Market Definition Bias in Studies of (Labor) Market Power

Jamie Fogel, Bernardo Modenesi, and Benjamin Scuderi

Preliminary and incomplete

March 2026

Abstract

This paper demonstrates two distinct and quantitatively important biases introduced by using an "incorrect" definition of market boundaries when attempting to make inferences about labor market power. The first source of bias, long recognized in the antitrust literature, stems from mismeasurement of relative firm size: the same firm will appear artificially dominant when markets are drawn too narrowly and artificially competitive when they are drawn too broadly. We derive a novel second source of bias, which we term *elasticity bias*, that generates statistical attenuation of estimates of key parameters that govern model-based conclusions about the size and distribution of markdowns across employers and markets. In simulations calibrated to Brazilian administrative data, we show that the second channel is an order of magnitude more important than the first. Further, we show that market definition bias can be large in empirically-relevant cases where the relative rate of misclassification may be modest, as with administrative labor–market boundaries such as industry/occupation–region cells adopted by virtually all existing studies. Relative to using administrative market definitions, using a data-driven, network-based market definition yields an approximately 48% smaller between-market elasticity, larger markdown dispersion, and overturns several qualitative conclusions about which workers are harmed by monopsony power. Finally, we propose a simple diagnostic that allows practitioners to pick among off-the-shelf classifications when using a data-driven one is infeasible.

1 Introduction

Large and widening wage gaps across firms have renewed interest in monopsony power in labor markets. A growing empirical literature measures this power by fitting sophisticated models of workers' labor supply to individual firms with administrative data that reveal firm and market structure (e.g., Manning, 2005; Berger et al., 2022; Felix, 2022). A core input into every such exercise is a market definition: researchers must partition jobs into the relevant sets of competitors that share a pool of workers. Markets are typically defined using "administrative" definitions — industry × region or occupation × region cells — however these definitions may not align with the actual sets of jobs which compete for the same workers. In this paper we show that these standard market definitions generate two distinct forms of misclassification bias. The first source of bias, long recognized in the antitrust literature as the *denominator problem*, stems from mismeasurement of relative firm size: the same firm will appear artificially dominant when markets are drawn too narrowly and artificially competitive when they are drawn too broadly. Moreover, even holding fixed the number of markets, misclassification may understate concentration in some markets and overstate it in others. We derive a novel second source of bias, which we term *elasticity bias*, that generates statistical attenuation of estimates of key parameters that govern model-based conclusions about the size and distribution of markdowns across employers and markets. These two sources of bias interact with each other and can lead to severely biased estimates of labor market power. We then build upon Fogel and Modenesi (2022) to provide a data-driven market definition that reduces both forms of bias and yields more accurate estimates of labor market power using Brazilian administrative linked employer-employee data.

To illustrate the practical implications of market definition bias, we first adapt the model of Berger et al. (2022) in which firms engage in Cournot-style competition and workers supply labor following a nested logit process. The nested logit preference structure gives rise to a partitioning of firms into specific markets, such that firms compete more strongly for workers with other firms in the same market, and less strongly with firms in other markets. The strength of competition is governed by two structural parameters: a within-market elasticity parameter and a between-market elasticity parameter. If a firm is small (in terms of employment or payroll) relative to its market, more of its competition for workers will come from other firms in the same market, and thus the firm-specific elasticity it faces will be closer to the within-market structural elasticity parameter. By contrast, if a firm is large relative to its market, the firm-specific labor supply curve it faces will depend upon substitution from firms in other markets, and thus its labor supply elasticity will be closer to the between-market structural elasticity. More formally, the labor supply elasticity faced by

an individual firm is a weighted average of the within-market and between-market elasticities, with the weight given by the firm's payroll-weighted employment share in its own market. In general, the within-market elasticity is larger than the between-market elasticity, as firms are more substitutable with, and thus face stronger competition from, firms in their own market.

Given this model and a set of market definitions, the payroll-weighted employment shares are observable, however the between-market and within-market elasticities must be estimated. Our estimation strategy follows Felix (2022) and consists of a two-step estimation strategy. The first step estimates the within-market elasticity as the coefficient from a regression of firm-level employment on firm-level wages, controlling for market fixed effects. The second step recovers the between-market elasticity from the coefficient from a regression of the first-step market fixed effect estimates on a market level nested logit "inclusive value" wage index, computed using the within-market elasticity from the first step. Elasticities are measured by wage sensitivity to labor-demand shocks: larger within-market wage movements indicate a higher within-market elasticity, while larger between-market wage movements indicate a higher between-market elasticity.

We then turn to the problem of characterizing the bias created by estimating these elasticities when firms are misclassified into the wrong markets. In general, misclassification causes the estimates of the (larger) within-market elasticity to be downward biased and the (smaller) between-market elasticity to be upward biased. In the limit, if jobs are classified into markets completely at random, the two elasticities' estimates converge to the same value. Intuitively, if firms are classified to markets at random, then there ceases to be any difference in the degree to which firms in the same market and firms in different markets exert competitive pressure.

We derive expressions for the bias in each elasticity parameter that depend on the degree of misclassification, where the degree of misclassification is measured as the extent to which the misclassified labor market partition distorts the true structure of wage dispersion. The bias is larger when between–market wage differentials are large, because assigning a job to the wrong market then flips high- and low-wage rankings and maximally distorts the wage gaps that identify the elasticities. In the special case when the two market definitions are nested, if the misclassified markets are coarser than the true markets then the within-market elasticity will be biased downwards towards the between-market elasticity. Conversely, if the misclassified markets are more granular than the true markets then the between-market elasticity will be biased upwards towards the within-market elasticity

In subsequent work, we will incorporate ex ante worker skill/productivity heterogeneity, which will allow us to incorporate the notion that some pairs of markets are more substitutable for each other and therefore exert more competitive pressure on each other. Intuitively, markets

are revealed to be more substitutable if they tend to hire from similar groups of workers. In general, ignoring the fact that some markets are closer to each other than others will cause researchers to overstate market power among jobs that belong to markets with close competitors.

We conduct a series of simulations to assess empirical magnitudes of misclassification bias in an idealized setting in which we know the correct market definitions and face no endogeneity concerns when estimating labor supply elasticity parameters. We calibrate labor supply elasticity parameters, labor market size distribution, firm size distributions, and wage distributions to match those found in the Brazilian administrative linked employer-employee data set RAIS during the 2009-2011 period. We then simulate labor demand shocks in order to obtain exogenous variation in wages that will allow us to recover unbiased estimates of labor supply elasticities following the identification strategy outlined above. We first estimate the within-market and between-market elasticities using the correct market definitions we used to generate the data. We then randomly misclassify an increasingly large fraction of jobs to the wrong markets and re-estimate the elasticities. We show that the simulations exactly match our theoretical predictions: as the misclassification rate increases, the two elasticities converge towards each other. Moreover, when more of the wage variance is between-market than within market, meaning that markets are more distinct, the degree of bias increases more quickly with the rate of misclassification. This is consistent with the notion that as markets grow more distinct, the consequences of misclassification rise. Given the theoretical importance of misclassification bias, we next turn to empirical applications that gauge its importance empirically.

A key challenge is that we do not know the true market definitions in typical real world settings.[1] Therefore, we compare commonly used "administrative" labor market definitions based on the intersection of occupation and geography with the "network-based" labor market definitions proposed by Fogel and Modenesi (2022). Fogel and Modenesi (2022) assign jobs to markets and workers to skill types according to a revealed preference argument by exploiting observed worker-job matching patterns. Intuitively, jobs that tend to hire similar workers are revealed to belong to similar markets and workers that tend to be hired by similar jobs are revealed to belong to similar worker skill types. While this draft focuses on the assignment of jobs to markets, in subsequent work we will use these revealed worker types to characterize worker heterogeneity. Fogel and Modenesi (2022) show that network-based labor markets outperform administrative labor markets in a variety of ways. For example, when they cluster workers into skill types, workers in the same skill type supply labor to a more concentrated set

[1] Indeed, a division of firms into discrete labor markets within a nested logit labor supply model is at best an approximation of reality.

of network-based labor markets than administrative labor markets. In this paper, we extend their market definition method to encompass monopsonistic competition and we estimate the network-based labor markets using Brazilian administrative data.

We estimate labor market elasticities — and by extension wage markdowns — by exploiting the early-1990s trade liberalization shock in Brazil following Felix (2022). Specifically, we follow the two-step identification strategy outlined above, but instrument for firm-level and market-level wages using firm-level and market-level trade shock exposure measures. We replicate Felix (2022)'s estimates of the two elasticity parameters, where labor markets are defined as the intersection of two-digit occupation and micro region.[2] We then produce estimates using the network-based labor markets. Using the network-based markets, we estimate a slightly larger within-market elasticity, and an approximately 48% smaller between-market elasticity. These results are consistent with significant misclassification in administrative labor markets relative to network-based labor markets.

Finally, we compute Herfindahl-Hirschman Indexes (HHIs) of market concentration and markdowns for each market, according to both the network-based and administrative market definitions. We find that more workers are employed in more heavily concentrated markets when we use network-based labor markets as opposed to administrative labor markets. We then combine the elasticity estimates from the previous step with the HHIs to compute the model-implied distribution of wage markdowns according to each market definition. We find significantly greater dispersion in markdowns, and a larger set of workers subject to large markdowns, when we employ network-based labor markets. This arises from three factors. First, according to the nested logit labor supply model with Cournot competition, the between-market labor supply elasticity serves as a lower bound on the labor supply elasticity faced by an individual firm (and thus an upper bound on the markdown), corresponding to the case when the firm is the only firm in its labor market. The network-based labor markets produce a significantly smaller between-market elasticity estimate, and thus a larger upper bound on the implied markdown. Second, by a symmetric argument, the within-market elasticity serves as an upper bound on firms' labor supply elasticities (and a lower bound on markdowns). Network-based labor markets yield a slightly larger within-market elasticity estimate and thus, when combined with the smaller between-market elasticity, a strictly wider range of possible markdowns. Third, recall that if a firm has greater market share, that means that the firm puts more weight on the between-market labor supply elasticity relative to the within-market elasticity. The fact that we generally find greater firm market shares (larger HHIs) using network-based definitions combined with the smaller between-market

[2]There are 558 micro regions in Brazil. They are approximately analogous to Commuting Zones in the U.S.

elasticity estimates is what explains the greater prevalence of large markdowns when using network-based market definitions.

Literature. Our paper is closely related to three distinct literatures. First, we follow a recent but growing literature estimating labor market power: Schubert et al. (2022), Azar et al. (2022), Azar et al. (2018), Azar et al. (2019), Roussille and Scuderi (2024), Felix (2022), Berger et al. (2022), Deb et al. (2024), Card et al. (2018), Yeh et al. (2022), Azar et al. (2022), Berger et al. (2024). Our model of the labor market is based on Berger et al. (2022) and our empirical strategy for estimating labor supply parameters follows Felix (2022). Second, there has been a growing realization that how one defines labor markets is critical for correctly estimating labor market power. Our method for defining labor markets builds on Fogel and Modenesi (2022) and is closely related to Nimczik (2018) and Jarosch et al. (2024). Relative to this literature, we (i) provide a data-driven method for assigning workers to heterogeneous worker types, and (ii) derive the implications of misclassification of jobs into markets for estimating labor supply elasticities. Finally, though we focus on labor markets, there are close ties to the IO literature on differentiated product markets and horizontal mergers: Farrell and Shapiro (2010), Baker and Bresnahan (1988), Miller and Weinberg (2017), Nevo (2001), Kaplow (2011), Kaplow (2010), Hovenkamp and Shapiro (2018), Harris and Simons (1991).

2 Model

We adopt a version of the basic framework of Berger et al. (2022) with ex ante homogenous workers. The model can be extended to allow for worker-level heterogeneity in skills/productivity, but we abstract from that case here for ease of exposition. We will return to this topic in subsequent drafts. A unit mass of individuals, indexed by i, match with a finite number of jobs, indexed by $j \in [1, \ldots, J]$. Jobs are nested within firms and may hire multiple workers at the same time. We assume hiring decisions are independent across jobs within firms; that is, there

A market definition is a partition of the set of J jobs into mutually-exclusive and exhaustive subsets. Denote the "true" market definition by Γ with elements γ: each job j belongs to one "true" labor market γ. For our purposes, labor market definition encodes the sets of jobs that compete most closely with one another over similar workers. Let $w_{j\gamma}$ denote the log wage paid by job j in market γ, and let $\ell_{j\gamma}$ denote log employment of job j in market γ.

Labor Supply. Workers make discrete choices over job opportunities, with indirect utility of working for job j in market γ equal to:

$$V_{ij\gamma} = w_{j\gamma} + \xi_{ij\gamma}, \tag{1}$$

where $\xi_{ij\gamma}$ is an idiosyncratic preference shock. We assume the $\xi_{ij\gamma}$ have the following joint distribution:

$$G\big(\{\xi_{ij\gamma}\}\big) = \exp\Bigg\{ -\sum_{\gamma} \Big[\sum_{j\in\gamma} \exp\big(-(1+\eta)\,\xi_{ij\gamma}\big) \Big]^{\frac{1+\theta}{1+\eta}} \Bigg\}. \tag{2}$$

which leads to a nested logit choice structure. In particular, the correlation of idiosyncratic utility shocks is higher for jobs in the same market γ. The parameter $\eta > 0$ governs the degree of substitution within markets (among jobs in the same γ), and $\theta > 0$ governs substitution between markets. As η increases, workers view jobs within a market as more substitutable (less differentiation within the nest). As θ increases, workers view markets themselves as closer substitutes (less segmentation between nests). The model is only sensible if $\eta \geq \theta$, and if $\theta = \eta$ the nesting structure collapses to a standard logit.

This model of worker preferences gives rise to the usual nested logit representation of labor supply to the firm:

$$\exp(\ell_{j\gamma}) = \frac{\exp((1+\eta)w_{j\gamma})}{\sum_{k\in\gamma}\exp((1+\eta)w_{k\gamma})} \times \frac{\big[\sum_{k\in\gamma}\exp((1+\eta)w_{k\gamma})\big]^{\frac{1+\theta}{1+\eta}}}{\sum_{\gamma'}\big[\sum_{k\in\gamma'}\exp((1+\eta)w_{k\gamma'})\big]^{\frac{1+\theta}{1+\eta}}} \tag{3}$$

Aggregating to the market level, we can write the total labor supply to market γ as:

$$\exp(\ell_{\gamma}) = \sum_{j\in\gamma}\exp(\ell_{j\gamma}) = \frac{\big[\sum_{k\in\gamma}\exp\big((1+\eta)w_{k\gamma}\big)\big]^{\frac{1+\theta}{1+\eta}}}{\sum_{\gamma'}\big[\sum_{k\in\gamma'}\exp\big((1+\eta)w_{k\gamma'}\big)\big]^{\frac{1+\theta}{1+\eta}}}\,, \tag{4}$$

For both expositional and conceptual reasons, it is useful to define the following wage indices, which are (up to scaling) equivalent to model-implied inclusive values (expected choice utilities):

$$\omega_{\gamma} \equiv \frac{1}{1+\eta}\,\log\Big(\sum_{k\in\gamma}\exp\big((1+\eta)w_{k\gamma}\big)\Big)$$
$$\omega \equiv -\log\Big(\sum_{\gamma'}\exp\big((1+\theta)\omega_{\gamma'}\big)\Big)$$

where ω_{γ} is the wage index for market γ and ω is an aggregate wage index. Given this

notation, the (log) labor supply curve to firm j in market γ is:

$$\ell_{j\gamma} = \omega \; + \; (\theta - \eta)\, \omega_\gamma \; + \; (1 + \eta)\, w_{j\gamma}\,, \tag{5}$$

and the (log) labor supply curve to market γ is:

$$\ell_\gamma = \omega \; + \; (1 + \theta)\, \omega_\gamma.$$

Labor supply to job j in market γ is increasing in its own wage $w_{j\gamma}$, decreasing in the wages of its rivals in market γ (since $\theta - \eta < 0$), and inherits a constant elasticity form thanks to the (nested) extreme value distributional assumption. Analogously, the inverse labor supply curve can be written

$$w_{j\gamma} = \lambda + \left(\frac{1}{\theta} - \frac{1}{\eta}\right) \lambda_\gamma + \frac{1}{\eta} \ell_{j\gamma}, \tag{6}$$

where λ_γ is a market-level labor supply index, and λ is an aggregate labor supply index.

At a high level, this model is vastly simplified relative to our intuition about how real workers trade off between employment opportunities. In particular, it is almost certainly over-restrictive to assume that substitution patterns between jobs can be described by just two preference parameters (η and θ) that are common to all workers. It is also over-restrictive to assume that workers' substitution patterns between pairs of jobs can be boiled down to asking whether the jobs are in the same "market" or not: more realistically, workers' preferences over firms likely vary smoothly or continuously. Even so, the basic model of preferences we have specified (which conforms to the standard model used by a number of studies) can be thought of as providing the best two-parameter approximation to the true pattern of cross-price elasticities for every pair of jobs. In other words, the simple model specified here is still useful as an approximation of the full $J \times J$ matrix of *diversion ratios* $D_{j \to k} \equiv \left| \frac{\partial \ell_k}{\partial w_j} / \frac{\partial \ell_j}{\partial w_j} \right| \exp(\ell_j - \ell_k)$. If we knew the diversion ratio (or, equivalently, the cross-price elasticities $\frac{\partial \ell_k}{\partial w_j}$) for each pair of firms, we could make precise statements about predicted markdowns without imposing market boundaries.

Labor Demand: Wage-Setting Conduct and Markdowns. Labor demand is governed by profit maximization, which gives rise to the generalized first-order condition:

$$\text{MRPL}_{j\gamma} = \text{MCL}^*_{j\gamma}.$$

This equation encodes the idea that profit-maximizing firms will choose employment and/or wage levels in job j that equate the marginal revenue product of labor ($\text{MRPL}_{j\gamma}$) to a generalized notion of the marginal cost of labor $\text{MCL}^*_{j\gamma}$, which differs depending upon the assumed model of firm wage-setting conduct. If firms engage in Cournot (quantity) competition, the FOC is:

$$\text{MRPL}_{j\gamma} = \exp(w_{j\gamma}) \cdot \left(1 + \frac{\partial w_{j\gamma}}{\partial \ell_{j\gamma}} \frac{\ell_{j\gamma}}{w_{j\gamma}}\right) = \exp(w_{j\gamma}) \cdot \left(1 + \varepsilon^{w\ell}_{j\gamma}\right).$$

If firms engage in Bertrand (price) competition, the FOC is:

$$\text{MRPL}_{j\gamma} = \exp(w_{j\gamma}) \cdot \left(1 + 1/\left[\frac{\partial \ell_{j\gamma}}{\partial w_{j\gamma}} \frac{w_{j\gamma}}{\ell_{j\gamma}}\right]\right) = \exp(w_{j\gamma}) \cdot \left(1 + 1/\varepsilon^{\ell w}_{j\gamma}\right).$$

In general $\varepsilon^{w\ell}_{j\gamma} \neq 1/\varepsilon^{\ell w}_{j\gamma}$, although there are special cases where this equality holds. One such case occurs if firms engage in *monopsonistic competition*, in which every firm behaves as if it were atomistic. Under monopsonistic competition, firms perceive $\varepsilon^{w\ell}_{j\gamma} = 1/\varepsilon^{\ell w}_{j\gamma} = 1/\eta$ (leading to proportional wage markdowns that are constant across firms). Going forward, we adopt Bertrand (price) competition as our primary assumption on firms' wage-setting conduct.

In equilibrium, it is straightforward to show that the inverse labor supply elasticity is a weighted average of the within- and between-market elasticities, where the weights are job j's payroll-weighted share of market γ employment:

$$\varepsilon^{\ell w}_{j\gamma} = \eta + (\theta - \eta)\, s_{j\gamma}$$

where $s_{j\gamma}$ is j's *payroll share* in market γ:

$$s_{j\gamma} \equiv \frac{\exp(w_{j\gamma}) \cdot \exp(\ell_{j\gamma})}{\sum_{k\in\gamma} \exp(w_{k\gamma}) \cdot \exp(\ell_{k\gamma})}. \tag{7}$$

The payroll-weighted market average elasticity of labor supply $\varepsilon^{\ell w}_{\gamma}$ is therefore:

$$\varepsilon^{\ell w}_{\gamma} = \textstyle\sum_{j\in\gamma} s_{j\gamma} \varepsilon^{\ell w}_{j\gamma} = \eta \underbrace{\left[\sum_{j\in\gamma} s_{j\gamma}\right]}_{=1} + (\theta - \eta) \underbrace{\left[\sum_{j\in\gamma} s^2_{j\gamma}\right]}_{\equiv \text{HHI}_\gamma} = \eta + (\theta - \eta)\ \text{HHI}_\gamma\,, \tag{8}$$

where HHI_γ is market γ's (payroll) Herfindahl index. Define firm j's wage markdown as the proportional gap between the MRPL and the wage paid by j:

$$\mu_{j\gamma} \equiv \frac{\text{MRPL}_{j\gamma} - \exp(w_{j\gamma})}{\exp(w_{j\gamma})}$$

Under Cournot competition, j's wage markdown is exactly equal to the inverse elasticity of labor supply to the firm, and the payroll-weighted market average wage markdown is equal to the payroll-weighted average elasticity:[3]

$$\mu_{j\gamma} = \Big[\eta + (\theta - \eta)\ s_{j\gamma}\Big]^{-1}, \text{ and } \mu_\gamma = \textstyle\sum_{j\in\gamma} s_{j\gamma}\mu_{j\gamma} \gtrsim \Big[\eta + (\theta - \eta)\ \text{HHI}_\gamma\Big]^{-1}. \tag{9}$$

These formulas encode the intuitive notion that larger firms are able to exercise a greater degree of wage-setting power than smaller firms. If firm j is very small in its market ($s_{j\gamma} \approx 0$), then $\varepsilon_{j\gamma} \approx \eta$ — its labor supply elasticity is essentially the elasticity of substitution among jobs within the market. In other words, if a small firm cuts its wages, its workers can readily find outside employment opportunities in the same market. Conversely, if firm j is the dominant employer in γ ($s_{j\gamma} \approx 1$), then $\varepsilon_{j\gamma} \approx \theta$ — its labor supply elasticity is essentially the elasticity of substitution between markets. In other words, if a large employer cuts its wages, outside employment opportunities at firms in the same market are relatively scarce.

Under a particular conduct assumption, the combination of market shares and labor supply parameters η and θ are enough to identify wage markdowns given firm-level observations of wages and employment. Specifically, in both the baseline Cournot and Bertrand models, markdown variation across firms depends upon the distribution of firm size within markets, and markdown variation across markets depends upon the between-market distribution of HHI. The extent to which variation in firm size and HHI translates into variation in markdowns is scaled by $\left(\frac{1}{\theta} - \frac{1}{\eta}\right)$:

$$\text{Var}\big(\mu_{j\gamma}\big) \ = \ \left(\frac{1}{\theta} - \frac{1}{\eta}\right)^2 \text{Var}\big(s_{j\gamma}\big), \qquad \text{Var}\big(\mu_\gamma\big) \ = \ \left(\frac{1}{\theta} - \frac{1}{\eta}\right)^2 \text{Var}\big(\text{HHI}_\gamma\big). \tag{10}$$

These equations illustrate the two key sources of market definition bias: one coming from mis-measurement of relative firm size ($s_{j\gamma}$, and one coming from mis-measurement of the elasticities η and θ. While researchers have long recognized the importance of the former, to our knowledge the consequences of the latter have not been rigorously explored.

3 Misclassification Bias in Elasticity Estimates

Thus far, we have described equilibrium outcomes given the correct segmentation of jobs into markets. We now consider the econometric problem of an analyst who does not know the true market classification and instead estimates the labor supply elasticities subject to some

[3]Note that μ_γ is an expectation of $\mu_{j\gamma}$ with respect to the payroll-share weights $s_{j\gamma}$. The lower bound of μ_γ in equation (15) follows from Jensen's inequality applied to the convex function $\mu_{j\gamma} = [\eta + (\theta - \eta)s_{j\gamma}]^{-1}$.

misclassification. Specifically, suppose that while the true market definition is given by a collection of partitions $\gamma \in \Gamma$, the researcher defines markets by some other grouping $m \in M$ — for instance, the intersection of occupation and geography.

How does misspecification of M relative to Γ affect estimated markdowns? There are two ways that misclassification will lead to bias in $\hat{\mu}$. First, misclassification can induce mismeasurement of job-level payroll shares ($s_{j\gamma}$) and therefore market-level HHIs (HHI_γ). We call this the "denominator problem." Second, misclassification can induce important biases in estimates of the elasticity parameters, η and θ. We call this "elasticity bias."

The Denominator Problem

The first source of measurement error has been widely acknowledged in prior studies of both product and labor market power. It arises because correct specification of both job-level market shares $s_{j\gamma} \equiv \frac{\exp(w_{j\gamma})\cdot\exp(\ell_{j\gamma})}{\sum_{k\in\gamma}\exp(w_{k\gamma})\cdot\exp(\ell_{k\gamma})}$ and the Herfindahl–Hirschman Index $\text{HHI}_\gamma = \sum_{j\in\gamma} s_{j\gamma}^2$, require a correct classification of which firms belong to the same market. Misclassifying jobs therefore contaminates the denominator of these expressions even when the numerator (the job's own payroll) is accurately measured. Drawing markets too broadly artificially deflates shares and HHIs, while drawing them too narrowly artificially inflates them. Moreover, misclassification may cause us to simultaneously overestimate market power in some markets while underestimating market power in other markets. Even if market boundaries are drawn at approximately the correct scale, misclassification of individual firms may drastically overstate or understate their own relative size.

Statistical Bias

While the denominator problem is a well-known issue, the statistical bias induced by using an incorrect market definition has not been as widely acknowledged. The remainder of this section describes this bias. We begin by describing a strategy for estimating η and θ, before turning to the implications of misclassification.

Intuitively, $\hat{\eta}$ is identified by comparing variation in within-market wage responses to within-market labor demand shocks. Larger covariances between wages and employment within markets imply larger estimates of $\hat{\eta}$. When using a market definition that is "too large," part of the variation used to estimate $\hat{\eta}$ comes from comparisons of wage responses across firms in different "true" markets, leading us to shrink $\hat{\eta}$ towards θ. $\hat{\theta}$ is estimated by comparing variation in between-market wage responses to between-market labor demand shocks. Larger covariances between wages and employment between markets imply larger estimates of $\hat{\theta}$. When using a market definition that is "too small," part of the variation used

to estimate $\hat{\theta}$ comes from comparisons of average wage responses of firms within the same "true" market, leading us to inflate $\hat{\theta}$ towards η.

3.1 Labor supply parameter estimation

For expositional purposes, we begin by assuming that the only source of unobserved heterogeneity is the presence of idiosyncratic taste shocks $\xi_{ij\gamma}$ that are uncorrelated with wages—in other words, there are no endogeneity issues. Under these (ideal) assumptions, we can use the structure of the model to motivate a pair of OLS regressions to estimate η and θ. This is the basic logic of Felix (2022)'s estimation strategy, although in practice estimation is accomplished via a pair of IV regressions to account for endogeneity. The logic described here for the OLS case extends to both the IV case and sophisticated estimation procedures.

Recall from equation 7 that the firm-specific labor supply curves can be written

$$\ell_{j\gamma} = \omega + (\theta - \eta)\,\omega_\gamma + (1+\eta)\,w_{j\gamma} \tag{11}$$

Defining a market-level fixed effect $\delta_\gamma \equiv \omega + (\theta-\eta)\,\omega_\gamma$ and assuming an orthogonal regression error term $\nu_{j\gamma}$ allows us to identify $1+\eta$ by regressing firm-level employment on firm-level wages, controlling for a market fixed effect.

$$\ell_{j\gamma} = \delta_\gamma + (1+\eta)\,w_{j\gamma} + \nu_{j\gamma} \tag{12}$$

Given $\hat{\eta}$, we can compute the market-level wage index: $\hat{\omega}_\gamma = \frac{1}{1+\hat{\eta}} \log\Big(\sum_{k\in\gamma} \exp\big((1+\hat{\eta})w_{k\gamma}\big)\Big)$ and estimate θ by regressing the market-level labor supply on the market-level wage aggregate

$$\ell_\gamma = \delta + (1+\theta)\,\hat{\omega}_\gamma + \nu_\gamma \tag{13}$$

which is the market-level labor supply curve derived above, with an additional regression error term.

If the market definition is correct, the first regression yields an unbiased estimate $\hat{\eta}$. The second regression yields an unbiased estimate $\hat{\theta}$ when both the correct market definition is used and $\hat{\omega}_\gamma$ is constructed using the true value of η.

3.2 Misclassification bias in $\hat{\eta}$

We derive the bias in the IV estimator of the within-market elasticity. The estimating equations implied by the structural model, demeaned within true markets $\gamma \in \Gamma$, are:

$$\ell_j - \bar{\ell}_{\gamma(j)} = (1+\eta)\,(w_j - \bar{w}_{\gamma(j)}) + \nu_j,$$
$$w_j - \bar{w}_{\gamma(j)} = \pi\,(z_j - \bar{z}_{\gamma(j)}) + (\zeta_j - \bar{\zeta}_{\gamma(j)}),$$

where z_j is the instrument (e.g. the tariff-change exposure of Felix 2022) and π is the first-stage coefficient. When the researcher instead uses the mismeasured market definition M, the IV estimator takes the form

$$\hat{\eta} = \frac{\mathrm{Cov}(\ell_j,\, z_j - \bar{z}_{m(j)})}{\mathrm{Cov}(w_j,\, z_j - \bar{z}_{m(j)})} - 1.$$

Plugging in the labor-supply equation and rearranging, the bias can be expressed as

$$\mathrm{Bias}(\hat{\eta}) = (\theta - \eta)\Big[1 - \beta_w + \Delta_w\Big],$$

where β_w is the leading attenuation term and Δ_w collects higher-order non-linear terms that are approximately zero in practice ($\Delta_w \approx 0$). To characterize β_w, define:

$$\rho_w \equiv \mathrm{Corr}\big(z_j - \bar{z}_{\gamma(j)},\; z_j - \bar{z}_{m(j)}\big),$$
$$\sigma^2_{\Gamma,w} \equiv \mathrm{Var}\big(z_j - \bar{z}_{\gamma(j)}\big), \qquad \sigma^2_{M,w} \equiv \mathrm{Var}\big(z_j - \bar{z}_{m(j)}\big),$$
$$\kappa_w \equiv \frac{\mathrm{Cov}\big(\bar{\zeta}_{\gamma(j)},\; z_j - \bar{z}_{m(j)}\big)}{\mathrm{Var}(z_j - \bar{z}_{m(j)})}, \qquad \tau_w \equiv \frac{\mathrm{Cov}\big(\zeta_j - \bar{\zeta}_{\gamma(j)},\; z_j - \bar{z}_{m(j)}\big)}{\mathrm{Var}(z_j - \bar{z}_{m(j)})}.$$

Here ρ_w is the within-market correlation of the instrument across the two definitions; $\sigma_{\Gamma,w}/\sigma_{M,w}$ is the ratio of within-market standard deviations; κ_w is a misclassification-induced omitted-variable bias term; and τ_w is a spurious overfitting term that is negligible when the instrument is exogenous within true markets. Assuming $\tau_w \approx 0$:

$$\beta_w = \frac{\pi\,\rho_w\,\frac{\sigma_{\Gamma,w}}{\sigma_{M,w}} + \tau_w}{\pi + \kappa_w + \tau_w} \;\approx\; \rho_w \cdot \frac{\sigma_{\Gamma,w}}{\sigma_{M,w}} \cdot \frac{\pi}{\pi + \kappa_w}. \tag{14}$$

This expression reveals three distinct channels through which misclassification attenuates $\hat{\eta}$ toward θ. First, instrument-ranking misalignment ($\rho_w < 1$): ρ_w falls below one when M fails to preserve the relative ranking of jobs within their true market γ with respect to the instrument. If a job receives a large demand shock relative to its correct market γ but a

small shock relative to the wrong market m, the within-market variation in z_j that should identify η is diluted. Second, variance contamination ($\sigma_{\Gamma,w} < \sigma_{M,w}$): when the mismeasured definition M is too coarse, $\bar{z}_{m(j)}$ absorbs less market-level variation, leaving more cross-market variation in $z_j - \bar{z}_{m(j)}$. The estimator then conflates within- and between-market variation, shrinking $\hat{\eta}$ toward θ. Third, misclassification-induced OVB ($\kappa_w \neq 0$): even if the instrument is valid conditional on the correct market fixed effects, using the wrong market FEs can induce correlation between the residualized instrument and the first-stage error $\bar{\zeta}_{\gamma(j)}$, biasing the IV estimator.

Setting $\Delta_w = 0$ and combining, the bias simplifies to:

$$\text{Bias}(\hat{\eta}) \approx (\theta - \eta)\left[1 - \underbrace{\rho_w \cdot \frac{\sigma_{\Gamma,w}}{\sigma_{M,w}} \cdot \frac{\pi}{\pi + \kappa_w}}_{\beta_w}\right].$$

Two important special cases arise when market definitions are nested. Formally, M is nested within Γ if for every $m \in M$ there exists $m' \in \Gamma$ such that $j \in m \Rightarrow j \in m'$. If M is nested within Γ (finer definition used incorrectly), then $\beta_w = 1$ and $\text{Bias}(\hat{\eta}) = 0$: defining markets too narrowly does not bias the within-market elasticity. If Γ is nested within M (too-coarse definition used), then $\beta_w < 1$ and the bias is positive, shrinking $\hat{\eta}$ toward θ. In this case the magnitude of the bias is increasing in the ratio $\sigma_{\Gamma,w}/\sigma_{M,w}$, which measures the degree of excess coarseness.

3.3 Misclassification bias in $\hat{\theta}$

The implications of misclassification for $\hat{\theta}$ follow analogously. The structural estimating equations at the market level are:

$$\begin{aligned}
\ell_\gamma &= \omega + (1+\theta)\,\omega_\gamma + \nu_\gamma, \\
\omega_\gamma &= \pi \bar{z}_\gamma + \bar{\zeta}_\gamma + \Delta_\gamma,
\end{aligned}$$

where ω_γ is the market-level wage index, $\bar{z}_\gamma$ is the market-level mean instrument, and Δ_γ is a nonlinear approximation error. Using the mismeasured market definition M, the IV estimator is

$$\hat{\theta} = \frac{\text{Cov}(\ell_m,\, \bar{z}_m)}{\text{Cov}(\omega_m,\, \bar{z}_m)} - 1.$$

The resulting bias expression is

$$\text{Bias}(\hat{\theta}) = (\eta - \theta)\Big[1 - \beta_b + \Delta_b\Big],$$

where β_b is the between-market attenuation coefficient and $\Delta_b \approx 0$ collects higher-order terms. Define the between-market analogues of the quantities introduced above:

$$\begin{aligned}
\rho_b &\equiv \text{Corr}\big(\bar{z}_{\gamma(j)},\ \bar{z}_{m(j)}\big), \\
\sigma^2_{\Gamma,b} &\equiv \text{Var}\big(\bar{z}_{\gamma(j)}\big), \qquad \sigma^2_{M,b} \equiv \text{Var}\big(\bar{z}_{m(j)}\big), \\
\kappa_b &\equiv \frac{\text{Cov}\big(\zeta_j - \bar{\zeta}_{\gamma(j)},\ \bar{z}_{m(j)}\big)}{\text{Var}(\bar{z}_{m(j)})}, \qquad \tau_b \equiv \frac{\text{Cov}\big(\bar{\zeta}_{\gamma(j)},\ \bar{z}_{m(j)}\big)}{\text{Var}(\bar{z}_{m(j)})}.
\end{aligned}$$

Here ρ_b is the correlation between the market-level instrument under the two definitions; κ_b is a misclassification-induced OVB term; and $\tau_b \approx 0$ is a spurious overfitting term. Assuming $\tau_b \approx 0$:

$$\beta_b = \frac{\pi\, \rho_b \frac{\sigma_{\Gamma,b}}{\sigma_{M,b}} + \tau_b}{\pi + \kappa_b + \tau_b} \ \approx\ \rho_b \cdot \frac{\sigma_{\Gamma,b}}{\sigma_{M,b}} \cdot \frac{\pi}{\pi + \kappa_b}.$$

This expression has the same three-channel structure as β_w, but the bias runs in the opposite direction — misclassification inflates $\hat{\theta}$ toward η. First, instrument-ranking misalignment ($\rho_b < 1$): ρ_b falls below one when M fails to preserve the ranking of markets by their average instrument value. Market-level variation in z_j that should identify θ is absorbed by the wrong market fixed effects. Second, variance contamination ($\sigma_{\Gamma,b} < \sigma_{M,b}$): when M is too granular, $\bar{z}_{m(j)}$ picks up within-market variation, inflating $\sigma_{M,b}$ relative to $\sigma_{\Gamma,b}$ and shrinking β_b below one. The estimator then conflates between- and within-market variation, pushing $\hat{\theta}$ toward η. Third, misclassification-induced OVB ($\kappa_b \neq 0$): using the wrong market partitioning can correlate the residualized between-market instrument with job-level shocks $\zeta_j - \bar{\zeta}_{\gamma(j)}$, biasing the IV estimator.

Setting $\Delta_b = 0$, the bias simplifies to:

$$\text{Bias}(\hat{\theta}) \approx (\eta - \theta)\left[1 - \underbrace{\rho_b \cdot \frac{\sigma_{\Gamma,b}}{\sigma_{M,b}} \cdot \frac{\pi}{\pi + \kappa_b}}_{\beta_b}\right].$$

The nesting cases apply in reverse relative to the $\hat{\eta}$ case. If M is nested within Γ (too-fine definition used for the between-market regression), then $\beta_b < 1$ and $\hat{\theta}$ is inflated toward η. If Γ is nested within M (too-coarse definition), then $\beta_b = 1$ and there is no bias. Intuitively,

defining markets too granularly splits firms that workers view as close substitutes into separate markets, treating true within-market variation as between-market variation and contaminating the estimate of θ.

3.4 Sign of Bias

The expressions derived above establish how the extent of misclassification governs the magnitude of bias in $\hat{\eta}$ and $\hat{\theta}$. A natural further question is whether the direction of the bias can be signed without strong additional assumptions. The following theorem answers this affirmatively.

Theorem 1 (Sign of Misclassification Bias)**.** *Maintain the modeling assumptions above and set* $\Delta_w = \Delta_b = 0$.
(i) ***OLS.*** *When* $z_j = w_j$,

$$\beta_w = \rho_w \cdot \frac{\sigma_{\Gamma,w}}{\sigma_{M,w}} \;\leq\; 1 \quad\Longrightarrow\quad \mathrm{Bias}(\hat{\eta}) = (\theta - \eta)\big[1 - \beta_w\big] \;\leq\; 0.$$

Similarly,

$$\beta_b = \rho_b \cdot \frac{\sigma_{\Gamma,b}}{\sigma_{M,b}} \;\leq\; 1 \quad\Longrightarrow\quad \mathrm{Bias}(\hat{\theta}) = (\eta - \theta)\big[1 - \beta_b\big] \;\geq\; 0.$$

(ii) ***IV.*** *When* $z_j \neq w_j$*, the same inequalities hold provided* $\tau_w = \tau_b = 0$ *and*

$$\mathrm{sign}(\pi) = \mathrm{sign}(\pi + \kappa_w) = \mathrm{sign}(\pi + \kappa_b),$$

where $\tau_w \equiv \mathrm{Cov}(\zeta_j - \bar{\zeta}_{\gamma(j)},\, z_j - \bar{z}_{m(j)}) / \mathrm{Var}(z_j - \bar{z}_{m(j)})$ *and* $\tau_b \equiv \mathrm{Cov}(\bar{\zeta}_{\gamma(j)},\, \bar{z}_{m(j)}) / \mathrm{Var}(\bar{z}_{m(j)})$ *are spurious overfitting terms that are typically negligible when the instrument is sufficiently strong.*

Proof sketch. The key inequalities $\beta_w \leq 1$ and $\beta_b \leq 1$ in the OLS case rest on a projection orthogonality argument. Let P_Γ and P_M be the orthogonal projection matrices onto the column spaces of the true and mismeasured market indicator matrices. Then $\beta_w \leq 1$ is equivalent to $E[\bar{w}_{\gamma(j)}(w_j - \bar{w}_{m(j)})] \geq 0$, i.e., $z' P_\Gamma (I - P_M) z \geq 0$. Decomposing the right factor as $z = P_\Gamma z + (I - P_\Gamma) z$, the cross term $z' P_\Gamma (I - P_M)(I - P_\Gamma) z$ vanishes because $P \cdot C \cdot (I - P) = 0$ for any projection P (since columns of PC lie in $\mathrm{col}(P)$, which is killed by $(I - P)$), while the remaining term equals $\|(I - P_M) P_\Gamma z\|^2 \geq 0$. The proof of $\beta_b \leq 1$ is symmetric. The IV extension requires $\mathrm{sign}(\pi) = \mathrm{sign}(\pi + \kappa_w) = \mathrm{sign}(\pi + \kappa_b)$ to preserve the direction of the first stage. A complete proof is provided in Appendix A. □

The upshot of Theorem 1 is sharp: *misclassification can only compress the estimated elasticity gap $\hat{\eta} - \hat{\theta}$ toward zero, never expand it.* It is impossible for misclassification to cause $\hat{\eta}$ to be biased upward or $\hat{\theta}$ to be biased downward. This one-directional result has a clean empirical implication: if one market definition yields a larger gap $\hat{\eta} - \hat{\theta}$ than another, the definition with the larger gap is, ceteris paribus, closer to the true market boundaries.

The IV conditions in part (ii) require that the sign of the net first-stage response $\pi + \kappa_w$ (resp. $\pi + \kappa_b$) agrees with the sign of π itself. Intuitively, κ_w measures the degree to which the mismeasured within-market instrument picks up true between-market unobserved shocks; as long as this "leakage" does not reverse the direction of the first stage—a condition that holds whenever the instrument is sufficiently strong relative to the market-level correlation—the sign result goes through. In our empirical application the within-market first-stage F-statistic exceeds 87 under all specifications (Table 5), providing reassurance that this condition holds for the within-market regression. The between-market first stage is considerably weaker (discussed further in Section 5), so the theorem provides a stronger guarantee for the sign of the bias in $\hat{\eta}$ than for $\hat{\theta}$ in that setting.

3.5 Simulations

To illustrate the practical relevance of the attenuation bias induced by using an incorrect market definition, we conduct a number of simulation exercises. For the purposes of these simulations, we assume that the marginal revenue product of each job, $\text{MRPL}_{j\gamma}$, is constant. Without loss of generality, we can decompose $\log(\text{MRPL}_{j\gamma})$ as the sum of market-γ- and firm-j-specific components, leading to the following representation of log wages:

$$w_{j\gamma} = \underbrace{\tau_\gamma + \tau_j}_{=\log(\text{MRPL}_{j\gamma})} - \log(1 + \mu_{j\gamma})$$

Our initial set of simulation exercises illustrate special cases of misclassification. When conducting each of these initial simulations, we set $\theta = 1$ and $\eta = 5$. Across simulations, the number of "true" markets is fixed at 125. The number of jobs in each market is determined by drawing a Poisson(40) random variable, so the expected number of jobs is 5,000. We generate the τ_γ as independent draws from a normal distribution with variance σ_b^2, and the τ_j as independent draws from a normal distribution with variance $1 - \sigma_b^2$, such that $\phi_\Gamma \approx \frac{1-\sigma_b^2}{\sigma_b^2}$. Given a set of jobs, markets, and productivity draws, we simulate equilibrium wage and employment levels as follows. We initialize the simulation by conjecturing an initial set of markdowns, $\mu_{j\gamma}^{(0)} = \frac{1}{\eta}$. At each iteration t of the simulation, we compute implied log wages at each job, $w_{j\gamma}^{(t)}$, by setting $w_{j\gamma}^{(t)} = \tau_\gamma + \tau_j - \log(1 + \mu_{j\gamma}^{(t-1)})$, and use them to construct the

market-level wage indices $\omega_\gamma^{(t)}$ and the overall wage index $\omega^{(t)}$. Given $w_{j\gamma}^{(t)}$, $\omega_\gamma^{(t)}$, and $\omega^{(t)}$, we compute implied log employment at each job as $\ell_{j\gamma}^{(t)} = \omega^{(t)} + (\theta - \eta)\omega\gamma^{(t)} + (1+\eta)w_{j\gamma}^{(t)}$. Finally, we use $\ell_{j\gamma}^{(t)}$ and $w_{j\gamma}^{(t)}$ to compute updated payroll shares $s_{j\gamma}^{(t)}$ and markdowns $\mu_{j\gamma}^{(t)}$. Once this process has iterated to convergence, we use the procedure outlined in Section 3.1 to recover estimates of η and θ under both the true market definition Γ and a series of misclassified market definitions M.

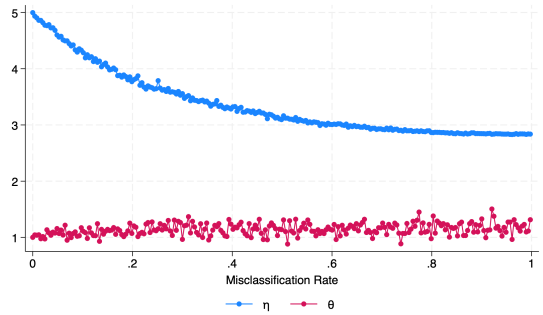
Figure 1 shows the results of these simulations. In Panels (a) and (b), we randomly perturb the market assignments of a varying fraction of jobs, ranging from 0 to 1, by drawing new market assignments randomly (according to the size distribution of the underlying true markets $\gamma \in \Gamma$). We label this fraction the "misclassification rate," although it is more accurately described as the fraction of jobs with perturbed labels. Panel (a) reports estimates from a simulation with $\sigma_b^2 = 0.5$, so $\phi_\Gamma \approx 1$. This scenario leads to relatively little bias in $\hat{\theta}$, but can induce significant bias in $\hat{\eta}$ even at relatively low levels of misclassifiction–for instance, when the "misclassification rate" is 20%, the difference between $\hat{\eta}$ and $\hat{\theta}$ is attenuated by roughly one third. Panel (b) reports estimates from a simulation with $\sigma_b^2 = 0.9$, so $\phi_\Gamma \approx 1/9$. This scenario reflects cases where wages are highly differentiated across markets, which causes enormous bias as a result of misclassification. When the "misclassification rate" is 20%, the difference between $\hat{\eta}$ and $\hat{\theta}$ is attenuated by nearly 95%. While the scenario considered by Panel (b) is unrealistically extreme, it helps to illustrate why misclassification bias may be large in realistic scenarios. A primary source of bias in this scenario comes from the fact that ρ_w falls substantially below one even when a small number of labels are perturbed. This is because markets are highly differentiated in this example, making it highly likely that a job's relative within-market rank (e.g. $w_{j\gamma} - \overline{w}_\gamma$) changes significantly when (erroneously) assigned to a new market. Similar dynamics can be at play even when markets are not as differentiated, so long as misclassification causes jobs that are high (low) wage relative to their true market γ to be assigned to markets m where they are low (high) wage.

Panels (c) and (d) consider cases where markets are either randomly subdivided or aggregated inappropriately, to illustrate the intuition of "nesting" cases. Panel (c) reports estimates from a simulation with $\sigma_b^2 = 0.25$ such that $\phi_\Gamma \approx 3$, where the incorrect market definition involves splitting true markets into increasingly smaller pieces. Panel (d) reports estimates from a simulation with $\sigma_b^2 = 0.75$ such that $\phi_\Gamma \approx 1/3$, where the incorrect market definition involves aggregating increasingly larger numbers of true markets together. As predicted by our theoretical results, using an incorrect market definition M that is nested within Γ (Panel (c)) yields an unbiased estimate $\hat{\eta}$, but can induce substantial bias in $\hat{\theta}$, while using an incorrect definition such that Γ is nested within M (Panel (d)) yields an unbiased estimate $\hat{\theta}$, but can induce substantial bias in $\hat{\eta}$. Section 4.3 reports the results of

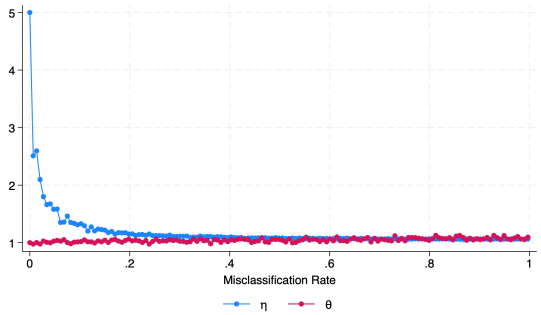
an additional simulation exercise calibrated to Brazilian administrative data.

Figure 1: Simulated elasticity estimates under misclassification: random relabeling and nesting cases

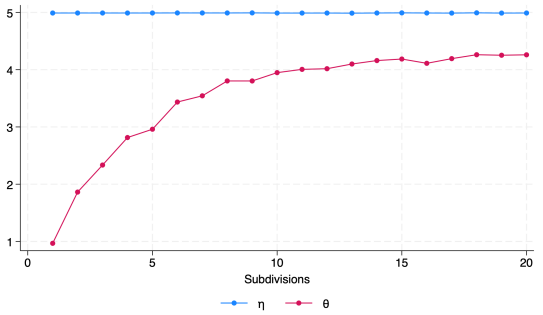
(a) Random misclassification, $\phi_\Gamma = 1$



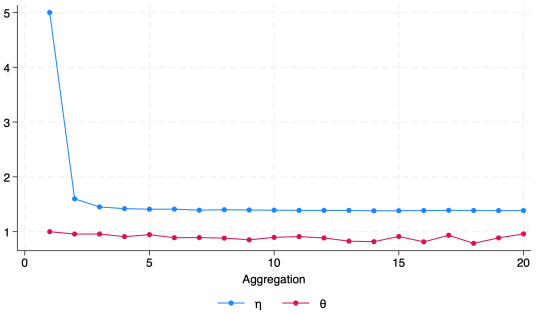
(b) Random misclassification, $\phi_\Gamma = 1/9$



(c) Nesting M within Γ, $\phi_\Gamma = 3$



(d) Nesting Γ within M, $\phi_\Gamma = 1/3$



Are the results of these simulation exercises realistic? While we provide direct evidence to answer this question in subsequent sections of this paper, we can already point to suggestive evidence from the literature. Berger et al. (2022) estimate a version of the model we specified on US data, using combinations of 3-digit industry NAICS codes and commuting zones to define labor markets. They report estimates $\hat{\eta} = 10.85$ and $\hat{\theta} = 0.42$. Deb et al. (2024) estimate an extremely similar model, using the same sample and identifying variation as Berger et al. (2022). However, Deb et al. (2024) use only industry classifications to define labor markets (no geography), estimating parameters separately for college- (H) and non-college-educated (L) workers. They report estimates $\hat{\eta}_H = 2.53$, $\hat{\theta}_H = 2.02$, $\hat{\eta}_L = 2.42$, and $\hat{\theta}_L = 1.85$.

3.6 Generalizing the bias beyond nested logit

The bias results derived in Sections 3.2–3.3 exploit the nested logit structure: the "true" market definition Γ enters through the two-level substitution hierarchy governed by η and

θ. In practice, however, the nested logit is at best an approximation to the full matrix of cross-wage elasticities—an observation we made explicitly in Section 2. Workers' substitution patterns between pairs of jobs likely vary smoothly rather than jumping discretely at market boundaries, and pairs of markets may differ in how substitutable they are for one another.

In this subsection we embed the same IV estimator in a more general mixed logit framework that replaces the two-parameter (η, θ) structure with a continuous correlation matrix P over worker tastes. We show that the bias of the IV estimator is governed by a single object—the gap $\Delta = P - P^G$ between the true taste-correlation matrix and the block-diagonal matrix implied by the IV grouping—and that the nested logit bias formulas of Sections 3.2–3.3 emerge as a special case.

3.6.1 A general mixed logit specification

Consider a generalization of the preference model in equation (1) in which worker i's indirect utility from job j is

$$V_{ij} \;=\; \beta\, w_j \;+\; \xi_j \;+\; \zeta_{ij} \;+\; \varepsilon_{ij},$$

where $\varepsilon_{ij} \sim \text{EV}_1$ (iid) as before, ξ_j is a job-specific fixed effect, and $\zeta_i = (\zeta_{i1}, \ldots, \zeta_{iJ})'$ is a vector of worker-specific taste shocks drawn from a distribution $G(\cdot)$ with

$$E[\zeta_{ij}] = 0, \qquad \text{Var}(\zeta_{ij}) = \sigma^2 \;\; \forall\, j, \qquad \text{Corr}(\zeta_{ij}, \zeta_{ik}) = P_{jk}.$$

The matrix P encodes the full substitution structure: P_{jk} is large when workers who like job j also tend to like job k, causing the two jobs to be close substitutes. In the nested logit of Section 2, this matrix takes the special form $P_{jk} = \mathbf{1}\{\gamma(j) = \gamma(k)\}$—a block-diagonal matrix with ones on each nest block—which enforces the binary same-market/different-market distinction. The general specification allows P to reflect arbitrary patterns of substitutability driven by geography, industry, occupation, amenities, or any other job characteristics.

Let $v_j \equiv \beta w_j + \xi_j$ denote the deterministic utility index, and define the base shares (the shares that would obtain absent heterogeneity) as $s_j^0 = \exp(v_j)/\sum_k \exp(v_k)$. Let s_j denote the observed (mixed logit) shares, which differ from s_j^0 due to the Jensen's inequality correction induced by the mixing distribution G.

3.6.2 Log-linear approximation to elasticities

We approximate the cross-wage elasticities using a log-linear expansion of the conditional choice probabilities around the no-heterogeneity baseline. Log-linearizing $\log s_j(\zeta)$ in ζ and

exponentiating yields

$$\frac{\partial \ell_j}{\partial w_k} \;\approx\; \beta\Big[\delta_{jk} \;-\; s_k \cdot c_{jk}\Big], \qquad c_{jk} \;\equiv\; \exp\big(\sigma^2 \widetilde{P}_{jk}\big), \tag{15}$$

where

$$\widetilde{P}_{jk} \;\equiv\; (e_j - s)^\top P\,(e_k - s) \;=\; P_{jk} - \bar{P}_j - \bar{P}_k + \bar{\bar{P}}$$

is the share-centered correlation, with $\bar{P}_j = \sum_l s_l P_{jl}$ and $\bar{\bar{P}} = \sum_{l,m} s_l s_m P_{lm}$.[4]

The centered correlation $\widetilde{P}_{jk}$ measures how much more (or less) correlated jobs j and k are relative to their share-weighted market averages. When $\sigma^2 = 0$, $c_{jk} = 1$ for all pairs and equation (33) reduces to the standard logit. When $\sigma^2 > 0$, c_{jk} varies with both j and k, breaking IIA in a pattern governed by the correlation structure P.

Connection to the nested logit. Under the nested logit, P is block-diagonal on the true nests Γ with $P_{jk} = \mathbf{1}\{\gamma(j) = \gamma(k)\}$, and the centered correlation reduces to

$$\widetilde{P}_{jk} = \begin{cases} 1 - 2S_\gamma + H_\Gamma & \text{if } \gamma(j) = \gamma(k), \\ H_\Gamma - S_{\gamma(j)} - S_{\gamma(k)} & \text{if } \gamma(j) \neq \gamma(k), \end{cases}$$

where $S_\gamma = \sum_{k \in \gamma} s_k$ is nest γ's aggregate share and $H_\Gamma = \sum_\gamma S_\gamma^2$ is the Herfindahl index of nest shares. The cross-wage elasticity then takes one of three values—own, same-market, and different-market—recovering the (η, θ) structure as a special case of the continuous P matrix.

3.6.3 The grouping-implied correlation and the misspecification matrix

The IV groups $m = 1, \ldots, M$ (whether administrative market definitions or any other grouping used for de-meaning) define an implicit "perfect substitutes" correlation structure:

$$P^G_{jk} \;=\; \mathbf{1}\{m(j) = m(k)\}.$$

This is the correlation matrix under which all jobs within a group are identical from the taste perspective: workers who like one job in a group like all jobs in that group equally. Define

[4]The log-linear approximation is the second step in a hierarchy of approximations to the mixed logit C matrix (the matrix of second moments of the conditional choice probabilities): a level-based delta method gives $c_{jk} \approx 1 + \sigma^2 \widetilde{P}_{jk}$, while a full Gaussian correction replaces $\sigma^2 P$ with a "shrunk" covariance $\Sigma_e = \sigma^2 P (I + 2\Omega\sigma^2 P)^{-1}$, where $\Omega = \mathrm{diag}(s) - ss'$ is the logit covariance matrix. The log-linear form strikes a balance between tractability and accuracy: it guarantees positivity of the implied C matrix and captures the Jensen's inequality correction that the linear delta method misses entirely.

the misspecification matrix

$$\Delta \;=\; P \;-\; P^G.$$

Its entries have a natural sign pattern. The diagonal entries are $\Delta_{jj} = 0$ for all j, since the own-taste variance is always correctly specified. For same-group pairs with $m(j) = m(k)$ and $j \neq k$, $\Delta_{jk} = P_{jk} - 1 \leq 0$, reflecting the fact that within-group correlations fall short of perfect substitutability. For cross-group pairs with $m(j) \neq m(k)$, $\Delta_{jk} = P_{jk} \geq 0$, since across-group correlations exceed the assumed value of zero.

A key property—which we verify below—is that the IV estimator is exactly unbiased when $\Delta = 0$. This is the natural generalization of the nested-logit result that correctly specified nests yield unbiased estimates: under the mixed logit, it is not the nest structure per se that matters, but whether the IV grouping captures the taste correlation structure.

3.6.4 Bias of the IV estimator

We work with the same indirect least squares estimator as in Section 3.2, de-meaning within the IV groups $B_1, \ldots, B_N$:

$$\hat{\beta} \;=\; \frac{\sum_j \ell_j \cdot (z_j - \bar{z}_{m(j)})}{\sum_j w_j \cdot (z_j - \bar{z}_{m(j)})},$$

where z_j is an instrument for w_j and $\bar{z}_{m(j)}$ is the mean of z_j within firm j's IV group.

Under the standard assumptions that z_j is iid with variance σ_z^2, independent across firms, and $w_j = \pi z_j + u_j$ with $u_j \perp z$, the first-stage covariance $E[w_k \cdot (z_j - \bar{z}_{m(j)})]$ is nonzero only for j and k in the same group. Substituting the log-linear elasticities (33) and simplifying (the algebra parallels the nested logit case with the general c_{jk} replacing the three-valued nested logit structure), we obtain:

Proposition 1 (General bias formula)**.** *Under the log-linear approximation (33) and the first-stage assumptions above,*

$$\hat{\beta} - \beta \;=\; -\frac{\beta}{J-N} \sum_{n=1}^{N} e^{\sigma^2 \gamma_n} \sum_{j \in B_n} s_j \left[e^{\sigma^2 \widetilde{\Delta}_{jj}} \;-\; \frac{1}{N_n} \sum_{l \in B_n} e^{\sigma^2 \widetilde{\Delta}_{lj}} \right], \tag{16}$$

where $\widetilde{\Delta}_{lj} = (e_l - s)^\top \Delta \, (e_j - s)$ is the share-centered misspecification, $\gamma_n = 1 - 2S_n + H_G$ depends only on group n's aggregate share $S_n = \sum_{j \in B_n} s_j$ and the Herfindahl index $H_G = \sum_n S_n^2$ of group shares, and $N_n = |B_n|$ is the number of firms in group n.

The bias (i) vanishes when $\Delta = 0$ (correctly specified groups), (ii) is always an attenuation ($\hat{\beta} < \beta$), and (iii) depends on Δ only through its within-group variation—the share-centered, group-demeaned misspecification.

Proof sketch. Under the log-linear approximation, the elasticity matrix is $E_{jk} = \beta[\delta_{jk} - s_k c_{jk}]$. Under the standard logit ($c_{jk} = 1$), the numerator of $\hat{\beta}$ reduces to $\beta(J - N)$ because the group-demeaning annihilates the constant cross-elasticity s_k. With heterogeneity, the "excess" correction $c_{jk} - 1$ survives the demeaning whenever c_{jk} varies with the row index j within a group—i.e., whenever IIA is violated within IV groups. Writing c_{jk} in terms of $\widetilde{P}_{jk} = \widetilde{P}^G_{jk} + \widetilde{\Delta}_{jk}$ and noting that $\widetilde{P}^G_{lj}$ is constant over l within each group (since P^G is block-diagonal), the P^G component factors out as the multiplicative prefactor $e^{\sigma^2 \gamma_n}$, and the within-group variation comes entirely from $\widetilde{\Delta}$. □

Intuition. The bias arises because the IV identifies β from the contrast between a firm's own-wage elasticity and the average cross-wage elasticity within its IV group. Under the standard logit, all cross-elasticities are equal ($c_{jk} = 1$), so this contrast is clean. With taste heterogeneity, the own-wage elasticity is damped by $c_{jj} > 1$ (workers who idiosyncratically prefer job j are less sensitive to its wage), but cross-elasticities c_{lj} for other firms l in the group are less damped because $\widetilde{P}_{lj} < \widetilde{P}_{jj}$ generically. The IV interprets this reduced contrast as a smaller β, producing attenuation.

3.6.5 Leading-order approximation

Expanding the exponentials in equation (38) to first order in σ^2 yields a more transparent expression:

$$\hat{\beta} - \beta \;\approx\; -\frac{\beta\,\sigma^2}{J - N} \sum_j s_j \Bigg[\underbrace{\left(1 - \bar{P}^n_j\right)}_{\substack{\text{within-group} \\ \text{imperfection}}} \;-\; \underbrace{\left(\bar{P}_j - \bar{\bar{P}}_n\right)}_{\substack{\text{across-group} \\ \text{leakage offset}}} \Bigg], \tag{17}$$

where $\bar{P}^n_j = \frac{1}{N_n} \sum_{l \in B_n} P_{lj}$ is firm j's average taste correlation with its IV group, $\bar{P}_j = \sum_l s_l P_{lj}$ is j's share-weighted average correlation with the entire market, and $\bar{\bar{P}}_n = \frac{1}{N_n} \sum_{l \in B_n} \bar{P}_l$ is the average market connectivity of firms in group n.

The first term, $(1 - \bar{P}^n_j)$, is large when firm j has low taste correlation with its group-mates. It captures the within-group shortfall: the IV grouping treats these firms as perfect substitutes, but workers do not view them that way. The second term, $(\bar{P}_j - \bar{\bar{P}}_n)$, provides a partial offset when firm j is more connected to the overall market than the average member of its group.

3.6.6 Decomposition into within- and across-group components

The misspecification matrix Δ decomposes naturally as

$$\Delta \;=\; \underbrace{\Delta^W}_{\text{within-group shortfall}} \;+\; \underbrace{\Delta^A}_{\text{across-group leakage}},$$

where $\Delta^W_{jk} = (P_{jk} - 1) \cdot \mathbf{1}\{m(j) = m(k),\, j \neq k\} \leq 0$ and $\Delta^A_{jk} = P_{jk} \cdot \mathbf{1}\{m(j) \neq m(k)\} \geq 0$. At leading order, the bias inherits this decomposition linearly:

$$\hat{\beta} - \beta \;\approx\; \underbrace{-\frac{\beta\sigma^2}{J-N}\,\mathrm{tr}\big(S\,\widetilde{G}\,\widetilde{\Delta}^W\big)}_{\text{within-group bias}} \;\; \underbrace{-\frac{\beta\sigma^2}{J-N}\,\mathrm{tr}\big(S\,\widetilde{G}\,\widetilde{\Delta}^A\big)}_{\text{across-group bias}},$$

where $S = \mathrm{diag}(s)$, $\widetilde{G} = I - \bar{G}$ is the within-group demeaning matrix, and $\widetilde{\Delta}^W$, $\widetilde{\Delta}^A$ are the share-centered versions of Δ^W and Δ^A. Both terms contribute negative (attenuating) bias.

3.6.7 Special cases

Nested logit. When P is block-diagonal on the true nests Γ—i.e., $P_{jk} = \mathbf{1}\{\gamma(j) = \gamma(k)\}$ for some within-nest correlation that we normalize to unity—the across-group leakage $\Delta^A = 0$ for any grouping M that is a refinement of Γ. The bias is then driven entirely by Δ^W, and equation (39) reduces to

$$\hat{\beta} - \beta \;\approx\; -\beta\sigma^2(1 - \rho_w) \cdot \frac{\sum_j s_j(1 - 1/N_{m(j)})}{J - N},$$

where ρ_w is the average within-nest correlation among group-mates. This recovers the structure of the attenuation term β_w in equation (24): the bias is proportional to $(1 - \rho_w)$, which measures the degree to which within-group correlations fall short of the perfect-substitute benchmark. When IV groups coincide with true nests, $\rho_w = 1$ and the bias vanishes, consistent with the nested case of Theorem 1.

Equicorrelated tastes. If $P_{jk} = \rho_0$ for all $j \neq k$, then $\widetilde{P}_{jk} = (1-\rho_0)(\delta_{jk} - s_j - s_k + H)$, and the entire correlation structure collapses to a single effective variance $\sigma^2_{\text{eff}} = \sigma^2(1 - \rho_0)$. The common component of taste (ρ_0) cancels from all substitution patterns because it shifts all jobs' utilities equally. The bias under equicorrelation with uniform shares is $\hat{\beta} - \beta \approx -\beta\sigma^2_{\text{eff}}/J$, which vanishes as $J \to \infty$.

Factor model. If $P_{jk} = (1-\gamma)\delta_{jk} + \gamma f_j' f_k$ for factor loadings $f_j \in \mathbb{R}^K$ (representing, e.g., industry, geography, or amenity bundles), then the leading-order bias depends on $\|\widetilde{f}_j - \widetilde{\bar{f}}_n\|^2$, where $\widetilde{f}_j = f_j - \sum_l s_l f_l$ is the demeaned loading and $\widetilde{\bar{f}}_n$ is the average loading in group n. Firms that are "outliers" in factor space relative to their IV group—e.g., a manufacturing firm grouped with service-sector firms—contribute most to the bias. This provides a concrete diagnostic: the bias is driven by the within-group dispersion of factor loadings, which can be computed from observables.

3.6.8 Relationship to the nested logit bias formulas

The bias formulas derived in Sections 3.2–3.3 can be mapped to the general framework as follows. The three channels identified in equation (24)—instrument-ranking misalignment (ρ_w), variance contamination ($\sigma_{\Gamma,w}/\sigma_{M,w}$), and misclassification-induced OVB (κ_w)—all arise from within-group variation in the $\widetilde{\Delta}$ matrix. The within-group imperfection term $(1 - \bar{P}_j^n)$ corresponds to $\rho_w < 1$: when within-group taste correlations fall short of 1, the IV group fails to preserve within-market rankings of the instrument. The variance contamination channel $(\sigma_{\Gamma,w}/\sigma_{M,w})$ is governed by the relative size of the centered correlation under Γ versus under M, since coarser groups inflate the denominator $\sigma_{M,w}$. Finally, across-group leakage ($\Delta^A \neq 0$) corresponds to $\kappa_w \neq 0$: taste correlations that cut across IV groups induce correlation between the residualized instrument and market-level unobservables.

The key additional insight from the general framework is that these three channels are unified by a single matrix $\Delta = P - P^G$, rather than being separate econometric objects. This has two practical advantages. First, it provides a diagnostic for IV group choice: given an estimate (or assumption) about P, one can compute $\mathrm{tr}(S\,\widetilde{G}\,\widetilde{\Delta})$ for candidate groupings and choose the one that minimizes this quantity. Second, it enables a bias correction: given a point estimate $\hat{P}$ from the mixed logit and the known grouping structure M, one can compute the leading-order bias from equation (39) and subtract it from $\hat{\beta}$.

3.6.9 Implications for the scale of bias

The general formula clarifies when the misclassification bias is quantitatively important. Table 1 summarizes the scaling of the bias under different degrees of alignment between the IV groups and the taste structure.

The table illustrates a sharp contrast between the second and fourth rows. When IV groups are aligned with the taste structure (even imperfectly), the bias scales with the within-cluster imperfection $(1-\rho_w)$ and can be made small by choosing groups that cluster close substitutes. When groups are orthogonal to the taste structure—as is plausible when administrative

Table 1: Scaling of misclassification bias under alternative group–taste alignments

Group–taste alignment	**Bias magnitude**	**Scaling**
Groups = true nests ($\Delta = 0$)	Zero	—
Groups $\subset$ nests (Δ^W small, $\Delta^A = 0$)	Small	$O\big(\sigma^2(1-\rho_w)/N_n\big)$
Groups aligned with taste clusters	Moderate	$O\big(\sigma^2(1-\rho_w)\big)$
Groups orthogonal to taste structure	Large	$O\big(\sigma^2(1-\rho_a)\big)$

industry×region cells fail to capture the true dimensions of worker substitutability—the bias scales with $(1-\rho_a)$, where ρ_a is the average across-cluster correlation. In the extreme case where groups are completely uninformative ($\rho_a \approx 0$), the bias approaches $\beta\sigma^2$, which does not vanish as the number of firms grows. This is consistent with the simulation evidence in Section 3.5 and Table 3, where administrative market definitions that are nearly orthogonal to the network-based markets (ARI ≈ 0) produce severely attenuated elasticity estimates, while the attenuation disappears when the correct (network-based) market definitions are used.

4 Defining Labor Markets Using Network Theory

Traditional approaches to labor market definition typically rely on administrative or geographic cells, such as industry–region or occupation–region partitions. While these variables undoubtedly capture some dimensions of worker skills, they may fail to capture the full set of latent factors (including skills, preferences, and information or other frictions) that govern workers' labor supply to jobs. In the Brazilian economy, Fogel and Modenesi (2022) show that approximately 40% of all worker-job transitions involve shifts between industries or even between 1- or 2-digit occupation codes in the short run, undermining the assumption that administrative variables convincingly reflect fixed structural parameters of labor markets. The authors show that workers within the same 4-digit occupation frequently compete for distinct sets of jobs demanding different tasks. Conversely, the same set of jobs often hire workers from a wide array of occupations. These facts indicate that worker skills and local labor markets are imperfectly proxied by administrative variables, and market definitions based on them may introduce biases into the estimation of labor market power.

Building on the insights and results from Fogel and Modenesi (2022), we adopt a data-driven approach to defining labor markets that better captures the empirical realities of worker-job mobility and competition. We leverage the intuition from Fogel and Modenesi (2022) that workers who frequently transition between particular jobs share similar skills,

and that jobs attracting similar workers require similar tasks, extending their methodology to explicitly incorporate imperfect labor market competition. Specifically, we replace their multinomial choice problem in which workers choose markets according to comparative advantage, with a nested logit model in which workers choose markets and jobs, thereby generating upward-sloping labor supply to individual jobs. This allows us to group jobs into markets and workers into skill and preference-based "worker types." While we do not currently incorporate ex ante worker heterogeneity, the worker types ι will be our definition of worker heterogeneity when we extend the model in subsequent drafts to incorporate the notion that markets that tend to hire from similar worker types are "closer", and thus are more affected by each others' labor demand choices.

4.1 Model and Estimation

Workers of skill type ι choose among jobs j nested within labor markets γ. We assume that the wage paid to worker i of type ι is the product of a worker type-market productivity term $\psi_{\iota\gamma}$ and a job j-specific wage per efficiency unit of labor. The indirect utility $V_{i\iota j\gamma}$ for worker i of type ι choosing job j within market γ incorporates worker heterogeneity through efficiency units of labor $\psi_{\iota\gamma}$, following the structure defined previously in equation (1):

$$V_{i\iota j\gamma} = \log(w_j \psi_{\iota\gamma}) + \xi_{ij\gamma}. \tag{18}$$

where the idiosyncratic error term $\xi_{ij\gamma}$ follows a nested logit distribution as specified in (2), which constitutes our market definition Assumption 1. The nested logit worker choice can be equivalently represented as a two-step decision: selecting a market and then a job within that market. To connect this labor supply framework to network theory models, we rearrange the nested logit probability into market-dependent (structural term) and job-dependent

components (observation-heterogeneity term):

$$\mathbb{P}(j \in \gamma | i \in \iota) = \underbrace{\frac{\left[\sum_{j\in\gamma}(\psi_{\iota\gamma} w_{j\gamma})^{1+\eta}\right]^{\frac{1+\theta}{1+\eta}}}{\sum_{\gamma'}\left[\sum_{j\in\gamma'}(\psi_{\iota\gamma} w_{j\gamma})^{1+\eta}\right]^{\frac{1+\theta}{1+\eta}}}}_{\mathbb{P}(\gamma|\iota)} \underbrace{\frac{(\psi_{\iota\gamma} w_{j\gamma})^{1+\eta}}{\sum_{j\in\gamma}(\psi_{\iota\gamma} w_{j\gamma})^{1+\eta}}}_{\mathbb{P}(j|\iota,\gamma)}$$

1st step: market choice — **2nd step**: job choice

$$= \underbrace{\frac{\left[\sum_{j\in\gamma}(\psi_{\iota\gamma} w_{j\gamma})^{1+\eta}\right]^{\frac{1+\theta}{1+\eta}-1}}{\sum_{\gamma'}\left[\sum_{j\in\gamma'}(\psi_{\iota\gamma} w_{j\gamma})^{1+\eta}\right]^{\frac{1+\theta}{1+\eta}}} \psi_{\iota\gamma}^{1+\eta}}_{=:\,\Omega_{\iota\gamma}} \underbrace{w_{j\gamma}^{1+\eta}}_{=:\,d_j} \tag{19}$$

Market component — Job component

The choice probability in (42) regards a single worker choice. In reality, however, we observe matches being dissolved and new matches being made. We formalize repeated job searches resulting from exogenous separation shocks, arriving following a Poisson process at rate d_i for worker i, making this our market definition Assumption 2.[5]. This makes the expected number of matches between worker i and job j over our sample period equal to:

$$d_i \cdot P(j \in \gamma | i \in \iota) = \Omega_{\iota\gamma} d_i d_j. \tag{20}$$

Let A_{ij} be the observed number of matches between worker i and job j, as an element of the matrix $\boldsymbol{A}$ of dimensions $N_W \times N_J$, containing information regarding all worker-job matches in the economy over a period of time. Let worker latent types and job latent markets be encoded in the latent membership vector $\boldsymbol{b}$. Following equation (43), the expected number of matches given types and markets is:

$$E[A_{ij}|\boldsymbol{b}] = \Omega_{\iota\gamma} d_i d_j,$$

Together, market definition Assumptions 1 and 2 imply (proof in B):

$$A_{ij}|\boldsymbol{b} \sim \text{Poisson}(\Omega_{\iota\gamma} d_i d_j). \tag{21}$$

[5]Despite being reliant on a Poisson separation process, its parameter varies by individual worker i, a common assumption in flexible namely non-parametric network theory, give it's flexibility (Peixoto, 2018). Furthermore, the Poisson assumption generates a network with a degree distribution that follows a power-law, suited to represent large real world sparse networks, like the labor market, which follows a power-law (Newman, 2018).

This framework above corresponds exactly to the bipartite degree-corrected Stochastic Block Model (dc-BiSBM), from network theory, where clustering into worker types and job markets is based on empirical match patterns. In principle, we could utilize the labor market network data in A_{ij} to choose the latent labor market partition $\boldsymbol{b}$ by maximizing the likelihood in (45). In practice, this is a combinatorial NP-hard optimization problem, so instead we estimate the latent types and markets by maximizing the Bayesian posterior[6]:

$$P(\boldsymbol{b}|A_{ij}) \propto P(A_{ij}|\boldsymbol{b})P(\boldsymbol{b}), \tag{22}$$

while selecting the optimal number of worker types and markets using the minimum description length (MDL) principle, which works as a classic Bayesian penalization term. In our case, it allows for more markets and types, i.e. model complexity, only if the model better explains the matches labor market network, i.e. if it comes with greater fit (for more details, see Fogel and Modenesi (2022)). Computational feasibility is ensured through Markov chain Monte Carlo (MCMC) methods, implemented efficiently using the *graph-tool* Python library[7].

4.2 Data

We utilize Brazil's RAIS (Relação Anual de Informações Sociais) dataset, an administrative employer–employee dataset that covers the universe of formally employed workers in Brazil. RAIS is essentially a census of formal jobs, collected annually by the Ministry of Labor. We use RAIS data from 1986-1990 to compute our estimates of network-based labor markets.

To define network-based labor markets by estimating the stochastic block model (SBM), we identify the giant component of the set of worker–job matches during the period immediately preceding the trade shock: 1986 to 1990. We define a job as an establishment- 4 digit occupation pair. For computational reasons, we restrict estimation to the three largest states in Brazil (which are contiguous): São Paulo, Rio de Janeiro, and Minas Gerais. We identify 3,737 unique markets.

Table 2 compares the size distributions of network-based and administrative (Occ2×Micro Region) labor markets within this three-state sample. Administrative markets are far more numerous (10,074 vs. 3,737) but far smaller on average, with a median employment of just 76 workers compared to 2,339 for network-based markets. The administrative markets also exhibit extreme skewness, with a standard deviation of employment (9,024) vastly exceeding the mean (1,037), reflecting many tiny markets. By contrast, the network-based

[6]The prior distributions $P(\boldsymbol{b})$ and $P(\Omega_{\iota\gamma})$, which was not shown above for simplicity, are chosen as maximum entropy priors, i.e., the most agnostic in terms the distribution of the sizes of markets and worker types.

[7]Source: `https://graph-tool.skewed.de/`. See Peixoto (2014) for details.

markets display a much more compact distribution. These patterns underscore the structural differences between the two approaches to market definition.

Table 2: Size distribution of labor markets: network-based (γ) vs. administrative (Occ2×Micro Region), three-state sample

	γ	Occ2 × Micro
Num. Markets	3,737	10,074
Mean	2,490	1,037
Std	2,485	9,024
Min	2	1
P25	1,944	17
P50	2,339	76
P95	2,804	337
Max	98,844	414,318

4.3 Initial Evidence on the Importance of Misclassification

To gauge the possible impact of misclassification we conduct an additional simulation exercise using RAIS data and our network-based estimates of γ. This exercise differs from the simulation and estimation procedures described in Section 3.5 in a number of ways. First, we set θ to 0.42 and η to 10.85, calibrated to match the estimates of Berger et al. (2022). Second, instead of estimating η and θ using cross-sectional variation for a single year, we estimate the parameters by regressing changes in (log) labor supply on changes in (log) wages, where we use real data for the "pre" year and simulated data for the "post" year. Additionally, we estimate the parameters via IV, rather than OLS.

In order to compute the simulation, we follow the following steps. First, we plug in the assumed values of η and θ to construct $\text{MRPL}^0_{j\gamma}$, the model-implied marginal revenue product in the base year, 1991. Then, we generate a series of additive firm- and market-level demand shocks, which we use to simulate a demand shock to each firm:

$$\log\left(\text{MRPL}^1_{j\gamma}\right) = \log\left(\text{MRPL}^0_{j\gamma}\right) + \pi Z_j + \zeta_\gamma + \zeta_j,$$

where

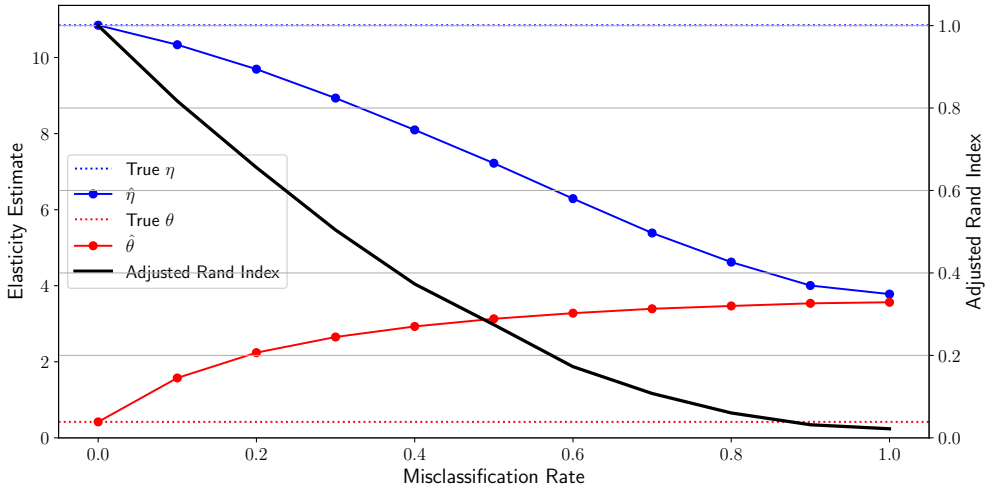
$$\zeta_\gamma \overset{iid}{\sim} N(0, \sigma_b^2), \quad \zeta_j \overset{iid}{\sim} N(0, \sigma_w^2), \quad Z_j \overset{iid}{\sim} \text{Bernoulli}(p_\gamma) \quad p_\gamma \overset{iid}{\sim} \text{Beta}(\alpha, \beta).$$

In this simulation exercise, Z_j stands in for a "realistic" instrumental variable—it is a component of the overall demand shock that is both observable and orthogonal to other factors influencing wages and employment. Variation in this instrument contains an important

between-market component: we first draw a market-specific shock probability p_γ from a Beta distribution, then draw firm-specific shocks according to the p_γ of each firms' market γ. Both the coefficient π and the parameters of the Beta(α, β) distribution determine the strength of the instrument, at both the within-market and between-market levels. Given a set of post-shock productivity measures ($\text{MRPL}^1_{j\gamma}$), we compute the post-shock equilibrium using the iterative procedure described in Section 3.5. Let Δx for any variable x denote the difference between the simulated post-shock value of the variable x^1 and the pre-shock value x^0 taken from the 1991 RAIS data. For each labor market definition $d \in D$, we estimate η by regressing $\Delta \ell_j$ on Δw_j and market d fixed effects, instrumenting Δw_j with Z_j. As before, we use the estimate $\hat{\eta}$ to construct $\hat{\omega}^0_d$ and $\hat{\omega}^1_d$. Finally, we estimate $\hat{\theta}$ by regressing $\Delta \ell_d$ on $\Delta \omega_d$, instrumenting $\Delta \omega_d$ with $\overline{Z}_d$, the average of Z_j across firms within each market d.

Figure 2 shows the first set of results of this simulation exercise, reporting estimated values of η and θ at progressively higher rates of random re-assignment of market labels. The true elasticity gap is $\eta - \theta = 10.85 - 0.42 = 10.43$. Like the initial simulation exercises, the results of this exercise indicate that bias can be substantial, even at relatively low rates of (random) misclassification. The Adjusted Rand Index (ARI) between the network-based γ definition and the Occ2×Region definition is just 0.007, indicating that these two market definitions are nearly orthogonal and that administrative markets are highly misclassified relative to network-based markets.

Figure 2: Estimated $\hat{\eta}$ and $\hat{\theta}$ under varying rates of random misclassification



While the results reported in Figure 2 confirm those of the prior simulation exercises described in Section 3.5, our simulation using RAIS data allows for further exploration of the empirical relevance of market definition bias. In particular, none of the simulation evidence

presented so far speaks to realistic patterns of misclassification—the exercises all rely either on completely random perturbation of market labels, or mechanical aggregation/sub-division of markets. In order to understand whether the patterns of bias suggested by these exercises are informative for understanding the likely biases of existing estimates of labor supply parameters, we repeat our estimation procedure using a broad range of commonly-used labor market definitions based on occupation (at 2-, 4-, and 5-digit levels), micro region, and 5-digit industry, both alone and interacted with geography. These estimates are reported in Table 3, alongside the Adjusted Rand Index (ARI) measuring the similarity of each definition to both γ and to the Occ2×Region definition. The first row of Table 3 reports baseline estimates using the "true" labor market definition γ, confirming that our estimation procedure recovers unbiased estimates of the structural elasticity parameters ($\hat{\eta} = 10.832 \approx \eta$, $\hat{\theta} = 0.454 \approx \theta$). The second row re-estimates using the Occ2×Micro Region definition; the estimates are severely attenuated: $\hat{\eta}$ shrinks to 5.03 and $\hat{\theta}$ inflates to 3.92, attenuating the true elasticity gap of 10.38 by approximately 89%. The remaining rows report estimates using other administrative definitions, all of which exhibit similarly severe attenuation. Notably, the ARI between γ and all administrative definitions is very close to zero (0.001–0.022), confirming that these definitions are nearly orthogonal to the network-based markets. The results reported in Table 3 indicate that, if anything, the practical impacts of misclassification bias may be more severe than indicated by many of our simulation results.

Table 3: Simulated elasticity estimates under alternative market definitions. DGP calibrated to $\eta = 10.85$, $\theta = 0.42$ (Berger et al., 2022). ARI columns report the Adjusted Rand Index measuring similarity to the true network-based definition (γ) and to the Occ2×Region definition.

Market Definition	$\hat{\eta}$	$\hat{\theta}$	ARI (γ)	ARI (Occ2×Region)
γ	10.832	0.454	1.000	0.007
Occupation (2-digit) × Micro Region	5.028	3.916	0.007	1.000
Occupation (2-digit)	4.799	2.870	0.001	0.241
Micro Region	4.844	3.444	0.001	0.074
Occupation (4-digit) × Micro Region	5.175	4.197	0.018	0.349
Occupation (4-digit)	4.946	3.776	0.004	0.159
Occupation (5-digit) × Micro Region	5.188	4.209	0.019	0.322
Occupation (5-digit)	4.954	3.812	0.004	0.153
Industry (5-digit) × Micro Region	5.389	3.996	0.022	0.108
Industry (5-digit)	5.055	3.344	0.005	0.040

5 Empirical Application: Brazilian Labor Market and Trade Shock

In section 3.1, we outlined a model-derived approach to estimating the labor supply elasticity parameters η and θ under the naive assumption that variation in wages is orthogonal to unobserved determinants of employment levels. Essentially, this strategy consisted of (i) identifying the within-market elasticity η by regressing firm-level employment on firm-level wages controlling for market fixed effects (equation 18), and (ii) identifying the between-market elasticity θ by regressing market-level employment on a market-level wage aggregate (equation 19).

In this section we follow the instrumental variables strategy of Felix (2022)[8], which leverages Brazil's trade liberalization from 1990 to 1994 as a source of plausibly exogenous variation in labor demand between firms and markets. In 1990, the Brazilian government announced a major reform reducing import tariffs across a wide range of industries, with the goal of being fully implemented by 1994. This reform was unanticipated and driven by external policy factors (Brazil's trade opening), not by conditions in any specific labor market. It generated a substantial shock: industries that had high pre-liberalization tariffs experienced large tariff cuts and consequently faced increased import competition, leading to contractions and even exit of many firms. At the same time, less protected industries saw smaller tariff reductions and hence milder impacts. Following Felix (2022)'s empirical strategy, we use firm-level variation in tariff exposure as an instrument for firm-level wages, and market-level tariff exposure as an instrument for market-level wages. While the expressions we derived in our prior analysis of market definition bias involved estimates of the labor supply curve $\ell_{j\gamma}(w)$, the empirical strategy outlined below involves estimating inverse labor supply curves $w_{j\gamma}(\ell)$. While this difference is important, estimates of between- and within-market inverse labor supply elasticities are subject to nearly identical biases due to misclassification.

5.1 Identification and Estimation of η

Felix (2022) shows that inverse labor supply to job j in market γ at time t can be written

$$\ln w_{j\gamma t} = \frac{1}{\eta} \ln \ell_{j\gamma t} + \delta_{\gamma t} + \epsilon_{j\gamma t}$$

where $\delta_{\gamma t}$ are market×year fixed effects and $\varepsilon_{j\gamma t} = \ln \xi_{j\gamma t}^{1+\eta}$ is the regression residual. Taking

[8] Felix (2022)'s identification builds upon Dix-Carneiro and Kovak (2017).

1991 to 1997 long differences this becomes the key estimating equation for $\frac{1}{\eta}$:

$$\text{[Second Stage]} \quad \Delta \ln w_{j\gamma} = \frac{1}{\eta} \Delta \ln \ell_{j\gamma} + \Delta \delta_{\gamma} + \Delta \epsilon_{j\gamma} \tag{23}$$

where $\Delta\delta_\gamma$ is a market fixed effect in the differenced regression. The fundamental threat to identification is that changes in labor supply to firm j in market γ ($\Delta \ln \ell_{j\gamma}$) may be correlated with changes in workers' idiosyncratic preferences for job j in market γ ($\Delta\epsilon_{j\gamma}$). Felix addresses this concern by instrumenting for $\Delta \ln \ell_{j\gamma}$ with the labor demand shock created by the change in import tariffs on job j's output sector, $\Delta \ln(1 + \tau_{s(j)})$, using the following first stage regression:

$$\text{[First Stage]} \quad \Delta \ln \ell_{j\gamma} = \lambda \, \Delta \ln\big(1 + \tau_{s(j)}\big) + \Delta d_{\gamma} + \nu_{j\gamma} \tag{24}$$

where Δd_γ is again a market fixed effect.

Estimating equations 48 and 49 requires measuring three model objects: the total units of labor $\ell_{j\gamma t}$ supplied to firm j in market γ at year t, the wage $w_{j\gamma t}$ paid by that firm–market pair, and the tariff shock to the firm. Total units of labor at firm j in year t, $\ell_{j\gamma t}$, are measured as the total number of workers employed at firm j in market γ during the entire month of December of year t. The wage at job j in market γ in time t, $w_{j\gamma t}$, is measured as firm j's *wage premium* in market γ from a regression of worker log December earnings on job $\times$ market fixed effects plus controls for worker characteristics, including age, education, and gender. Finally, the tariff shock to job j is the policy-induced change in import tariffs in job j's output sector:

$$\Delta \ln\big(1 + \tau_{s(j)}\big) \equiv - \ln\left(\frac{1 + \tau_{s(j),1994}}{1 + \tau_{s(j),1990}}\right).$$

5.2 Identification and Estimation of θ

Identification of θ follows analogously at the market-level. Felix (2022) derives the following estimating equation for the gap between workers' key elasticities of substitution $\left(\frac{1}{\theta} - \frac{1}{\eta}\right)$, from which $\frac{1}{\theta}$ can be backed out:

$$\text{[Second Stage]} \quad \Delta\delta_{\gamma} = \alpha + \left(\frac{1}{\theta} - \frac{1}{\eta}\right) \Delta \ln L_{\gamma} + \Delta\varepsilon_{\gamma} \tag{25}$$

where the constant α absorbs country-level wage component changes common to all markets,[9] $\Delta \ln L_\gamma$ is the change in the market-level CES labor supply index described below, and $\Delta\epsilon_\gamma = \Delta \ln \xi_{j\gamma}^{1+\theta}$ is the market-level regression residual, which also has a structural interpretation as the (scaled) change in workers' idiosyncratic preferences for market γ.

The fundamental threat to identification of $\left(\frac{1}{\theta} - \frac{1}{\eta}\right)$ in regression equation 51 is that changes in the taste-adjusted labor supplied to market γ ($\Delta \ln L_\gamma$) are correlated with changes in workers' idiosyncratic preferences for market γ ($\Delta\varepsilon_\gamma = \Delta \ln \xi_{j\gamma}^{1+\theta}$). Felix addresses this concern by instrumenting for the market-level change in labor supply with ΔICE_γ, the market-level policy-induced import competition exposure shock. ΔICE_γ is measured as a weighted average of job-level tariff changes for each market γ. The market-level first stage regression is thus:

$$\text{[First Stage]} \quad \Delta \ln L_\gamma = \tilde{\alpha} + \lambda \Delta ICE_\gamma + \Delta\nu_\gamma \tag{26}$$

where $\tilde{\alpha}$ is a constant, and $\Delta\nu_\gamma$ is a regression residual. The between-market inverse elasticity $\frac{1}{\theta}$ is then backed out by summing the estimate $\left(\frac{1}{\theta} - \frac{1}{\eta}\right)$ from equation 51 and the estimate of $\frac{1}{\eta}$ from equation 48.

To estimate equations 51 and 52, we need to measure three objects: $\Delta\delta_\gamma$, the market-level component of the firm-level wage, and $\Delta \ln L_\gamma$, the market-level change in the CES labor supply index; and the market-level import competition exposure measure ΔICE_γ. We measure $\Delta\delta_\gamma$ as the market fixed effect from regression equation 48 and compute $\Delta \ln L_\gamma$ given our point-estimate for $\frac{1}{\eta}$ as follows:

$$\Delta \ln L_\gamma = \Delta \ln \left\{ \left[\sum_{\gamma \in \Gamma} \left(\xi_{j\gamma} / \ell_{j\gamma} \right)^{\frac{1+\eta}{\eta}} \right]^{\frac{\eta}{1+\eta}} \right\} \tag{27}$$

where the taste-shifters $\xi_{j\gamma}$ are calculated as $\xi_{j\gamma t} = (1+\eta)\exp(\nu_{j\gamma t})$, where $\nu_{j\gamma t}$ are the residuals from a regression of $[\ln w_{j\gamma t} - (1/\eta)\ln \ell_{j\gamma t}]$ on a market fixed effect and we plug in the estimate $\frac{1}{\eta}$ from above.

5.3 Results

In this section, we implement the empirical strategy described above, using two-stage least squares (2SLS) to estimate the labor supply elasticities $\frac{1}{\eta}$ and $\frac{1}{\theta}$ under various market definitions. Our central empirical objective is to compare estimates produced using traditional

[9] (i.e., $\alpha = \frac{1}{\theta}\Delta \ln \left(\frac{L}{\ell}\right) + \Delta \ln W$).

administrative labor markets to those generated using network-based labor markets identified via the stochastic block model. We use RAIS data from 1991 and 1997 to compute long differences in employment and wages when estimating the labor supply parameters. For the purposes of calculating market-level average markdowns, we calculate payroll HHIs using data from 1997. We follow Felix (2022) as closely as possible and restrict to private sector workers aged 18 to 65. Following Felix (2022), we define administrative labor markets as two-digit occupation by microregion cells, while adopting the network-based definitions—denoted γ—estimated using RAIS data from 1986-1990.

Table 5 presents our elasticity estimates. Column 1 presents our estimates using the full RAIS sample under administrative (Occ2×Micro Region) market definitions. Column 2 repeats this exercise, restricting the sample to the three largest Brazilian states—São Paulo, Rio de Janeiro, and Minas Gerais—which are the same states used in the stochastic block model analysis. The estimates remain nearly unchanged, indicating that our state restriction does not substantively alter the results under administrative market definitions. Column 3 implements our main specification, using the same three-state sample but replacing the administrative market definitions with the network-based markets.

The contrast between columns 2 and 3 is striking. When using the network-based markets, we estimate a slightly higher value of $\hat{\eta}$ and a substantially lower value of $\hat{\theta}$. This pattern aligns closely with the predictions from our model of misclassification bias: if administrative markets contain greater misclassification than network-based markets, we would expect to see $\hat{\eta}$ biased downward and $\hat{\theta}$ biased upward under the administrative definitions. The larger value of $\hat{\eta}$ and smaller value of $\hat{\theta}$ under the γ specification suggest that the network-based markets exhibit less misclassification and are thus better aligned with the theoretical structure of the model. We interpret this as evidence that the stochastic block model yields labor market definitions that more accurately reflect true patterns of worker substitution and job competition.

As shown in equation (15), the resulting markdowns μ_γ depend not only on η and θ but also on market concentration—specifically, the payroll HHI of each market. This corresponds to what we earlier referred to as the "denominator problem": under our Bertrand model, the labor supply elasticity facing each job is a weighted average of η and θ, where the weights are determined by the job's share of its market. As a result, the estimated markdowns depend critically on both the estimated elasticities and the distribution of HHIs across markets.

Figures 3 and 4 present the distribution of market-level HHIs under the three market definitions in Table 5. In the unweighted distribution (Figure 3), the network-based markets (blue) exhibit a broader and less concentrated set of HHIs, with few markets exhibiting perfect concentration. By contrast, the administrative markets—particularly in the full Brazil sample

(green)—contain a large mass of markets with an HHI of one, which are typically very small markets containing only a handful of firms or even a single firm. These highly concentrated but tiny markets are of limited economic significance. When we weight the HHI distributions by market-level employment (Figure 4), the pattern reverses. The network-based markets exhibit a rightward shift in the distribution, indicating greater concentration in economically important markets. This suggests that while administrative markets include many highly concentrated but economically trivial groups, the network-based definitions identify a smaller number of concentrated markets that carry greater aggregate weight in the economy.

Figures 5 and 6 present the resulting distributions of wage markdowns under the various market definitions, computed following equation 15, plugging the estimated inverse elasticities and HHIs. The differences between the administrative and network-based definitions are striking. In the unweighted distribution (Figure 5), the administrative market definitions produce a narrow and compressed distribution of markdowns. This is expected: since the markdown depends on the elasticity of labor supply to the firm, and this elasticity is bounded between $\frac{1}{\eta}$ and $\frac{1}{\theta}$, a smaller gap between $\widehat{\frac{1}{\eta}}$ and $\widehat{\frac{1}{\theta}}$ mechanically limits the possible range of markdowns. By contrast, the network-based markets exhibit a wider dispersion of markdowns, reflecting both a larger spread between $\widehat{\frac{1}{\eta}}$ and $\widehat{\frac{1}{\theta}}$ and a richer distribution of HHIs. The difference becomes even more pronounced in the weighted distribution (Figure 6). Here, the administrative definitions yield very little variation in markdowns across economically significant markets. The maximum markdowns are modest, and most of the mass is clustered in a narrow band. By contrast, the network-based markets produce a broader distribution with heavier mass in the upper tail—indicating more markets with substantial monopsony power.

These findings underscore the importance of market classification in studies of labor market power. Misclassification biases both the estimated elasticities and the market concentration measures that determine markdowns. In our empirical application, the effects on elasticities appear to be the more consequential of the two, with the administrative definitions yielding both larger $\widehat{\frac{1}{\eta}}$ and smaller $\widehat{\frac{1}{\theta}}$, thereby understating the scope for wage-setting power. The use of administrative market definitions leads to a systematic understatement of markdowns and a severe compression in their distribution, which in turn may lead to misleading policy conclusions about the prevalence and magnitude of monopsony in the labor market.

Table 4: Labor supply elasticity estimates ($\hat{\eta}$ and $\hat{\theta}$) using IV

Market Definition Sample	Admin Full (1)	Admin 3 State (2)	Network 3 State (3)
$\hat{\eta}$	0.712	0.709	0.811
	(0.078)	(0.079)	(0.063)
[First stage F]	91.7	87.2	211.7
$\hat{\theta}$	0.548	0.512	0.264
	(0.339)	(0.356)	(0.418)
[First stage F]	2.4	2.3	0.2
Observations	961,370	571,923	463,671
Firms	368,770	223,138	174,816
Markets	19,125	7,927	3,727

3 State: São Paulo, Rio de Janeiro, and Minas Gerais

"Administrative" labor markets: 2-digit occupation × micro region (558)

η and θ regressions weighted by job size and market size, respectively.

Figure 3: Distribution of HHIs by market definition (unweighted)

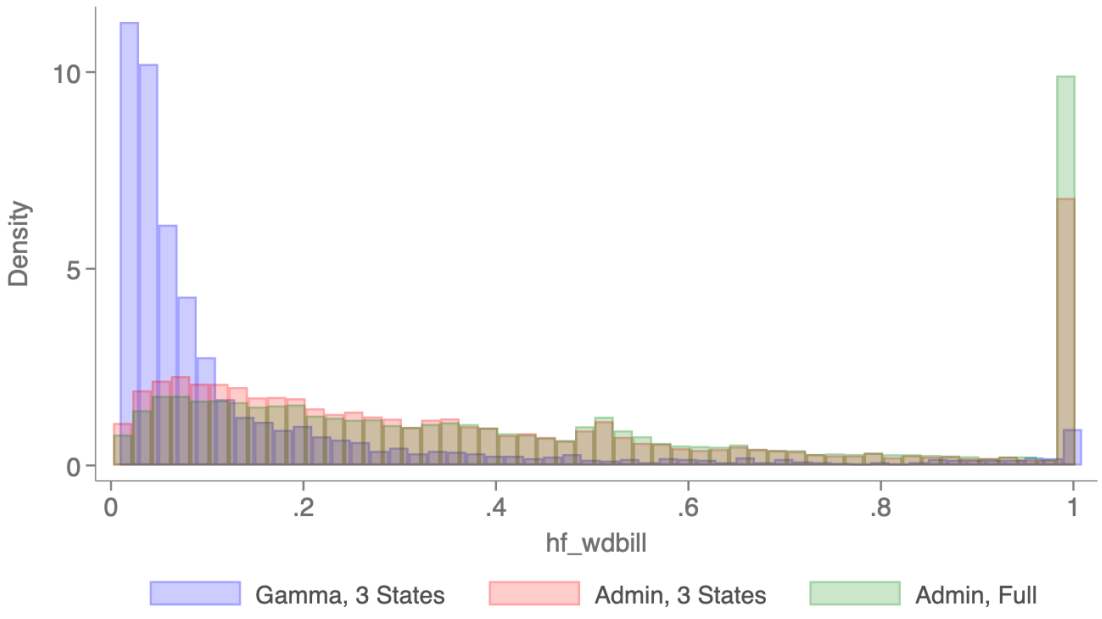


Figure 4: Distribution of HHIs by market definition (weighted by market employment)

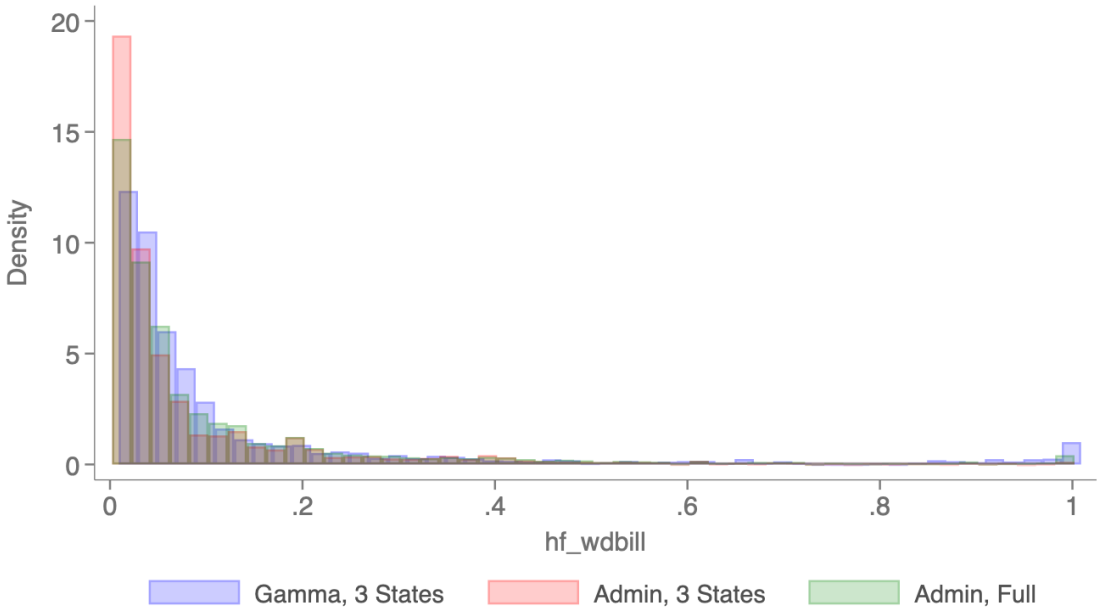


Figure 5: Distribution of wage markdowns by market definition (unweighted, linear scale)

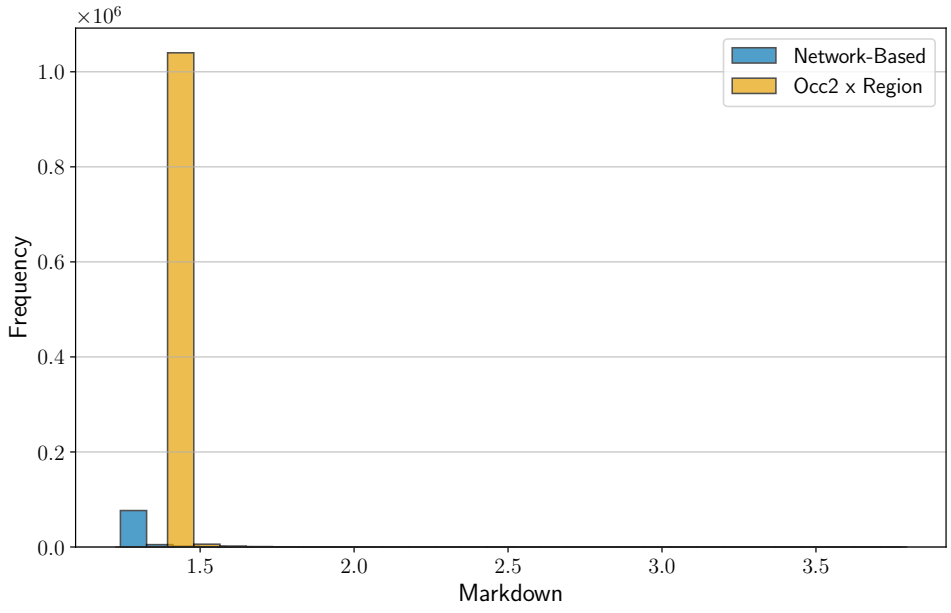


Figure 6: Distribution of wage markdowns by market definition (weighted by job payroll, linear scale)

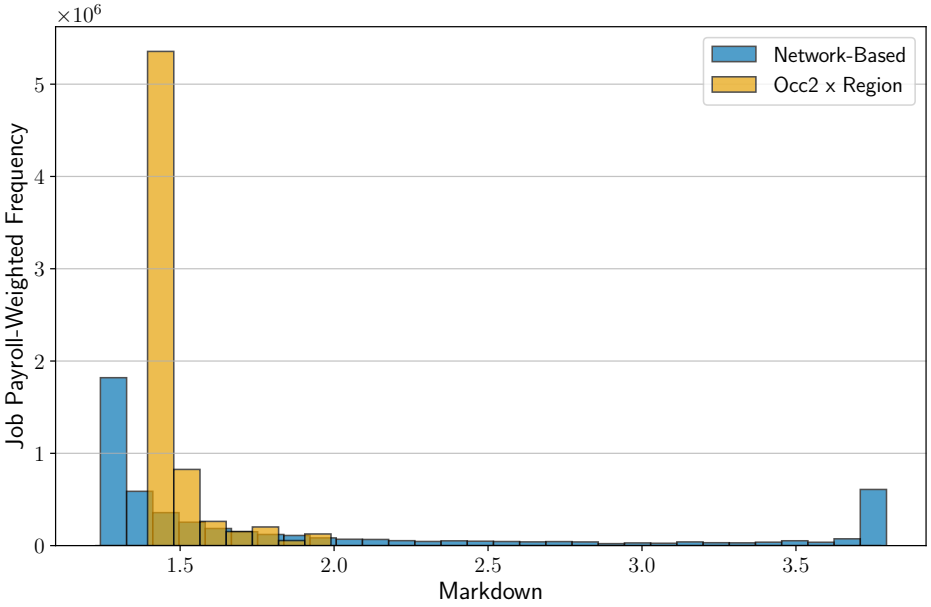
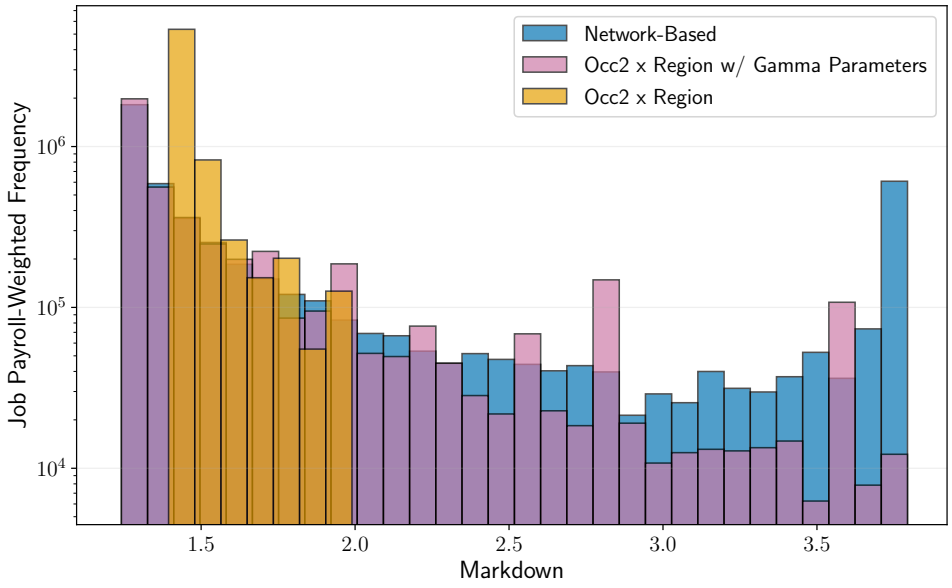


Figure 7: Markdown counterfactuals (weighted by job payroll, log scale)



References

Azar, Jose A., Ioana Marinescu, Marshall I. Steinbaum, and Bledi Taska, "Concentration in US Labor Markets: Evidence From Online Vacancy Data," Working Paper 24395, National Bureau of Economic Research March 2018.

Azar, Jose, Ioana Marinescu, and Marshall Steinbaum, "Measuring Labor Market Power Two Ways," 2019.

Azar, José, Steven T Berry, and Ioana Marinescu, "Estimating Labor Market Power," NBER Working Paper 30365 August 2022.

Baker, Jonathan B and Timothy F Bresnahan, "Estimating the residual demand curve facing a single firm," *International Journal of Industrial Organization*, 1988, *6* (3), 283–300.

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

_ , _ , Andreas R Kostøl, and Simon Mongey, "An anatomy of monopsony: Search frictions, amenities, and bargaining in concentrated markets," *NBER Macroeconomics Annual*, 2024, *38* (1), 1–47.

Card, David, Ana Rute Cardoso, Joerg Heining, and Patrick Kline, "Firms and Labor Market Inequality: Evidence and Some Theory," *Journal of Labor Economics*, 2018, *36* (1), 13–70.

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

Dix-Carneiro, Rafael and Brian K Kovak, "Trade liberalization and regional dynamics," *American Economic Review*, 2017, *107* (10), 2908–2946.

Farrell, Joseph and Carl Shapiro, "Antitrust evaluation of horizontal mergers: An economic alternative to market definition," *The BE Journal of Theoretical Economics*, 2010, *10* (1), 000010220219351704 1563.

Felix, Mayara, "Trade, Labor Market Concentration, and Wages," 2022.

Fogel, Jamie and Bernardo Modenesi, "What is a Labor Market? Classifying Workers and Jobs Using Network Theory," 2022.

Harris, Barry C and Joseph J Simons, "Focusing market definition: How much substitution is necessary," *J. Reprints Antitrust L. & Econ.*, 1991, *21*, 151.

Hovenkamp, Herbert and Carl Shapiro, "Horizontal mergers, market structure, and burdens of proof," *The Yale Law Journal*, 2018, pp. 1996–2025.

Jarosch, Gregor, Jan Sebastian Nimczik, and Isaac Sorkin, "Granular Search, Market Structure, and Wages," *The Review of Economic Studies*, 01 2024, *91* (6), 3569–3607.

Kaplow, Louis, "Why (ever) define markets," *Harv. L. Rev.*, 2010, *124*, 437.

_ , "Market definition and the merger guidelines," *Review of Industrial Organization*, 2011, *39*, 107–125.

Manning, Alan, *Monopsony in Motion*, Princeton University Press, 2005.

Miller, Nathan H and Matthew C Weinberg, "Understanding the price effects of the MillerCoors joint venture," *Econometrica*, 2017, *85* (6), 1763–1791.

Nevo, Aviv, "Measuring market power in the ready-to-eat cereal industry," *Econometrica*, 2001, *69* (2), 307–342.

Newman, M., *Networks*, OUP Oxford, 2018.

Nimczik, Jan Sebastian, "Job Mobility Networks and Endogenous Labor Markets," 2018.

Peixoto, Tiago P, "Efficient Monte Carlo and greedy heuristic for the inference of stochastic block models," *Physical Review E*, 2014, *89* (1), 012804.

Peixoto, Tiago P., "Nonparametric weighted stochastic block models," *Phys. Rev. E*, Jan 2018, *97*, 012306.

Roussille, Nina and Benjamin Scuderi, "Bidding for Talent: A Test of Conduct in a High-Wage Labor Market," 2024.

Schubert, Gregor, Anna Stansbury, and Bledi Taska, "Employer Concentration and Outside Options," Working Paper 2022.

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

A Proof of Theorem 1

We prove the OLS case ($z_j = w_j$). The IV extension follows the same algebra with wages replaced by first-stage fitted values; the sign conditions $\text{sign}(\pi) = \text{sign}(\pi + \kappa_w) = \text{sign}(\pi + \kappa_b)$ ensure the inequalities are preserved.

Notation. Let N denote the number of jobs. Collect observations in the vector $z \in \mathbb{R}^N$. Let D_Γ and D_M be the $N \times |\Gamma|$ and $N \times |M|$ matrices of market indicators for the true and mismeasured definitions, respectively, so that $[D_\Gamma]_{j\gamma} = \mathbf{1}\{j \in \gamma\}$ and similarly for D_M. Define the orthogonal projection matrices

$$P_\Gamma = D_\Gamma (D_\Gamma' D_\Gamma)^{-1} D_\Gamma', \qquad P_M = D_M (D_M' D_M)^{-1} D_M'.$$

Both P_Γ and P_M are symmetric and idempotent ($P^2 = P$, $P' = P$). In vector notation, $[P_\Gamma z]_j = \bar{z}_{\gamma(j)}$ and $[(I - P_M)z]_j = z_j - \bar{z}_{m(j)}$.

Lemma (Projection orthogonality). *For any symmetric idempotent matrix P and any conformable matrix C,*

$$P\,C\,(I - P) \;=\; 0.$$

Proof. Since $PD = D$ (each column of D is in $\text{col}(P)$ by construction), we have $(I-P)D = 0$, hence $D'(I - P) = 0$. Every row of PC lies in the row space of D (because $PC = D(D'D)^{-1}D'C =: D\tilde{C}$ for some matrix $\tilde{C}$). Therefore each row of PC times $(I - P)$ equals a row of $\tilde{C}$ times $D'(I - P) = 0$. □

Part (i): $\beta_w \leq 1$.

In the OLS case, $\beta_w = \text{Cov}(w_j - \bar{w}_{\gamma(j)},\, w_j - \bar{w}_{m(j)}) / \text{Var}(w_j - \bar{w}_{m(j)})$. The condition $\beta_w \leq 1$ is equivalent to

$$\text{Var}(w_j - \bar{w}_{m(j)}) - \text{Cov}(w_j - \bar{w}_{\gamma(j)},\, w_j - \bar{w}_{m(j)}) \;\geq\; 0,$$

which simplifies (using $\text{Cov}(\bar{w}_{m(j)},\, w_j - \bar{w}_{m(j)}) = 0$) to

$$E\big[\bar{w}_{\gamma(j)}\,(w_j - \bar{w}_{m(j)})\big] \;\geq\; 0.$$

In matrix form this is $\frac{1}{N} z' P_\Gamma (I - P_M) z \geq 0$. Decompose $z = P_\Gamma z + (I - P_\Gamma)z$ on the right:

$$\frac{1}{N} z' P_\Gamma (I - P_M) z = \underbrace{\frac{1}{N} z' P_\Gamma (I - P_M) P_\Gamma z}_{=:\,A} + \underbrace{\frac{1}{N} z' P_\Gamma (I - P_M)(I - P_\Gamma) z}_{=:\,B}.$$

By the Lemma (with $P = P_\Gamma$ and $C = I - P_M$), $B = 0$. For A, since P_Γ and $(I - P_M)$ are

symmetric and $(I - P_M)$ is idempotent:

$$A = \frac{1}{N}(P_\Gamma z)'(I - P_M)(P_\Gamma z) = \frac{1}{N}(P_\Gamma z)'(I - P_M)'(I - P_M)(P_\Gamma z) = \frac{1}{N}\big\|(I - P_M)P_\Gamma z\big\|^2 \;\geq\; 0.$$

Hence $E[\bar{w}_{\gamma(j)}(w_j - \bar{w}_{m(j)})] \geq 0$, which gives $\beta_w \leq 1$ and therefore

$$\text{Bias}(\hat{\eta}) = (\theta - \eta)(1 - \beta_w) \leq 0. \qquad (\eta \geq \theta \text{ and } \beta_w \leq 1)$$

Part (ii): $\beta_b \leq 1$.

$\beta_b = \text{Cov}(\bar{z}_{\gamma(j)}, \bar{z}_{m(j)}) / \text{Var}(\bar{z}_{m(j)})$. The condition $\beta_b \leq 1$ is equivalent to

$$E\big[(\bar{z}_{m(j)} - \bar{z}_{\gamma(j)})\,\bar{z}_{m(j)}\big] \;\geq\; 0.$$

Since $E[(w_j - \bar{w}_{m(j)})\bar{w}_{m(j)}] = 0$, this is further equivalent to

$$E\big[(w_j - \bar{w}_{\gamma(j)})\,\bar{w}_{m(j)}\big] \;\geq\; 0,$$

i.e., $\frac{1}{N}z'(I - P_\Gamma)P_M z \geq 0$. Decompose $z = P_M z + (I - P_M)z$ on the left:

$$\frac{1}{N}z'(I - P_\Gamma)P_M z = \underbrace{\frac{1}{N}z'P_M(I - P_\Gamma)P_M z}_{=:\,A'} + \underbrace{\frac{1}{N}z'(I - P_M)(I - P_\Gamma)P_M z}_{=:\,B'}.$$

By the Lemma (with $P = P_M$ and $C = I - P_\Gamma$), $P_M(I - P_\Gamma)(I - P_M) = 0$; taking the transpose, $(I - P_M)(I - P_\Gamma)P_M = 0$, so $B' = 0$. For A':

$$A' = \frac{1}{N}(P_M z)'(I - P_\Gamma)(P_M z) = \frac{1}{N}\big\|(I - P_\Gamma)P_M z\big\|^2 \;\geq\; 0.$$

Hence $\beta_b \leq 1$ and therefore

$$\text{Bias}(\hat{\theta}) = (\eta - \theta)(1 - \beta_b) \geq 0. \qquad (\eta \geq \theta \text{ and } \beta_b \leq 1) \qquad \square$$

B Proof that A_{ij} follows a Poisson distribution

If an individual worker i only searched for a job once, then the probability of worker i matching with job j would be equal to $\mathbb{P}_{ij} = d_j \Omega_{\iota\gamma}$ and A_{ij} would follow a Bernoulli distribution:

$$A_{ij} \sim Bernoulli(d_j \Omega_{\iota\gamma}).$$

However, since worker i searches for jobs $c_i \equiv \sum_{t=1}^{T} c_{it}$ times, A_{ij} is actually the sum of c_i Bernoulli random variables, and is therefore a Binomial random variable. Conditional on knowing c_i,

$$A_{ij}|c_i \sim Binomial(c_i, d_j\Omega_{\iota\gamma}).$$

However, we still need to take into account the fact that c_i is a Poisson-distributed random variable with arrival rate d_i. Consequently, the unconditional distribution of A_{ij} is Poisson as well:

$$A_{ij} \sim Poisson(d_i d_j \Omega_{\iota\gamma}).$$

We prove this fact by multiplying the conditional density of $A_{ij}|c_i$ by the marginal density of c_i to get the joint density of A_{ij} and c_i, and then integrating out c_i.

$$P(A_{ij}, c_i) = \underbrace{P(A_{ij}|c_i)}_{Bin(c_i, d_j\Omega_{\iota\gamma})} \quad \times \quad \underbrace{P(c_i)}_{Poisson(d_i)}$$

Deriving the joint distribution:

$$P(A_{ij}, c_i) = \binom{c_i}{A_{ij}} (d_j\Omega_{\iota\gamma})^{Aij}(1 - d_j\Omega_{\iota\gamma})^{c_i - Aij} \times \frac{d_i^{c_i}\exp(-d_i)}{c_i!}$$

We want to find out the marginal distribution of A_{ij}:

$$\begin{aligned}
P(A_{ij}) &= \sum_{c_i=0}^{\infty} P(A_{ij}, c_i) \\
&= \sum_{c_i=0}^{\infty} \binom{c_i}{A_{ij}} (d_j\Omega_{\iota\gamma})^{Aij}(1 - d_j\Omega_{\iota\gamma})^{c_i - Aij} \times \frac{d_i^{c_i}\exp(-d_i)}{c_i!} \\
&= \sum_{c_i=0}^{\infty} \frac{c_i!}{A_{ij}!(di - A_{ij})!} (d_j\Omega_{\iota\gamma})^{Aij}(1 - d_j\Omega_{\iota\gamma})^{c_i - Aij} \times \frac{d_i^{c_i}\exp(-d_i)}{c_i!} \\
&= \frac{(d_j\Omega_{\iota\gamma})^{Aij}\exp(-d_i)}{A_{ij}!} \sum_{c_i=0}^{\infty} \frac{1}{(di - A_{ij})!} (1 - d_j\Omega_{\iota\gamma})^{c_i - Aij} d_i^{c_i}
\end{aligned}$$

If the summation term is equal to

$$\sum_{c_i=0}^{\infty} \frac{1}{(di - A_{ij})!}(1 - d_j\Omega_{\iota\gamma})^{c_i - Aij} d_i^{c_i} = d_i^{A_{ij}} \exp\left(d_i(1 - d_j\Omega_{\iota\gamma})\right) \tag{28}$$

then $P(A_{ij}) = \frac{(d_i d_j \Omega_{\iota\gamma})^{Aij} \exp(-d_i d_j \Omega_{\iota\gamma})}{A_{ij}!}$, i.e. A_{ij} would be Poisson distributed:

$$A_{ij} \sim Poisson(d_i d_j \Omega_{\iota\gamma})$$

Proving (54) is equivalent to proving the following equality:

$$1 = \frac{1}{d_i^{A_{ij}} \exp\left(d_i(1 - d_j\Omega_{\iota\gamma})\right)} \sum_{c_i=0}^{\infty} \frac{1}{(di - A_{ij})!}(1 - d_j\Omega_{\iota\gamma})^{c_i - Aij} d_i^{c_i}$$

Proof:

$$d_i^{-A_{ij}} \exp\left(-d_i(1 - d_j\Omega_{\iota\gamma})\right) \sum_{c_i=0}^{\infty} \frac{1}{(di - A_{ij})!}(1 - d_j\Omega_{\iota\gamma})^{c_i - Aij} d_i^{c_i} =$$

$$= \sum_{c_i=0}^{\infty} \frac{\exp\left(-d_i(1 - d_j\Omega_{\iota\gamma})\right)}{(di - A_{ij})!}(1 - d_j\Omega_{\iota\gamma})^{c_i - Aij} d_i^{c_i - A_{ij}}$$

$$= \sum_{c_i=0}^{\infty} \frac{\exp\left(-d_i(1 - d_j\Omega_{\iota\gamma})\right)}{(di - A_{ij})!}(d_i(1 - d_j\Omega_{\iota\gamma}))^{c_i - Aij}$$

We assume $\lambda = d_i(1 - d_j\Omega_{\iota\gamma})$ for simplicity and we apply a change of variables $z = c_i - A_{ij}$

$= \sum_{z=0}^{\infty} \frac{\exp(-\lambda)}{z!}\lambda^z$, knowing that in our problem $c_i \geq A_{ij}$, i.e. $z \geq 0$.

$= 1$

Since we have the p.d.f. of a Poisson r.v. inside the summation, i.e. $z \sim Poisson(\lambda)$. $\square$

Therefore, we have

$$A_{ij} \sim Poisson(d_i d_j \Omega_{\iota\gamma}).$$