Referee Report on "Labor Market Power, Self-Employment, and Development" by F. Amodio, P. Medina and M. Morlacco

Romain FERNEX

April 26, 2026

1 Summary

This paper studies the role of self-employment in labor markets, with a focus on low-income countries. Using linked administrative data on Peruvian manufacturing firms combined with household survey data, the authors document how self-employment interacts with employer concentration and labor market power. They estimate firm-level wage markdowns by instrumenting employment with shocks from a rural electrification program, then embed these reduced-form estimates into a general equilibrium model of oligopsonistic local labor markets with Roy-type worker sorting between wage employment and self-employment. Their main argument is that a higher local density of self-employment mitigates wage markdowns in strongly concentrated markets, implying that pro-competitive policies aimed at expanding wage employment may be less effective than expected because they simultaneously reduce the option value of self-employment and can strengthen firms' wage-setting power. The authors argue that this balancing role of self-employment may help explain why self-employment remains persistently large in developing countries, even though classic development theories predict that industrialization should gradually replace it with formal wage employment.

The paper contributes to three main strands of the labor economics literature. First, it challenges the traditional dual-economy view by showing that formal wage employment and informal self-employment are closely linked within local markets through workers' occupational choices, rather than operating as fully segmented markets. Second, it contributes to the oligopsony literature by documenting sizable wage markdowns in Peruvian manufacturing and by showing how employer concentration and aggregate labor supply elasticity jointly

determine wage-setting power once self-employment is modeled as an explicit outside option. Third, it speaks to the literature on earnings gaps and informality by studying how comparative and absolute advantage shape earnings differences between wage workers and the self-employed, and by interpreting informal self-employment as a buffer against labor-market power. This is in line with prior works on the relation between self-employment and labour market frictions (Poschke 2025).

2 Main Comments

My first criticism concerns the firm-level IV strategy used by the authors to identify wage markdowns. The exclusion restriction requires that the PER×EC shock affects wages only through its impact on labor demand, and not through any other channel. The authors dismiss the main threat to this assumption, namely that a positive productivity shock could be shared with workers as rents, on the grounds that union density in Peru is very low throughout the sample period. While the literature supports the idea that unionization facilitates rent-sharing, it also shows clearly that rent-sharing is not exclusive to unionized firms. In particular, Card, Devicienti and Maida (Card, Devicienti, and Maida 2014) demonstrate, using matched Italian employer-employee data, that firms with higher price-cost margins share rents more generously with their workers, exhibiting quasi-rent wage elasticities of 0.059 against 0.024 for low-margin firms, a pattern that holds independently of union coverage. This is directly relevant here because the PER electrification shock, by relaxing a binding input constraint, is likely to raise productivity and quasi-rents most for the most electricity-constrained firms, precisely the group whose wages the IV strategy treats as moving only along with labor supply. The Peruvian institutional environment further undermines the authors' dismissal of rent-sharing: the Decreto Legislativo No. 892 of 1996 mandates that manufacturing firms with more than twenty employees distribute a minimum of ten percent of pre-tax profits to their workers, providing a direct statutory channel through which productivity gains translate into wage increases irrespective of union membership. If a non-negligible fraction of the wage response to the PER shock reflects this statutory profit-sharing or voluntary rent-sharing at high-margin firms, the IV estimate of the inverse labor supply elasticity conflates movements along the labor supply curve with changes in the firm's wage-setting rule, potentially biasing measured markdowns upward and overstating the degree of labor market power underlying the paper's main empirical and counterfactual conclusions. As a direct robustness check, I would encourage the authors to run a standard rent-sharing regression of log wages on log value added per worker, instrumented with the PER shock, separately for high- and low-EC firms, controlling for firm and market×year fixed effects. A

significant positive coefficient, larger for high-EC firms, would indicate the exclusion restriction is likely violated and the IV estimates require reinterpretation.

Aside from the empirical design, I take issue with the model fit results shown in Table 4 of the paper. The authors acknowledge in their model validation exercise that the structural model substantially overpredicts the sensitivity of both wage and self-employment earnings to employer concentration: the model implies log earnings–log HHI slopes of approximately -1.36 and -1.17 respectively, while the data show coefficients of -0.14 and -0.11. In fact, footnote 5 of Online Appendix Section D.4 shows matching this data moment tightly drags down overall model fit as the authors disclose that the earnings-HHI correlation moments are downweighted by a factor of thirty in the MSM loss function precisely because of the poor fit. I conceed that the markdown estimates themselves appear stable across alternative entry-cost specifications, as shown in Online Appendix Table A.9, but this exercise speaks to the parametric form of fixed costs and not to the slope misfit, which the authors never directly target. The tenfold gap therefore remains concerning and cannot be dismissed as a minor imperfection since it shapes the paper's central counterfactual claim that changes in markdowns explain up to ninety-nine percent of the cross-market variation in the wage impact of industrial policies. Because markdowns in the model respond far more to changes in concentration than they do in the data, this figure likely reflects an artifact of the model's overreaction to HHI variation rather than a genuine structural feature of Peruvian labor markets, a concern particularly acute for markets at the extremes of the concentration distribution, where model and data diverge most sharply. The same misfit applies symmetrically to self-employment earnings: by overpredicting the negative HHI gradient, the model sets the baseline value of the self-employment outside option too low in concentrated markets, inflating the apparent cost of pro-competitive policies that reduce self-employment and reinforcing the paper's policy warning in a way that may be more model-driven than data-driven. To assess the robustness of the key results, I believe the authors should also provide model fit estimates when the earnings-HHI regression moments are not downweighted in the MSM estimation and rerun the policy counterfactuals of Table A.13 as a robustness check.

3 Minor Comments

Aside from the above critiques, some of the claims made by the authors would gain from being more nuanced. This is notably the case of the external validity claim in their conclusion, where they suggest their framework generalizes cleanly to other developing economies. It is undeniable that their proposed mechanism bears some resemblance to the view of self-

employment as an intermediate state between unemployment and wage employment, which has gradually taken root in the literature with the rise of the gig economy (Boeri et al. 2020). Nonetheless, some studies invert this paradigm by showing that wage employment may actually serve as an outside option for the self-employed, enabling them to safely experiment with independent work and learn their potential as entrepreneurs (Dillon and Stanton 2017). This distinction matters for external validity: while self-employment is rising again in developed economies, a significant share of it appears driven by positive selection and preference for flexibility rather than by labor market frictions, making the mechanisms identified here considerably less portable to high-income settings than the authors imply. More generally, the institutional and labor market conditions specific to Peruvian manufacturing, including high informality and limited union coverage, are sufficiently particular that the authors should be more circumspect about the breadth of settings to which their framework and quantitative conclusions apply.

4 Paper Extensions

Using their model the authors show extensively that adopting pro-competitive policies and raising productivity in wage employment may yield attenuated welfare gains, due to the fact that more workers enter wage employment, which in turn lowers the value of self-employment as an outside option. They illustrate this with two policies that raise productivity in wage employment: Peru's Youth Job Training Program (PROJoven) and a rural roads program aimed at improving firm productivity and market access. They get convincing results for both. What they never directly examine, however, is the symmetric case: what happens when the shock targets self-employment specifically?

Based on the microfoundations of their model, a policy that facilitates entrepreneurship should strengthen the value of the outside option represented by self-employment and, in turn, raise wages most strongly in concentrated markets where employers have sizable wage-setting power. If it does not, this means the value of self-employment is likely overestimated by the model and its role in wage-setting may be more tenuous than the model suggests. To study this empirically, I propose an empirical design leveraging the geographic expansion of the Cajas Municipales (CMAC) microfinance network across Peru over the period 2004–2011, which introduces plausibly exogenous variation in the value of self-employment.

4.1 On the Cajas Municipales and Peruvian Microfinance

The Cajas Municipales de Ahorro y Crédito (CMACs) are microfinance institutions owned by municipal governments but governed by a board composed of members of civil society, originally meant to provide credit to people with limited access to formal financing. The first CMAC opened its doors in 1982 and twelve more were founded by 1992. The institutional roster has been broadly stable since, with a few mergers and one liquidation bringing the count to ten today.

What truly expanded over time was the geographic footprint of the network. Until 2002 each CMAC was restricted to operating in its home region under a complementarity rule enforced by FEPCMAC, the federation of CMACs. The 2002 deregulation lifted this territorial restriction and triggered a nationwide branch rollout. Based on SBS data[1], the number of CMAC offices rose from 136 in December 2001 to 296 in December 2005 and 900 by December 2010, with the steepest expansion phase taking place between 2002 and 2011, which roughly corresponds to the timeframe considered in Amodio et al.'s paper. As of the early 2010s, the CMACs represented 40.6% of the Peruvian microfinance industry (Jaramillo 2013), making them a prime source of financing for the credit-constrained workers who form the core demographic behind the paper's dynamics.

4.2 Empirical Design

To test the hypothesis above empirically, I rely on two datasets: the ENAHO panel data, which the main paper already uses, and SBS administrative data tracking the number of CMAC offices in each Peruvian department at monthly frequency, which I aggregate to annual to align with ENAHO. The SBS data give an exact count of CMAC offices per department-year, which I use as a proxy for local access to microenterprise credit and thus for the relative attractiveness of self-employment in each market. One important caveat is that the SBS reports CMAC offices at the department level whereas Amodio et al.'s local labor markets are defined at the province level. I therefore assign each province the CMAC intensity of its parent department, which likely introduces measurement error in the treatment but preserves cross-departmental variation. I then estimate the following design on the population of wage-employed workers in order to measure the response of log wages to the expansion of the CMACs:

$$\log(w_{ipt}) = \beta_1 \, CMAC_{p,t} + \beta_2 \, CMAC_{p,t} \times \widetilde{HHI}_p^{pre} + \gamma X'_{ipt} + \alpha_p + \lambda_t + \varepsilon_{ipt}$$

[1]SBS is the Superintendencia de Banca, Seguros y AFP, the Peruvian regulator of the financial system. Office counts are computed from SBS report C-1201 "Distribución de Oficinas por Zona Geográfica".

Where:

- $\log(w_{ipt})$ is the log hourly wage of worker i in province p in year t
- $CMAC_{p,t}$ is a continuous treatment variable defined as the logarithm of one plus the number of offices per 100,000 working-age residents, which is built using SBS data and household location information in the ENAHO
- $\widetilde{HHI}_p^{pre}$ is the employer wage-bill HHI used in the source paper. It is computed only on pre-period data and kept constant across years, so that the variation in the interaction term comes entirely from the CMAC expansion. I use this interaction term to check whether wages indeed increase more in highly concentrated markets compared to less concentrated ones, in accordance with the mechanism described in the structural model.
- X_{ipt} are individual controls including age, age squared, gender, education and a dummy for living in an urban area
- α_p and λ_t are fixed effects for provinces and years respectively.

Lastly I cluster errors at the department level, which is where the treatment actually varies. If β_2 is strongly positive, this signifies that making self-employment more valuable indeed raises wages disproportionately in concentrated markets, in line with the mechanism described by Amodio et al. If β_2 is null or negative, this goes against one of the core model predictions about the role of self-employment in concentrated markets, which means the structural results should be interpreted with caution. While β_1 on its own captures the average wage response to credit access, it is β_2 that discriminates Amodio et al.'s mechanism from generic credit-supply stories.

4.3 Limitations and potential robustness checks

The design I proposed above does have some genuine flaws : First, there may be an endogeneity issue if CMACs did not expand randomly. They may target departments with growing demand or shifting employment composition, which would bias both coefficients. It is possible, however, to correct for this using the following shift-share instrument. The shares are time-invariant baseline weights computed in 2001, before the deregulation, capturing each province's geographic exposure to each of the 13 incumbent CMACs through the inverse distance from province p to CMAC c's headquarters. The shifts are each CMAC's leave-one-out national expansion trajectory, that is, the cumulative number of branches opened by CMAC c outside province p up to year t. Combining the two, the predicted CMAC exposure of

province p in year t is

$$Z_{p,t} = \sum_{c=1}^{13} \text{share}_{p,c}^{2001} \times \text{branches}_{c,t}^{-p}$$

which I use to instrument $CMAC_{p,t}$ in 2SLS. The logic is that each component of $Z_{p,t}$ is plausibly exogenous to local wage dynamics in p: the shares are fixed before the 2002 deregulation that triggered the rollout and therefore cannot respond to subsequent shocks, while the leave-one-out shifts are driven by each CMAC's institution-specific factors such as capital base, management capacity, and federation-level initiatives, rather than by conditions in any single destination province. Even if a CMAC strategically targets growing markets near its headquarters, the timing of its national push is plausibly orthogonal to the wage trajectory of a distant province whose only link to that CMAC is geographic proximity.

There may also be spatial spillover if workers choose to commute across departments in response to the expansion of the program. This is unlikely due to both the limited mobility across markets and regions documented by Amodio et al., who report that roughly 85% of worker moves stay within geographic area, and to the fact that people who benefit most from microfinance have very modest means and are thus less likely to migrate across regions.

5 References

Boeri, Tito et al. (2020). "Solo Self-Employment and Alternative Work Arrangements: A Cross-Country Perspective on the Changing Composition of Jobs". In: *Journal of Economic Perspectives* 34.1, pp. 170–195. DOI: 10.1257/jep.34.1.170. URL: https://www.aeaweb.org/articles?id=10.1257/jep.34.1.170.

Card, David, Francesco Devicienti, and Agata Maida (2014). "Rent-Sharing, Holdup, and Wages: Evidence from Matched Panel Data". In: *The Review of Economic Studies* 81.1, pp. 84–111. URL: https://www.jstor.org/stable/43551667.

Dillon, Eleanor Wiske and Christopher Stanton (2017). *Self-Employment Dynamics and the Returns to Entrepreneurship*. NBER Working Paper 23168. National Bureau of Economic Research. URL: https://www.nber.org/papers/w23168.

Jaramillo, Miguel (2013). *Peruvian Microfinance's Caja Municipal Model*. ELLA Case Study Brief. Lima, Peru: Group for the Analysis of Development (GRADE).

Poschke, Markus (2025). "Wage Employment, Unemployment and Self-Employment across Countries". In: *Journal of Monetary Economics* 149, p. 103684. ISSN: 0304-3932. DOI: 10.1016/j.jmoneco.2024.103684. URL: https://www.sciencedirect.com/science/article/pii/S0304393224001375.