the source names appear only in the construction steps.

Table A.1. Sample Construction

Step	Restriction	Observations removed	Observations remaining
0	Raw Admin3 extract, 2003–2017	—	931,327,200
1	Workers aged 18–65	326,269,218	605,057,982
2	Employment-status codes 201–204 retained	327,936,532	277,121,450
3	Positive earnings; days and hours observed; residence district observed and not code 999	7,577,310	269,544,140
4	FEOR merge (master observations unmatched)	0	269,544,140
5	Invalid occupation codes	99,462	269,444,678
6	Invalid industry codes	2,265,850	267,178,828
7	Budapest districts excluded	45,899,869	221,278,959
	Cleaned analysis panel		**221,278,959**
8	Complete cases on the Step 1 controls	48,633,813	172,645,146
	Step 1 estimation sample		**172,645,146**
9	Singleton firms, not estimable with firm fixed effects	11,869	172,633,277
	Step 2 estimation sample		**172,633,277**

Note: Counts follow the production order in `rerun1_clean_and_steps12_FIXED.do` and the corresponding Rerun1 log. The employment-status filter retains `fogvisz1` codes 201–204 in one operation, so records carrying the public-work relationship code 215 have no separate removal count. Step 3 combines 1,212,098 observations with missing earnings or district, 606,166 observations with district code 999, and 5,759,046 observations failing the positive-earnings/days/hours requirement. Occupation and industry deletions are counted before Budapest is removed.

A.2. Estimation Samples

Steps 1 and 2 estimate on complete cases for the controls in the wage equation. Relative to the cleaned analysis panel, Step 1 excludes 48,633,813 observations. Missing industry codes account for 48,207,034 of these observations (21.8% of the cleaned panel and 99.1% of the Step 1 gap). The remaining 426,779 observations are valid-industry records that lack the quasi-education and/or occupation controls; the number is a residual union and should not be added to marginal missingness counts.

Step 2 additionally excludes 11,869 singleton observations, approximately 0.007% of the Step 1 sample, for which a firm fixed effect cannot be estimated. Section V reports both the comparison re-estimated on the Step 2 sample and the design re-estimated without the industry control.

A.3. The Eligible-Wage Sample

Two second-stage covariates—the public-sector employment share and employer concentration—are constructed on a wider non-Budapest extract of 223,286,902 eligible wage records. It applies the same age, employment-status, earnings, working-time, and residence-district restrictions as the cleaned panel, but it is formed before occupation and industry codes restrict the sample. The public-sector share defined in A.1 is computed on this universe, so its denominator is all retained records with fogvisz1 codes 201–204 rather than the subset carrying valid occupation and industry codes.

The apparent 357,369-observation discrepancy arises because the occupation and industry deletion counts in Table A.1 are national counts recorded before Budapest is removed, whereas 223,286,902 already excludes Budapest. Of the 2,365,312 observations removed sequentially for invalid occupation or industry codes nationwide, 357,369 are in Budapest. Thus the non-Budapest invalid-code deletion is 2,007,943, and the samples reconcile exactly:

$$223{,}286{,}902 - 2{,}007{,}943 = 221{,}278{,}959.$$

The eligible-wage sample and cleaned analysis panel therefore use the same age, employment-status, wage, working-time, and geographic restrictions; they differ only because the former does not require valid occupation or industry codes.

A.4. The Quasi-Education Proxy

Education is not observed as a qualification record in the analysis extract. The Step 1 and Step 2 control `max_educ2` is instead constructed from the occupation histories in `admin3_feor.dta`. The FEOR file is merged on worker and month (`anon` and `t`). For years before 2010 the construction uses the FEOR 2003 fields `feor1_2003_2`, `feor1_2003_3` and `feor1_2003_4`; from 2010 onward it uses the corresponding FEOR 2008 fields.

Each worker-month is first assigned to one of four latent education-requirement levels. At the broad one-digit level, FEOR groups 1–2 map to level 4, group 3 to level 3, groups 4–8 to level 2, and group 9 to level 1. Pre-specified three- and four-digit occupation exceptions then override this broad mapping where the detailed occupation implies a different requirement; the exact exception lists are in `rerun1_clean_and_steps12_FIXED.do` and are fixed independently of wages or the estimated gender gaps. Within each worker's observed sequence, the code takes the cumulative maximum of this four-level assignment through month t. It then collapses levels 1–2 to `max_educ2 = 1` (elementary/vocational), level 3 to `max_educ2 = 2` (secondary), and level 4 to `max_educ2 = 3` (tertiary).

The proxy is missing when no FEOR record observed up to that worker-month can be mapped to a latent level; this affects about 0.25% of the cleaned panel. `max_educ2` should therefore be interpreted as a cumulative occupation-implied education requirement, not directly observed educational attainment. Because `max_educ2` is a cumulative maximum over a worker's observed history whereas the occupation control `occ_h2` is the contemporaneous code, the two are correlated but not redundant: the former retains information about occupational attainment that the current job need not reflect. The proxy also shares its FEOR source with the district female higher-skill occupation share, but the two enter at different levels: `max_educ2` is a worker-month control in Steps 1 and 2, whereas `female_highskill_share` is a district-level covariate in Step 3.

Table A.2. Descriptive Statistics, Cleaned Analysis Panel

	Men	Women	Total
Panel A. Sample structure			
Worker–firm–month observations	114,512,444	106,766,515	221,278,959
Years covered		2003–2017 (15 years)	
Distinct firms		589,926	
Districts (*járás*)		174	
District-year cells		2,610	
District-by-gender cells		348	
Female observations	—	106,766,515	106,766,515 (48.2%)
Observations per district-gender cell (mean / median / min / max): 635,859 / 435,779 / 108,896 / 3,339,170			
Panel B. Continuous variables			
N	114,512,444	106,766,515	221,278,959
Mean log hourly wage	6.448	6.356	6.404
SD log hourly wage	0.652	0.585	0.622
*p*10	5.675	5.676	5.676
*p*50	6.377	6.310	6.342
*p*90	7.320	7.101	7.213
Mean age	39.9	41.1	40.5
SD age	11.0	10.6	10.8
Median age	39.0	42.0	40.0
Panel C. Composition (column percentages)			
Education			
Elementary / vocational	64.29	44.25	54.62
Secondary	13.40	31.56	22.17
Tertiary	22.31	24.18	23.21
Occupation skill group			
Higher-skill occupations	30.74	55.73	42.80
Lower-skill occupations	69.26	44.27	57.20
Industry (NACE section)			
A Agriculture	5.45	2.27	4.07
C Mining	0.36	0.08	0.24
D Manufacturing	32.72	27.77	30.57
E Utilities	3.15	1.31	2.35
F Construction	9.30	1.77	6.03
G Trade, repair	13.33	19.60	16.05
H Accommodation, food	2.21	4.43	3.17
I Transport, storage, post	11.79	6.16	9.35
J Financial intermediation	1.14	3.77	2.28
K Real estate, business services	10.36	11.73	10.96
L Public administration, defence	3.86	6.55	5.03
M Education	1.64	4.14	2.73
N Health and social work	1.85	7.30	4.21
O Other community, personal services	2.75	3.11	2.90

Note: Cleaned Admin3 panel, 2003–2017, Budapest excluded. Higher-skill occupations are harmonised FEOR codes below 34, the same threshold used to construct the female higher-skill share in the second stage. Industry sections B, P and Q are omitted (each below 0.1%). Education composition excludes observations without the quasi-education proxy (about 0.25%). Occupation percentages exclude 525 observations with a missing FEOR code (217 men and 308 women). Industry percentages exclude 48,207,034 observations with a missing NACE/TEÁOR code (21.8% of the cleaned panel).

B. District Characteristics: Construction, Sources and Diagnostics

This appendix documents the ten district characteristics entering the second stage. Table B.1 gives the construction of each measure, the original variable code in its source database, and the mechanism it is intended to index. Table B.2 reports the collinearity diagnostics that governed the screening described in Section III.D, and Table B.3 reports the pairwise correlations among the retained measures. Two source databases are used: the Admin3 linked employer–employee panel, from which the employment-based measures are constructed directly, and T-STAR, the KSH settlement-level statistical database harmonised by HUN-REN KRTK, whose records are summed from settlements to districts before ratios are formed.

Table B.1. Second-Stage Covariates: Construction, Sources, and Sample Basis

Variable	Construction	Source and original variable code	Dimension indexed
Workforce composition			
Female higher-skill occupation share	Employed women in higher-skill occupations divided by all employed women; higher-skill is harmonised FEOR below 34, the threshold used in Table 1	Admin3 cleaned analysis panel: `occ_h2`; `female = 1 - ferfi`	Local supply of skilled female labour
Female share of registered jobseekers	Female registered jobseekers divided by all registered jobseekers	T-STAR: `TAAS003` (female jobseekers), `TAAS001` (total)	Women's outside options relative to men's
Job-opportunity and employer structure			
Log district employment	Log of district employment; main measure from T-STAR, cross-checked against Admin3 residence-based worker counts (correlation 0.994 in logs)	T-STAR: `al_03`; Admin3: unique workers by `jaras_al`	Local market thickness
Employer concentration (HHI)	For each district-year, the employed residents attached to each firm give firm shares of district employment; the index is the sum of squared shares, then averaged over 2003–2017	Admin3: `vallazon`, `jaras_al`	Employer concentration, one observable source of employer wage-setting power (Section I.B)

Continued on next page

Table B.1 continued

Variable	Construction	Source and original variable code	Dimension indexed
Industry concentration (HHI)	Herfindahl index of one-digit TEÁOR employment shares, computed among observations with a valid industry code; district-year values averaged over 2003–2017	Admin3: `industry1`	Breadth of the local industry mix
Institutional and socio-economic environment			
Public-sector employment share	Share of employment in a public legal relationship: civil servants, public employees and armed services. Employees of state-owned enterprises work under ordinary contracts and count as private	Admin3: `fogvisz1` = 202 / 203 / 204	Public employment as a contextual local institution (Section I.C)
Nursery places per 100 children aged 0–2	Nursery places across all providers, divided by the permanent population aged 0–2	T-STAR: `TAAJ302` "bölcsődei férőhelyek összesen"; denominator `TAAA201` "Állandó népességből a 0–2 évesek száma"	Local formal childcare capacity
Sewage-connected dwelling share	Dwellings connected to the public sewage network divided by the housing stock	T-STAR: `TAAH108` "Közcsatornahálózatba bekapcsolt lakások száma"; denominator "Lakásállomány"	Local development and public-service provision
Log crimes per 1,000 inhabitants	Log of registered crimes recorded by place of commission, per 1,000 inhabitants	T-STAR: `TAAX000` "Ismertté vált közvádas bűncselekmények"	Local social conditions
Net migration per 1,000 inhabitants	Permanent in-migration minus permanent out-migration, per 1,000 inhabitants	T-STAR: `TAAB102` "Állandó odavándorlások", `TAAB103` "Állandó elvándorlások"	Local population change through migration

Note: $N = 174$ non-Budapest districts. T-STAR records are settlement-level and are summed to districts before ratios are formed; district-year values are then averaged over 2003–2017. The T-STAR employment measure is unavailable for 2004–2006, so its district average uses the twelve available years. The Admin3-based measures are not all built on the same rows: the public-sector share and employer concentration use the wider eligible-wage sample of 223,286,902 observations, formed before occupation and industry codes restrict the analysis panel; the female higher-skill share uses the cleaned analysis panel; industry concentration is computed only on observations carrying a valid industry code. The recorded share of armed-services employment shifts to a higher level from 2012 onwards; Section V reports the second stage with those observations excluded from the public-sector measure.

Table B.2. Variance Inflation Factors, Final Specification

Variable	VIF
Female higher-skill occupation share	4.94
Public-sector employment share	3.48
Employer concentration (HHI)	3.46
Female share of registered jobseekers	2.94
Log district employment	2.70
Net migration per 1,000 inhabitants	2.61
Industry concentration (HHI)	2.06
Nursery places per 100 children aged 0–2	1.91
Sewage-connected dwelling share	1.87
Log crimes per 1,000 inhabitants	1.39
Mean	**2.74**

Note: $N = 174$ non-Budapest districts. Factors are computed on the unweighted design matrix, as is conventional. No covariate exceeds the warning threshold of five described in Section III.D.

Table B.3. Pairwise Correlations Among the Second-Stage Covariates

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
(1) Female higher-skill share	1.00									
(2) Female jobseeker share	0.36	1.00								
(3) Log district employment	0.62	0.50	1.00							
(4) Employer HHI	−0.21	**−0.71**	−0.47	1.00						
(5) Industry HHI	−0.58	−0.14	−0.20	−0.03	1.00					
(6) Public-sector share	0.51	−0.32	0.12	0.46	−0.53	1.00				
(7) Nursery places	0.42	0.47	0.58	−0.40	−0.08	0.08	1.00			
(8) Sewage share	0.46	0.34	0.48	−0.47	−0.01	0.01	0.48	1.00		
(9) Log crimes	0.42	0.17	0.45	−0.06	−0.28	0.29	0.27	0.22	1.00	
(10) Net migration	0.47	0.57	0.34	−0.61	−0.25	−0.19	0.18	0.36	0.06	1.00

Note: $N = 174$ non-Budapest districts; Pearson correlations, unweighted across districts. One pair exceeds the inspection threshold of 0.7 in absolute value described in Section III.D: employer concentration and the female share of registered jobseekers, at -0.715 (shown in bold). Both carry variance inflation factors below the warning threshold, 3.46 and 2.94 (Table B.2), and both are retained. The largest factor in Table B.2 belongs to the female higher-skill occupation share, whose largest single correlation is 0.624.

C. Second-Stage Stepwise Build-Up

This appendix reports the full entry paths behind Table 4. Each covariate enters cumulatively, in the order in which its domain is introduced in Section III.D: the two workforce-composition measures first, then the four job-opportunity and employer-structure measures, then the four institutional and socio-economic measures. Table C.1 gives the path for the firm-adjusted gap and Table C.2 the path for the baseline gap. The two tables use identical covariates, weights, and inference and differ only in the dependent variable. Columns (2), (6), and (10) of Table C.1 are Models 1, 2, and 3 of Table 4, and column (10) of Table C.2 is Model 4.

Two features of the paths bear on the discussion in Section IV.B. The workforce-composition measures lose their associations at identifiable steps rather than gradually: the female high-skill occupation share when industry concentration enters in column (5), and the female share of registered jobseekers when employer concentration enters in column (4). Several associations that appear in intermediate columns also fail to survive the full specification—employer concentration in columns (4) and (5) of Table C.1, and employer concentration and nursery provision in columns (7) to (9) of Table C.2, the latter discussed in Section VI. The paths are therefore reported in full rather than summarised.

Table C.1. Stepwise Build-Up of the Second Stage: Firm-Adjusted Gap

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Female high-skill occupation share	0.080^{***}	0.110^{***}	0.171^{***}	0.163^{***}	0.043	−0.008	−0.009	0.020	0.019	−0.028
	(0.017)	(0.020)	(0.027)	(0.028)	(0.031)	(0.029)	(0.029)	(0.032)	(0.032)	(0.042)
Female share of registered jobseekers		-0.177^{***}	-0.126^{**}	−0.030	−0.040	0.062	0.073	0.040	0.041	0.021
		(0.058)	(0.055)	(0.097)	(0.079)	(0.063)	(0.065)	(0.066)	(0.067)	(0.071)
Log district employment			-0.010^{***}	-0.008^{**}	−0.004	-0.007^{**}	-0.006^{**}	-0.006^{**}	-0.006^{*}	−0.004
			(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Employer concentration (HHI)				0.333^{**}	0.249^{*}	0.057	0.046	−0.057	−0.054	0.049
				(0.167)	(0.134)	(0.168)	(0.173)	(0.184)	(0.190)	(0.155)
Industry concentration (HHI)					-0.279^{***}	-0.206^{***}	-0.200^{***}	-0.175^{***}	-0.176^{***}	-0.160^{**}
					(0.049)	(0.058)	(0.062)	(0.065)	(0.066)	(0.067)
Public-sector employment share						0.213^{***}	0.225^{***}	0.217^{***}	0.219^{***}	0.267^{***}
						(0.069)	(0.079)	(0.076)	(0.073)	(0.090)
Nursery places per 100 children aged 0–2							−0.000	0.000	0.000	0.000
							(0.000)	(0.000)	(0.000)	(0.000)
Sewage-connected dwelling share								-0.024^{***}	-0.024^{**}	-0.023^{**}
								(0.009)	(0.009)	(0.009)
Log crimes per 1,000 inhabitants									−0.001	0.002
									(0.006)	(0.006)
Net migration per 1,000 inhabitants										0.001
										(0.001)
Constant	-0.176^{***}	-0.106^{***}	-0.064^{**}	-0.135^{**}	−0.050	-0.100^{**}	-0.113^{**}	-0.102^{**}	-0.101^{**}	-0.105^{**}
	(0.010)	(0.024)	(0.027)	(0.055)	(0.046)	(0.040)	(0.046)	(0.043)	(0.045)	(0.044)
R^2	0.087	0.150	0.205	0.231	0.380	0.434	0.436	0.457	0.457	0.474
Districts	174	174	174	174	174	174	174	174	174	174

Note: District-level weighted least-squares regressions for $N = 174$ non-Budapest districts, weighted by district employment; heteroskedasticity-robust standard errors are reported in parentheses. The dependent variable is the firm-adjusted gap recovered from equation (III.4). Covariates enter cumulatively from left to right. Columns (2), (6), and (10) correspond to Models 1, 2, and 3 of Table 4. $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

Table C.2. Stepwise Build-Up of the Second Stage: Baseline Gap

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Female high-skill occupation share	0.204***	0.252***	0.330***	0.330***	0.175***	0.086	0.079	0.123**	0.124**	0.016
	(0.032)	(0.035)	(0.050)	(0.054)	(0.062)	(0.057)	(0.053)	(0.060)	(0.060)	(0.083)
Female share of registered jobseekers		−0.290**	−0.225**	−0.216	−0.229	−0.049	0.041	−0.011	−0.017	−0.061
		(0.120)	(0.111)	(0.178)	(0.160)	(0.128)	(0.127)	(0.131)	(0.135)	(0.137)
Log district employment			−0.013*	−0.012	−0.008	−0.012	−0.007	−0.006	−0.008	−0.005
			(0.007)	(0.008)	(0.007)	(0.007)	(0.008)	(0.007)	(0.008)	(0.008)
Employer concentration (HHI)				0.031	−0.077	−0.415	−0.502	−0.663*	−0.676*	−0.443
				(0.313)	(0.293)	(0.363)	(0.356)	(0.385)	(0.396)	(0.342)
Industry concentration (HHI)					−0.358***	−0.231*	−0.177	−0.139	−0.136	−0.102
					(0.097)	(0.118)	(0.121)	(0.130)	(0.131)	(0.129)
Public-sector employment share						0.375**	0.474***	0.462***	0.454***	0.564***
						(0.153)	(0.163)	(0.158)	(0.152)	(0.173)
Nursery places per 100 children aged 0–2							−0.002***	−0.002**	−0.002**	−0.001
							(0.001)	(0.001)	(0.001)	(0.001)
Sewage-connected dwelling share								−0.038*	−0.038*	−0.036*
								(0.021)	(0.021)	(0.021)
Log crimes per 1,000 inhabitants									0.006	0.013
									(0.013)	(0.014)
Net migration per 1,000 inhabitants										0.002*
										(0.001)
Constant	−0.256***	−0.140***	−0.088	−0.094	0.015	−0.073	−0.176	−0.159	−0.166	−0.176*
	(0.017)	(0.052)	(0.067)	(0.118)	(0.113)	(0.097)	(0.108)	(0.107)	(0.108)	(0.106)
R^2	0.156	0.203	0.228	0.228	0.297	0.343	0.372	0.387	0.388	0.412
Districts	174	174	174	174	174	174	174	174	174	174

Note: District-level weighted least-squares regressions for $N = 174$ non-Budapest districts, weighted by district employment; heteroskedasticity-robust standard errors are reported in parentheses. The dependent variable is the baseline gap recovered from equation (III.2). Covariates and entry order are identical to Table C.1. Column (10) corresponds to Model 4 of Table 4. $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

D. Private-Sector Re-estimation

This appendix reports the samples and distributions behind the private-sector re-estimation. Steps 1 and 2 are re-estimated on ordinary employment relationships alone, excluding the civil-service, public-employee, and armed-services categories defined in Section II.C, and the second stage is then run on the resulting district gaps. The second-stage estimates are in Model 5 of Table 4 and the moments in columns (3) and (4) of Table 3; what this appendix adds is the size of the estimation samples, the shape of the two distributions, and the district-by-district agreement with the main sample.

The private-sector Step 1 sample contains 157,238,548 worker–firm–month observations, and Step 2, after 11,787 singleton observations are dropped, 157,226,761. Relative to the main estimation samples of 172,645,146 and 172,633,277 reported in Section II.B, restricting to ordinary employment relationships removes 15,406,598 observations, or 8.9%.

That share is well below the district-average public-sector share of 0.201 reported in Table 2, and the difference is a property of the analysis panel rather than of the districts. Public-sector records are six times as likely as ordinary employment records to lack a usable industry code, 66.4% against 11.0%, and since the analysis panel differs from the wider eligible-wage sample only in requiring valid codes, the requirement alone reduces the public-sector share of the sample from about a fifth to about a twelfth (Table E.3). The dependent variable in the main specification is accordingly estimated on a panel that is already only about 9% public, which is worth keeping in view when reading Model 5 of Table 4: the private-sector re-estimation confirms the main result on a modestly smaller sample rather than transforming it.

The private-sector distributions have the same shape as their main-sample counterparts, with slightly more dispersion both before and after firm fixed effects, and the compression is 67.0% against 68.8%. Panel B of Figure D.1 shows that the agreement with the main sample is uniform across the range rather than an average: no district departs materially from the 45-degree line, and the private-sector estimates are the more dispersed of the two, lying below the line at the bottom of the range and above it at the top.

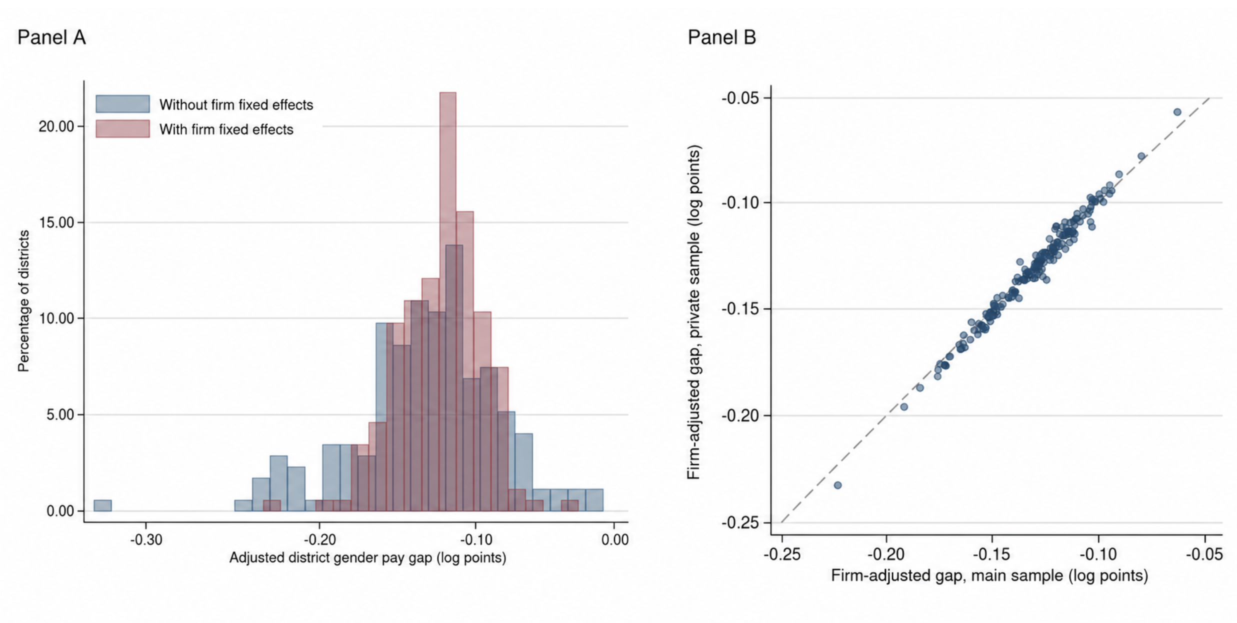

Figure D.1. Private-Sector Re-estimation

Panel A: Distribution before and after firm fixed effects, private-sector sample.
Panel B: Firm-adjusted gap, private sample against main sample.

Note: Panel A gives the unweighted distribution of the private-sector estimates across 174 non-Budapest districts, in bins of 0.01 log points; the corresponding main-sample distributions are in Figure 1. Panel B plots the private-sector firm-adjusted estimate against its main-sample counterpart, with the 45-degree line. The two sets of estimates correlate at 0.991, and the baseline estimates at 0.994.

E. Estimation Samples and Industry-Code Coverage

The industry code is the largest single restriction on the estimation sample: it is missing or invalid for 22.5% of the eligible-wage sample, and its incidence varies sharply across employment relationships and across years. This appendix reports four things: what happens to the dispersion result when the estimation samples and the industry control are varied (E.1); what the second stage looks like when observations without an industry code are restored (E.2); how coverage is distributed across legal relationships and over time (E.3); and what happens when the district coverage rate is itself entered as a control (E.4).

E.1. Dispersion under Alternative Samples and Controls

Table E.1. Dispersion of the Recovered District Gaps under Alternative Samples and Controls

Specification	**Estimation sample**	**Baseline mean**	**Baseline SD**	**Firm-adjusted mean**	**Firm-adjusted SD**	**Variance reduction**
M: Main specification	172,645,146 172,633,277	−0.1474	0.0404	−0.1309	0.0226	68.757%
M': M on the Step 2 sample	172,633,277	[·]	[·]	−0.1309	0.0226	68.759%
C: Industry control removed	172,633,277	−0.1373	0.0449	−0.1309	0.0226	74.698%
S: C, missing-code records restored	220,653,668	−0.1176	0.0443	−0.1116	0.0231	72.758%

Note: Estimation samples are worker–firm–month observations; where two figures are given, they are the Step 1 and Step 2 samples. M is the specification of Section III. M' re-estimates Step 1 on the Step 2 sample, so that the two steps are exactly rather than nearly sample-matched. C removes the individual-level industry control on the same valid-industry sample. S extends C to the observations that carry no usable industry code, adding 48,020,391 observations. Moments are unweighted across the 174 districts. Moving from C to S reduces the variance reduction by 1.939 percentage points. The correlations between the C and S district gaps are 0.971 for the baseline estimates and 0.919 for the firm-adjusted estimates.

The sample mismatch between the two steps is immaterial. Step 2 drops 11,869 singleton observations, 0.007% of the Step 1 sample; re-estimating Step 1 on that sample moves the variance reduction by 0.002 percentage points.

Removing the industry control raises the estimated compression, and does so in the baseline step alone. The firm-adjusted district gaps are virtually unchanged—the largest absolute difference across the 174 districts is 0.00013 log points, 0.58% of their standard deviation, and the two series correlate at 0.999999—while the baseline dispersion rises from 0.0404 to 0.0449. Decomposing the 5.94-point movement from 68.757% to 74.698% by holding the firm-adjusted dispersion at its main-specification value attributes 5.943 points to the baseline step and -0.002 to the firm-adjusted step. The industry indicators are not exactly collinear with firm identity, since a firm's industry code occasionally changes across years, but the estimate reported in Section IV is conservative in a precise sense: what it attributes to the control set is compression that firm fixed effects would otherwise perform.

Restoring the missing-code observations lowers the compression and shifts the level. The variance reduction falls by 1.9 percentage points and the mean baseline gap narrows by about two log points, from -0.137 to -0.118. The direction is the one anticipated in footnote 1 of Section III.A: districts and workers with less complete coverage display smaller adjusted female penalties.

E.2. The Second Stage under Restored Coverage

The two covariates other than the public-sector share that carry the main specification are stable. Industry concentration is essentially unchanged, -0.160 against -0.157, and the sewage share is slightly stronger and more precisely estimated. The public-sector coefficient falls by about a quarter, from 0.267 to 0.195, and remains positive and distinguishable from zero at the 5% level; this is the largest movement any check in Section V produces on a coefficient the paper relies on, and it is reported as such rather than folded into a statement of stability.

Two coefficients the main specification does not rely on move more. Employer concentration rises from 0.049 to 0.264 and approaches conventional significance ($p = 0.063$), and the female high-skill share changes sign, from -0.028 to 0.029, though neither is distinguishable from zero in either column. The overall fit improves markedly, from $R^2 = 0.474$ to 0.581. Both patterns are consistent with the restored observations being concentrated in districts that the covariate set already describes well; E.4 shows this is largely true.

Removing the public-sector share from the specification raises the female high-skill coefficient from -0.028 to 0.058 (0.038) under C and from 0.029 to 0.093 (0.032) under S, where it becomes distinguishable from zero, and strengthens industry concentration from -0.160 to -0.250 (0.056) under C and from -0.157 to -0.222 (0.044) under S. The suppression relationship documented for the main specification therefore survives the coverage change intact, and in the restored-coverage sample it is more visible rather than less.

Table E.2. Second Stage on the Firm-Adjusted District Gaps, Spec C against Spec S

	(1) Spec C	**(2) Spec S**
Female high-skill share	−0.028	0.029
	(0.042)	(0.036)
Female jobseeker share	0.022	−0.017
	(0.071)	(0.064)
Log employment size	−0.004	−0.005
	(0.003)	(0.003)
Employer concentration (HHI)	0.049	0.264
	(0.155)	(0.141)
Industry concentration (HHI)	**−0.160**	**−0.157**
	(0.067)	(0.056)
Public-sector employment share	**0.267**	**0.195**
	(0.090)	(0.077)
Nursery places per 100 children aged 0–2	0.000	0.000
	(0.000)	(0.000)
Sewage-connected dwelling share	**−0.023**	**−0.027**
	(0.009)	(0.009)
Log recorded crime	0.002	0.000
	(0.007)	(0.006)
Net migration rate	0.001	0.000
	(0.001)	(0.001)
Constant	**−0.105**	−0.071
	(0.044)	(0.040)
R^2	0.474	0.581
F statistic	19.28	37.31
Districts	174	174

Note: Employment-weighted least-squares regressions on 174 districts; heteroskedasticity-robust standard errors are reported in parentheses. Bold values mark coefficients distinguishable from zero at the 5% level. Column (1) reproduces the second stage of the main specification: because the firm-adjusted district gaps under M and C differ by at most 0.00013 log points, as shown in Table E.1, the two second stages are numerically the same. The public-sector coefficient of 0.267 in column (1) is the figure reported in Table 4. On the baseline gaps, the same two columns give a public-sector coefficient of 0.666 (0.195) under C and 0.647 (0.190) under S, with R^2 values of 0.440 and 0.465.

E.3. Coverage by Legal Relationship and by Year

Public-sector records are six times as likely as ordinary employment records to lack a usable industry code: 66.4% against 11.0%. Because the cleaned analysis panel differs from the eligible-wage sample only in requiring valid occupation and industry codes (Appendix A.3), and the occupation code is almost universally present, the industry requirement alone reduces the public-sector share of the sample from about a fifth to about a twelfth. The district gaps estimated in Section III are therefore recovered from a panel that is roughly 91% private-sector employment before any private-sector restriction is imposed, which is the context in which Model 5 of Table 4 should be read.

Table E.3. Industry-Code Coverage by Employment Legal Relationship

Legal relationship (*fogvisz1*)	No usable industry code	Usable code	Total
Ordinary employment (201)	19,516,014 (11.02%)	157,620,949 (88.98%)	177,136,963
Civil servants and related officials (202)	3,708,035 (44.10%)	4,699,717 (55.90%)	8,407,752
Public employees (203)	25,060,464 (73.86%)	8,870,037 (26.14%)	33,930,501
Armed services (204)	1,866,680 (48.97%)	1,945,006 (51.03%)	3,811,686
All public (202–204)	**30,635,179 (66.38%)**	15,514,760 (33.62%)	46,149,939
Total	50,151,193 (22.46%)	173,135,709 (77.54%)	223,286,902

Note: Eligible-wage sample, $N = 223{,}286{,}902$; percentages are calculated within each row. A record enters the Step 1 estimation sample only with a usable industry code. Both a missing code and an invalid code disqualify it, and both are counted here. The district-average figure of 0.2008 on the eligible-wage sample corresponds to the 0.201 reported in Table 2. The occupation code is present for 99.96% of the eligible-wage sample and for over 99.99% of every public category, so the shortfall is specific to the industry code. The public-sector share of the eligible-wage sample is 20.67% pooled and 0.2008 as a district average; after the code requirements, it is 8.96% pooled and 0.0755 as a district average.

Memo: Coverage by period: share of worker-months with a usable industry code (%).

Period	Public (202–204)	Ordinary (201)	All
2003–2011	26.01	88.82	75.91
2012–2016	53.51	90.02	82.39
2017	0.01	85.42	67.84

Public-sector coverage improves substantially in the second half of the panel, from about a quarter to about a half, and then fails almost completely in the final year: in 2017, 3,268,626 of 3,269,031 public-sector worker-months carry no usable industry code. Ordinary employment shows no comparable break—coverage falls by 4.6 percentage points between 2012–2016 and 2017, to 85.4%, which is within the range observed earlier in the panel and equal to the 2003 figure. The final-year failure is therefore specific to public-sector records. It leaves the 2017 wave contributing almost no public-sector observations to Step 1, but the district industry Herfindahl, which rests mainly on private-sector composition, is unaffected in any material way, and no re-estimation is implied. This makes concrete the statement in footnote 1 of Section III.A that missingness is highest in the earliest and the final years of the panel.

E.4. *Conditioning on the District Coverage Rate*

Because coverage is so unevenly distributed across legal relationships, the district industry-missingness rate and the district public-sector share are close to the same variable. They correlate at 0.705, and regressing the missingness rate on the four covariates that carry the second stage returns a public-sector coefficient of 79.6 (7.1) with $R^2 = 0.706$: a district ten percentage points more public has an industry-missingness rate about eight percentage points higher.

Entering the missingness rate directly into the Spec S specification therefore does not isolate a coverage channel so much as split a single underlying quantity. It has a large effect. The public-sector coefficient falls from 0.195 (0.077) to 0.050 (0.065) and is no longer distinguishable from zero ($p = 0.44$), while the missingness rate itself enters at 0.0028 (0.0008), $p < 0.001$, and the fit rises to $R^2 = 0.632$. Industry concentration weakens from -0.157 to -0.116 (0.057) but remains significant; the sewage share weakens from -0.027 to -0.018 (0.008) and also remains significant.

Table E.3 shows that industry-code missingness is substantially more prevalent among public-sector records: public records lack an industry code six times as often as ordinary records, and the district missingness rate is largely a mechanical consequence of the district's public-sector composition. The missingness rate is accordingly a post-treatment variable with respect to public employment, and conditioning on it removes most of the variation the public-sector coefficient is estimated from. The specification cannot be used to argue that the public-sector result is an artefact of coverage; equally, it cannot be used to rule that possibility out, because the two are not separately identified in this design.

What can be said is narrower and should be said in exactly that form: the public-sector coefficient is the one second-stage result whose interpretation is sensitive to how the missing-code observations are handled, and it should be reported with that qualification rather than as a settled magnitude. The results the paper leans on most heavily—the dispersion compression of Section IV.A and the industry-concentration coefficient—are unaffected by every variant in this appendix.

F. The Definition of the Public-Sector Share

The public-sector share is the strongest covariate in the second stage and, after Appendix E, the one whose magnitude the robustness programme qualifies. It is built from the employment legal-relationship code rather than from firm ownership, which Admin3 does not record. Two features of that construction warrant examination.

The first is testable. The share counts armed services (code 204) as public alongside civil servants (202) and public employees (203). Armed-services employment is the component of the public sector least well suited to a residence-based geography, since garrison and training locations need not coincide with the districts in which service personnel are registered as resident, and Section II.C notes that its recorded share shifts to a higher level from 2012 onwards. This appendix re-estimates the second stage with the public sector defined as codes 202–203 only, retaining code-204 records in the denominator and in every wage sample.

The second is not testable within these data and is recorded here as a limitation. Employees of state-owned enterprises work under ordinary private-law employment contracts (code 201) and are therefore counted as private throughout. The measure is accordingly a measure of *public-sector employment relationships*, not of public ownership, and it understates the public presence in districts whose large employers are state-owned. Section VI returns to what this implies for the interpretation of the coefficient.

F.1. The Composition of the Public Categories

Table F.1. Employment Legal Relationships in the Eligible-Wage Sample

$N = 223{,}286{,}902$ worker-months, 2003–2017, districts outside Budapest.

Legal relationship (*fogvisz1*)	Records	% of sample	% of public	Industry code present
201 Ordinary employment	177,136,963	79.33	—	88.98%
202 Civil servants and related officials	8,407,752	3.77	18.22	55.90%
203 Public employees	33,930,501	15.20	73.52	26.14%
204 Armed services	3,811,686	1.71	8.26	51.03%
All public (202–204)	**46,149,939**	**20.67**	**100.00**	**33.62%**

Note: Counts are from the eligible-wage sample of Appendix A.3, on which the district public-sector share is computed. Codes above 204—public work (*közfoglalkoztatás*, 215), self-employment, cooperative membership, and the other non-standard relationships—are outside the sample altogether, as described in Appendix A.1.

The definitional question concerns a small component. Armed services are 1.7% of the eligible-wage sample and 8.3% of public-sector records; public employees (203) alone account for nearly three quarters of the public category. If the second-stage result were sensitive to the treatment of code 204, it would be sensitive to a component too small to carry it, which is why the check is reported rather than assumed.

F.2. How Code 204 Enters the Data

Two features of the code-204 records bear on whether excluding them should be expected to matter, and both cut against the concern that motivates the check.

Table F.2. Prevalence of Code 204 by Year

Year	Share of eligible worker-months (%)	Year	Share (%)
2003	0.93	2011	0.94
2004	1.00	2012	2.72
2005	0.95	2013	2.72
2006	0.87	2014	2.66
2007	0.78	2015	2.60
2008	0.69	2016	2.59
2009	0.75	2017	2.66
2010	2.45	**Pooled**	**1.71**

Note: Armed-services records average 0.86% of eligible worker-months from 2003 to 2009 and in 2011, and 2.66% from 2012 onwards; 2010 is an isolated year at the later level. The pooled share over 2003–2017 is 1.71%, and the 2013–2017 share is 2.65%.

Table F.3. District Distribution of Code 204 and of the Two Public-Sector Shares

Across the 174 districts.

	Mean	SD	Min	p25	Median	p75	Max
Code-204 share	0.0171	0.0075	0.0061	0.0118	0.0154	0.0207	0.0433
Public share, 202–204	0.2008	0.0412	0.1264	—	—	—	0.3481
Public share, 202–203	0.1837	0.0374	0.1185	—	—	—	0.3261

Note: The two public-sector definitions correlate at 0.987; the code-204 share correlates with the 202–204 share at 0.574.

The coverage of code 204 changes level once, and the change is not confined to the early years. The share is at the higher level in 2010, returns to the earlier level in 2011, and then shifts permanently from 2012 onwards. The series displays a level-shift pattern consistent with a recording or classification change rather than a gradual trend. Because the second-stage covariate averages district-year means with equal weight across years, it affects every district in the same direction.

Armed-services records are not geographically concentrated. All 174 districts have a positive code-204 share, the interquartile range is 0.012 to 0.021, and the maximum is only 2.5 times the mean. This is the pattern a residence-based identifier produces: service personnel are recorded in the districts where they live, not where they are garrisoned, so the misallocation the check was designed to detect does not arise in a concentrated form. It is also why the two definitions correlate at 0.987—subtracting code 204 removes a component that is close to a common proportion of each district's public share rather than a few districts' worth of it.

F.3. The Second Stage under the Alternative Definition

Table F.4. Public-Sector Coefficient under Both Definitions, Four Outcomes

Employment-weighted least squares on 174 districts; robust standard errors in parentheses.

Dependent variable	(1) Public = 202–204	(2) Public = 202–203	Raw change	Per-SD change
Firm-adjusted gap (main)	0.267^{***} (0.090)	0.285^{***} (0.096)	+6.6%	−3.2%
Firm-adjusted gap, private sample	0.324^{***} (0.095)	0.352^{***} (0.102)	+8.8%	−1.2%
Baseline gap	0.564^{***} (0.173)	0.556^{***} (0.188)	−1.3%	−10.4%
Baseline gap, private sample	0.525^{***} (0.185)	0.506^{**} (0.201)	−3.7%	−12.6%

Note: The nine remaining covariates and the weight are those of Table 4. Column (2) replaces the public-sector share with the same quantity computed over codes 202–203, leaving code-204 records in the denominator and in every wage and outcome sample; the alternative covariate is therefore a definitional change to one regressor and nothing else. The per-SD column compares the effect of a one-standard-deviation move in each share, using the standard deviations in Table F.3. The rebuilt 202–204 share reproduces the stored covariate to ten decimal places, so column (1) is Table 4 exactly, and the identity linking the two shares to the code-204 share holds to 1.5×10^{-8}. The corresponding R^2 values are: firm-adjusted, $0.474 \to 0.467$; firm-adjusted private, $0.474 \to 0.469$; baseline, $0.412 \to 0.393$; and baseline private, $0.382 \to 0.364$. $^{***}p < 0.01$; $^{**}p < 0.05$.

The firm-adjusted public-sector association is not materially sensitive to excluding armed-services records. The baseline-gap estimates show somewhat greater sensitivity in per-standard-deviation terms. Nevertheless, all four coefficients retain their signs and remain statistically significant under the alternative definition: three are significant at the 1% level, while the private-sample baseline estimate is significant at the 5% level. Model fit is marginally lower under the alternative definition in all four cases.

The two columns should not be compared coefficient by coefficient. The 202–203 share is smaller in every district—its mean is 91.5% of the 202–204 mean and its standard deviation 90.8% of it—so the same underlying association must attach a larger coefficient to the narrower measure. The rise from 0.267 to 0.285 in the main specification is the arithmetic of that rescaling and is not evidence that including armed services attenuated the estimate. In per-SD terms, the two are within 3.2% of each other, which is a smaller movement than any other check in the robustness programme produces on this coefficient. The private-sample column behaves the same way, at 1.2%.

The baseline-gap rows move more in per-SD terms, by about 11%, and in the opposite direction from the firm-adjusted rows. This is consistent with what Section IV.A establishes about the two series: the baseline gaps carry the firm-composition variation that firm fixed effects remove, and that variation is more sensitive to which employment relationships are counted as public. The paper's argument rests on the firm-adjusted rows, where the sensitivity is smallest.

G. Second-Stage Robustness

This appendix reports the second stage under alternative inference procedures, under an earlier measurement window for the district characteristics, on samples that omit the most influential districts, and with the public-sector share excluded from the covariate set. Table G.1 covers the first two; Table G.2 the last two. All specifications use the covariates and weights of Model 3 in Table 4 and differ only as the column or row headings state.

Table G.1. Alternative Inference and Start-of-Period Covariates

	(1) Robust SE	(2) County-clustered	(3) Wild bootstrap p-value	(4) 2003–2007 covariates
Female high-skill occupation share	−0.028	−0.028	—	−0.040
Female share of registered jobseekers	0.021	0.021	—	−0.007
Log district employment	−0.004	−0.004	0.346	−0.004
Employer concentration (HHI)	0.049	0.049	—	0.016
Industry concentration (HHI)	-0.160^{**}	-0.160^{**}	0.066	-0.114^{**}
Public-sector employment share	0.267^{***}	0.267^{**}	0.005	0.288^{***}
Nursery places per 100 children aged 0–2	0.000	0.000	—	0.000
Sewage-connected dwelling share	-0.023^{**}	-0.023^{**}	0.054	-0.017^{**}
Log crimes per 1,000 inhabitants	0.002	0.002	—	0.007
Net migration per 1,000 inhabitants	0.001	0.001	0.215	0.001^{**}
R^2	0.474	0.474	—	0.503
Districts	174	174	174	174

Note: The dependent variable is the firm-adjusted gap; all columns are weighted by district employment. Column (1) repeats Model 3 of Table 4. Column (2) clusters standard errors on the nineteen counties, which leaves the point estimates unchanged. Column (3) reports wild-cluster bootstrap p-values (9,999 replications, Rademacher weights, null imposed), computed for the coefficients distinguishable from zero under robust standard errors and, for reference, for log district employment and net migration. Column (4) measures all covariates over 2003–2007 rather than the full period; because the T-STAR employment measure is unavailable for 2004–2006, district size in this column averages 2003 and 2007 only, and employer concentration is held at its full-period value. $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

Table G.2. Influential Districts and the Public-Sector Share

Sample or specification	Public-sector share	Industry HHI	Employer HHI	Districts	R^2
Panel A. Firm-adjusted gap					
Full sample	0.267^{***} (0.090)	-0.160^{**} (0.067)	0.049 (0.155)	174	0.474
Omitting Paksi	0.186^{***} (0.053)	-0.220^{***} (0.042)	0.132 (0.141)	173	0.519
Omitting Paksi, Téti, and Szentgotthárdi	0.188^{***} (0.054)	-0.220^{***} (0.043)	0.131 (0.142)	171	0.518
Omitting the five largest Cook's distances	0.200^{***} (0.055)	-0.203^{***} (0.043)	0.157 (0.131)	169	0.525
Public-sector share excluded	—	-0.250^{***} (0.056)	0.167 (0.143)	174	0.409
Public-sector share excluded, omitting Paksi	—	-0.287^{***} (0.044)	0.220 (0.140)	173	0.485
Panel B. Alternative dependent variables					
Baseline gap, full sample	0.564^{***} (0.173)	-0.102 (0.129)	-0.443 (0.342)	174	0.412
Baseline gap, omitting Paksi	0.422^{***} (0.123)	-0.206^{**} (0.094)	-0.298 (0.338)	173	0.420
Private-sector gap, full sample	0.324^{***} (0.095)	-0.173^{**} (0.071)	0.090 (0.168)	174	0.474
Private-sector gap, omitting Paksi	0.237^{***} (0.056)	-0.237^{***} (0.045)	0.179 (0.153)	173	0.515
Memo: leave-one-out across all 174 districts, firm-adjusted gap					
Median	0.267				
Standard deviation	0.008				
Minimum / maximum	0.186 / 0.290				
Change from the full sample, omitting any district other than Paksi	-4.2% to $+6.6\%$				

Note: Each row re-estimates the specification of Table 4 on the stated sample, weighted by district employment, with heteroskedasticity-robust standard errors in parentheses; only the three coefficients discussed in the text are shown. The five districts with the largest Cook's distances are Paksi, Miskolci, Szigetszentmiklósi, Nagykátai, and Dunaújvárosi; of these, only Paksi and Dunaújvárosi have outlying recovered gaps, the other three being influential through leverage rather than through the dependent variable. The last two rows of Panel A drop the public-sector share from the covariate set and are otherwise identical to the first and second. The memo block summarises 174 regressions, each omitting one district. $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

References

John M. Abowd, Francis Kramarz, and David N. Margolis. High wage workers and high wage firms. *Econometrica*, 67(2):251–333, 1999. doi: 10.1111/1468-0262.00020.

Fredrik Andersson, Simon Burgess, and Julia I. Lane. Cities, matching and the productivity gains of agglomeration. *Journal of Urban Economics*, 61(1):112–128, 2007. doi: 10.1016/j.jue.2006.06.005.

José Azar, Ioana Marinescu, and Marshall Steinbaum. Labor market concentration. *Journal of Human Resources*, 57(S):S167–S199, 2022. doi: 10.3368/jhr.monopsony.1218-9914r1.

Marianne Bertrand, Claudia Goldin, and Lawrence F Katz. Dynamics of the gender gap for young professionals in the financial and corporate sectors. *American Economic Journal: Applied Economics*, 2(3):228–255, 2010. doi: 10.1257/app.2.3.228.

David Blackaby, Philip Murphy, Nigel O'Leary, and Anita Staneva. Regional pay? the public/private sector pay differential. *Regional Studies*, 52(4):477–489, 2018. doi: 10.1080/00343404.2017.1331295.

Francine D. Blau and Lawrence M. Kahn. The gender wage gap: Extent, trends, and explanations. *Journal of Economic Literature*, 55(3):789–865, 2017. doi: 10.1257/jel.20160995.

Benjamin Bruns. Changes in workplace heterogeneity and how they widen the gender wage gap. *American Economic Journal: Applied Economics*, 11(2):74–113, 2019. doi: 10.1257/app.20160664.

Gábor Békés. Location of manufacturing fdi in hungary: How important are business-to-business relationships? Technical Report June 2004, Central European University, 2004.

David Card, Jörg Heining, and Patrick Kline. Workplace heterogeneity and the rise of west german wage inequality. *The Quarterly Journal of Economics*, 128(3):967–1015, 2013. doi: 10.1093/qje/qjt006.

David Card, Ana Rute Cardoso, and Patrick Kline. Bargaining, sorting, and the gender wage gap: Quantifying the impact of firms on the relative pay of women. *The Quarterly Journal of Economics*, 131(2):633–686, 2016. doi: 10.1093/qje/qjv038.

David Card, Jesse Rothstein, and Moises Yi. Location, location, location. *American Economic Journal: Applied Economics*, 17(1):297–336, 2025. doi: 10.1257/app.20220427.

Youngjoo Cha and Kim A. Weeden. Overwork and the slow convergence in the gender gap in wages. *American Sociological Review*, 79(3):457–484, 2014. doi: 10.1177/0003122414528936.

Pierre-Philippe Combes, Gilles Duranton, and Laurent Gobillon. Spatial wage disparities: Sorting matters! *Journal of Urban Economics*, 63(2):723–742, 2008. doi: 10.1016/j.jue.2007.04.004.

Jorge de la Roca and Diego Puga. Learning by working in big cities. *The Review of Economic Studies*, 84(1): 106–142, 2017. doi: 10.1093/restud/rdw031.

Kim V.L. England. Suburban pink collar ghettos: The spatial entrapment of women? *Annals of the Association of American Geographers*, 83(2):225–242, 1993. doi: 10.1111/j.1467-8306.1993.tb01933.x.

Michaela Fuchs, Anja Rossen, Antje Weyh, and Gabriele Wydra-Somaggio. Where do women earn more than men? explaining regional differences in the gender pay gap. *Journal of Regional Science*, 61(5): 1065–1086, 2021. doi: 10.1111/jors.12532.

Edward L. Glaeser and David C. Maré. Cities and skills. *Journal of Labor Economics*, 19(2):316–342, 2001. doi: 10.1086/319563.

Edward L. Glaeser, Matt Resseger, and Kristina Tobio. Inequality in cities. *Journal of Regional Science*, 49 (4):617–646, 2009. doi: 10.1111/j.1467-9787.2009.00627.x.

Claudia Goldin. A grand gender convergence: Its last chapter. *American Economic Review*, 104(4):1091–1119, 2014. doi: 10.1257/aer.104.4.1091.

Boris Hirsch, Thorsten Schank, and Claus Schnabel. Differences in labor supply to monopsonistic firms and the gender pay gap: An empirical analysis using linked employer-employee data from germany. *Journal of Labor Economics*, 28(2):291–330, 2010. doi: 10.1086/651208.

Boris Hirsch, Marion König, and Joachim Möller. Is there a gap in the gap? regional differences in the gender pay gap. *Scottish Journal of Political Economy*, 60(4):412–439, 2013. doi: 10.1111/sjpe.12017.

Gábor Kertesi and János Köllő. Fighting "low equilibria" by doubling the minimum wage? hungary's experiment. IZA Discussion Paper 970, IZA Institute of Labor Economics, 2003.

Eva Kiss. Foreign direct investment in hungary: Industry and its spatial effects. *Eastern European Economics*, 45(1):6–28, 2007. doi: 10.2753/eee0012-8775450101.

Henrik Kleven, Camille Landais, and Jakob Egholt Søgaard. Children and gender inequality: Evidence from denmark. *American Economic Journal: Applied Economics*, 11(4):181–209, 2019. doi: 10.1257/app.20180010.

Thomas Le Barbanchon, Roland Rathelot, and Alexandra Roulet. Gender differences in job search: Trading off commute against wage. *The Quarterly Journal of Economics*, 136(1):381–426, 2021. doi: 10.1093/qje/qjaa033.

Anna Lovász and Ágnes Szabó-Morvai. Childcare availability and maternal labor supply in a setting of high potential impact. *Empirical Economics*, 56(6):2127–2165, 2019. doi: 10.1007/s00181-018-1423-x.

Anna Lovász. Competition and the gender wage gap: New evidence from linked employer-employee data in hungary, 1986-2003. Technical Report Budapest Working Papers on the Labour Market BWP 2008/4, Institute of Economics, Hungarian Academy of Sciences, 2008.

Janice Fanning Madden. Why women work closer to home. *Urban Studies*, 18(2):181–194, 1981. doi: 10.1080/00420988120080341.

Aleksandra Majchrowska and Paweł Strawiński. Regional differences in gender wage gaps in poland: New estimates based on harmonized data for wages. *Central European Journal of Economic Modelling and Econometrics*, 8(2):115–141, 2016.

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

Jaan Masso, Jaanika Meriküll, and Priit Vahter. The role of firms in the gender wage gap. *Journal of Comparative Economics*, 50(2):454–473, 2022. doi: 10.1016/j.jce.2021.10.001.

Jacob Mincer and Solomon Polachek. Family investments in human capital: Earnings of women. *Journal of Political Economy*, 82(2, Part 2):S76–S108, 1974. doi: 10.1086/260293.

Ines P. Murillo Huertas, Raul Ramos, and Hipolito Simon. Regional differences in the gender wage gap in spain. *Social Indicators Research*, 134(3):981–1008, 2017. doi: 10.1007/s11205-016-1461-8.

Martina Mysíková. Gender wage gap in the czech republic and central european countries. *Prague Economic Papers*, 21(3):328–346, 2012. doi: 10.18267/j.pep.427.

Natascha Nisic. Smaller differences in bigger cities? assessing the regional dimension of the gender wage gap. *European Sociological Review*, 33(2):292–304, 2017. doi: 10.1093/esr/jcx037.

Claudia Olivetti and Barbara Petrongolo. The evolution of gender gaps in industrialized countries. *Annual Review of Economics*, 8(1):405–434, 2016. doi: 10.1146/annurev-economics-080614-115329.

Andrew M. Penner, Trond Petersen, Are Skeie Hermansen, Anthony Rainey, István Boza, Marta M. Elvira, Olivier Godechot, Martin Hällsten, Lasse Folke Henriksen, Feng Hou, Aleksandra Kanjuo Mrčela, Joe King, Naomi Kodama, Tali Kristal, Alena Křížková, Zoltán Lippényi, Silvia Maja Melzer, Eunmi Mun, Paula Apascaritei, Dustin Avent-Holt, Nina Bandelj, Gergely Hajdu, Jiwook Jung, Andreja Poje, Halil Sabanci, Mirna Safi, Matthew Soener, Donald Tomaskovic-Devey, and Zaibu Tufail. Within-job gender pay inequality in 15 countries. *Nature Human Behaviour*, 7(2):184–189, 2023. doi: 10.1038/s41562-022-01470-z.

Barbara Petrongolo and Maddalena Ronchi. Gender gaps and the structure of local labor markets. *Labour Economics*, 64:101819, 2020. doi: 10.1016/j.labeco.2020.101819.

Rita Pető. Foreign-owned firms and the gender wage gap: Does cultural transmission matter? Technical Report KRTK-KTI Working Paper 2025/9, Institute of Economics, Centre for Economic and Regional Studies, 2025.

Euan Phimister. Urban effects on participation and wages: Are there gender differences? *Journal of Urban Economics*, 58(3):513–536, 2005. doi: 10.1016/j.jue.2005.08.006.

Melanie A. Rapino and Thomas J. Cooke. Commuting, gender roles, and entrapment: A national study utilizing spatial fixed effects and control groups. *The Professional Geographer*, 63(2):277–294, 2011. doi: 10.1080/00330124.2010.547790.

Jørn Rattsø and Hildegunn E. Stokke. Private-public wage gap and return to experience: Role of geography, gender and education. *Regional Science and Urban Economics*, 84:103571, 2020. doi: 10.1016/j.regsciurbeco.2020.103571.

Olav Sorenson and Michael S. Dahl. Geography, joint choices, and the reproduction of gender inequality. *American Sociological Review*, 81(5):900–920, 2016. doi: 10.1177/0003122416656360.

Álmos Telegdy. Public wage spillovers: The role of individual characteristics and employer wage policies. *Labour Economics*, 55:116–129, 2018. doi: 10.1016/j.labeco.2018.08.008.

Christopher H. Wheeler. Search, sorting, and urban agglomeration. *Journal of Labor Economics*, 19(4):879–899, 2001. doi: 10.1086/322823.