Monetary Policy Effects on Wage Inequality Between and Within Firms

Christian Moser* Columbia Business School

Farzad Saidi† Stockholm School of Economics & CEPR

Benjamin Wirth‡ IAB Nuremberg

February 15, 2018

Abstract

We study pass-through of negative interest rates to workers' wages in the European currency union from 2014-16. To this end, we construct a novel dataset combining administrative linked employer-employee (IAB-LIAB) data with proprietary syndicated loans (Dealscan) and executive compensation (BoardEx) data from Germany. To identify monetary policy effects on wage inequality between and within firms, we exploit the interaction of nominal interest rate movements around the zero lower bound with variation in pre-determined balance sheet exposure of banks and their lending relations with firms. We find significant increases in credit supply of affected banks, leading to increased within-firm inequality. The results are driven by executives and highest-paid employees reaping relatively greater benefits from positive firm-level credit supply shocks. At the same time, low-paying firms increase their average pay in response to easier credit access.

JEL classification: D22, G21, G31, G32, J31

Keywords: credit supply, monetary policy, negative policy rates, wage inequality, risk taking, financial constraints

*Columbia University, Graduate School of Business, Uris Hall 819, 3022 Broadway, New York, NY 10027, USA. E-mail: c.moser@columbia.edu

†Stockholm School of Economics, Swedish House of Finance, Drottninggatan 98, 111 60 Stockholm, Sweden. E-mail: Farzad.Saidi@hhs.se

‡Institute for Employment Research (IAB), Regensburger Str. 104, 90478 Nuremberg, Germany. E-mail: Benjamin.Wirth@iab.de

1 Introduction

There is increasing interest in employer heterogeneity as an important determinant of worker outcomes, including employment risk and earnings inequality (Abowd et al. 1999, Card et al. 2013, Alvarez et al. 2018). An emerging lesson from these studies is that a significant share of the rise in earnings inequality is driven by greater between-firm pay differences over time. Yet the drivers of these

We study pass-through of negative interest rates to workers' wages in the European currency union from 2014-16. To this end, we construct a novel dataset combining administrative linked employer-employee (IAB-LIAB) data with proprietary syndicated loans (Dealscan) and executive compensation (BoardEx) data from Germany. To identify monetary policy effects on wage inequality between and within firms, we exploit the interaction of nominal interest rate movements around the zero lower bound with variation in predetermined balance sheet exposure of banks and their lending relations with firms.

We find significant increases in credit supply of affected banks, leading to increased within-firm inequality. The results are driven by executives and highest-paid employees reaping relatively greater benefits from positive firm-level credit supply shocks. At the same time, low-paying firms increase their average pay in response to easier credit access.

Our work builds on Heider, Saidi, and Schepens (2017), who show that the transmission of negative, rather than non-negative, rates depends on banks' funding structure. Their empirical identification is centered on the idea that banks are reluctant to pass on negative rates to their depositors. Therefore, banks that rely more on deposit funding—as opposed to interbank lending—experience a negative shock to their net worth, which induces bank risk taking. Heider, Saidi, and Schepens (2017) show that this risk taking by high-deposit, rather than low-deposit, banks takes on the form of more lending to riskier firms.

2 Testable Hypotheses

We hypothesize that firms that are financially constrained set wages in response to such constraints. Financial constraints can, in turn, be relaxed through credit provision by banks, while credit supply is partly governed by monetary policy.

Variations in firms' financial constraints and their subsequent risk taking may have an effect on the distribution of wages within firms if one assumes that firms are constrained in their ability to offer optimal wages, typically characterized by workers' marginal product of labor, and differentially so across workers according to their position in the wage distribution.

Denote the wedge between her wage and marginal product of labor for a top earner at firm j by $\Delta_j^{top} = wage_j^{top} - MPL^{top}$ and accordingly for a bottom earner $\Delta_j^{bottom} = wage_j^{bottom} - MPL^{bottom}$. When a firm's financial constraints are relaxed, it will optimally reduce the wedge first which implies the greatest marginal benefit in terms of returns to labor. If this holds true for top rather than bottom earners, then the relaxation of financial constraints should lead to greater within-firm wage inequality. A potential reason for why firms may adjust wages of top, rather than bottom, earners upward following the relaxation of financial constraints is that the latter's labor market is more competitive and higher wages are an important means of attracting rare talent.

Conversely, when a firm's financial constraints become tighter, the opposite may hold true, depending on whether the marginal benefit in terms of returns to labor is constant across the two scenarios. This would be consistent with the finding of Guiso, Pistaferri, and Schivardi (2013) that firms borrow from their workers when they are financially constrained due to lacking financial development in their location, and more so from workers who are relatively wealthy.

Motivated by the idea that riskier firms exhibit higher within-firm wage inequality, we furthermore explore the role of firm-level risk for inequality. In doing so, we maintain the premise that while risky firms are more likely to be financially constrained, not all financially constrained firms are necessarily risky. To the extent that relaxing financial constraints leads

to riskier investment, our empirical evidence is consistent with the above-mentioned notion that weaker financial constraints can lead to greater within-firm wage inequality.

Firm-level risk can govern the distribution of wages through at least two channels. We distinguish between workers' compensation that may be differentially linked to firm-level volatility as a parameter (in their contract) according to the position of workers in the wage distribution, and the time variation in firms' cash flows as implied by their riskiness.

First, if within-firm risk increases, employees of that firm may be differentially affected as long as their compensation is differentially linked to firm-level risk. For instance, a firm's CEO may be compensated through stock options the value of which is increasing in risk, whereas more regular (think blue-collar) workers' compensation is orthogonal to firm-level volatility. In the more extreme case of an excessive increase in risk putting a firm into dire straits, employees at the bottom of the firm's career hierarchy are more likely to be laid off. These considerations altogether suggest that a within-firm increase in risk translates into higher within-firm wage inequality (across groups).

When examining this channel, we are constrained by the fact that we do not necessarily observe variable compensation in the data. Furthermore, total pay outcomes under variable-wage schemes do not only depend on the variable-wage rate set by the firm but also on workers' effort. In the above example, we assume that workers do not adjust their working hours, which is an assumption that we need to test empirically.

Second, higher firm-level risk implies greater variation in firms' cash flows over time. Fagereng, Guiso, and Pistaferri (2017) document that workers are partially insured against firm-level shocks, in that there exists non-zero pass-through of such shocks to workers' wages. In particular, they show that firms offer less insurance to workers with higher wealth, particularly against permanent, rather than transitory, firm-level shocks.

The implications of this second channel for the distribution of wages within firms depends on the differential pass-through of positive vs. negative firm-level shocks to workers according to their position in the wage distribution. In particular, while top earners' wages may covary

more with firm-level shocks *overall* (Fagereng, Guiso, and Pistaferri (2017)), this may be primarily due to the greater pass-through of positive firm-level outcomes. This, in turn, would imply greater within-firm wage inequality.

Similarly, bottom earners may be affected by greater pass-through of negative firm-level outcomes, in that they are laid off more frequently following the realization of such negative firm-level outcomes. Note that this implies a sample bias in Fagereng, Guiso, and Pistaferri (2017) if they limit their sample to workers that remain employed at the same firm undergoing any shocks.

3 Empirical Strategy and Implementation

Using the matched employer-employee data, we are able to build a worker panel that records working stints of worker i at any firm j in year t. From the firms' point of view, we will separately investigate their wage-setting behavior for new hires (extensive margin) and the wage development of previously hired workers (intensive margin).

At any given point in time, an individual i can be employed or unemployed. Furthermore, if individual i is employed at firm j in year t, then there exist three possibilities regarding worker i's employment status in the previous year $t-1$, namely: she could have been unemployed, employed at the same firm j, *and/or* employed by another firm that is not j (the qualification *and* refers to cases where worker i switched jobs in $t-1$).

Within a given worker-year, we keep the main job j. By main job, we define the job held by worker i in December. If there are multiple jobs that month, we keep the one with the highest nominal wage.

Then, for each worker, we expand the dataset to include zero wages (unemployment spells) in $t-1$ $(t+1)$ if worker i is employed by firm j in year t but not in $t-1$ $(t+1)$, unless they re-join the same firm j as in $t-2$ $(t+2)$, e.g., maternity leave. Finally, we drop all observations associated with public, finance, and utilities firms.

When analyzing the extensive margin, we keep observations of each worker i at firm j in year t only if she did not work for firm j before. In order to be able to compare wages across employees within and between firms, we annualize the recorded wage by using information on within-year employment spells. At the extensive margin, we may have previously unemployed individuals joining a firm, or they switch firms. We classify an individual as unemployed in a given year t if she does not receive any wage that year. Again, we annualize the recorded wage by using information on within-year employment spells.

When we zoom in on the intensive margin, we limit the sample to currently employed workers i at firm j in year t that already worked at the same firm in year $t-1$.

Furthermore, whenever the unit of observation is at the worker level ijt, we control for (a proxy for) part-time-employment status on the right-hand side of our regression specification.

Main specifications: Our dependent variables are functions of individual wages or the distribution thereof. For our scrutiny of the intensive margin and switching employees at the extensive margin, we can estimate the following specification:

$$wage_{ijt} = \beta_1 Decile_{it-1} \times Shock_{jt} + \beta_2 Decile_{it-1} + \beta_3 X_{ijt} + \mu_{jt} + \theta_{ij} + \epsilon_{ijt}, \qquad (1)$$

where $wage_{ijt}$ is worker i's annualized wage at firm j in year t, $Decile_{it-1}$ is a vector for each decile according to the position of her wage in the distribution of firm j or elsewhere in year $t-1$, $Shock_{jt}$ is a firm-level change at j in year t, X_{ijt} are other controls, such as an indicator for part-time employment, μ_{jt} denotes firm-time fixed effects, and θ_{ij} are worker-firm fixed effects. Standard errors are clustered at the firm level.

Note that in the above setting, one can further interact our explanatory variables with an indicator for whether individual i worked at the same firm j in $t-1$. This serves the purpose of distinguishing between the intensive-margin effect and a potential extensive-margin effect stemming from employees switching firms.

In particular, for non-switchers, it holds that firm j associated with both $Decile_{it-1}$ and the fixed effects μ_{jt} are identical. Therefore, we set $\mu_{jt} = \mu_{jt-1}$ for all workers i that are

unemployed in period t but were employed by firm j in period $t-1$.

What is more, θ_{ij} may be replaced by θ_i if worker i is never observed to switch firms, or if she does switch firms, but leaves each job after at most one year.

When we wish to analyze within-firm wage inequality for new hires, we observe each worker-firm pair ij only once, so that our measure of within-firm wage inequality is collapsed to the firm-year level jt. We then estimate the following specification:

$$inequality_{jt} = \beta_1 Shock_{jt} + \beta_2 X_{jt} + \mu_j + \eta_t + \epsilon_{jt}, \tag{2}$$

where $inequality_{jt}$ is a measure of within-firm wage inequality (possibly limited to full-time employees) at firm j in year t, $Shock_{jt}$ is a firm-level change at j in year t, X_{jt} are other firm-level controls, μ_j denotes firm fixed effects, and η_t are year fixed effects. Standard errors are clustered at the firm level.

Alternatively, the analysis could also be conducted at the ijt level. In doing so, one concentrates on all new hires of a given firm j in year t, and worker i may occur multiple times in the sample if she is observed to start multiple jobs during the run-time of the data. As $Decile_{it-1}$ is only defined for switchers, but not for previously unemployed individuals, one has to modify (1) for example as follows:

$$wage_{ijt} = \beta_1 Skill_{it-1} \times Shock_{jt} + \beta_2 Skill_{it-1} + \beta_3 X_{ijt} + \mu_{jt}(+\theta_i) + \epsilon_{ijt}, \tag{3}$$

where $wage_{ijt}$ is newly hired worker i's annualized wage at firm j in year t, $Skill_{it-1}$ is a pre-determined characteristic of worker i in year $t-1$, e.g., skill, $Shock_{jt}$ is a firm-level change at j in year t, X_{ijt} are other controls, such as an indicator for part-time employment, μ_{jt} denotes firm-time fixed effects, and θ_i are worker fixed effects that can be included only for workers that switch jobs at least once. Standard errors are clustered at the firm level.

3.1 Monetary-policy transmission as a source of firm-level variation

As a shock to firm-level risk, we use the introduction of negative policy rates in the euro area and, thus, Germany. Heider, Saidi, and Schepens (2017) show that the introduction of negative policy rates enabled high-deposit, rather than low-deposit, banks to finance riskier firms.

Our empirical identification exploits this positive credit supply shock for risky projects of firms that are in relationships with high-deposit, rather than low-deposit, banks after the implementation of the negative policy rate in June 2014. Using the matched employer-employee data, this enables us to trace the income-inequality implications of higher firm-level risk for German workers within affected firms as well as between affected and non-affected firms. We use syndicated-loan data from DealScan to sort firms into those in financing relationships with high-deposit banks (treatment group) vs. low-deposit banks (control group).

$Shock_{jt}$ can thus be understood as an interaction term between $Treatment_j$ and $After_t$, where $Treatment_j$ is a variable that indicates the treatment (intensity) as a function of firm j's pre-2014 banking relationships (in particular the degree of deposit funding of those banks) and $After_t$ is a post dummy for the period from June 2014 onwards.

3.2 Mechanism underlying the risk-inequality link:

We hypothesize that changes in firms' wage distribution react to the relaxation of financial constraints for firm-level risk taking. There are at least two channels through firm-level risk may affect wage inequality.

First, the relaxation of financial constraints for risky investments may lead to higher *realized* volatility, e.g., sales-growth volatility. This may capture the effect that a higher volatility has on the compensation of different types of workers in a given firm-year cluster,

e.g., volatility enters the wage bill as a parameter, as is the case for the value of stock options which are granted to some, but not all, employees within a firm.

Alternatively, a higher volatility may imply that firm-level outcomes, such as sales growth between year t and $t-1$, are more volatile *over time*, i.e., higher positive and/or negative sales-growth realizations. If risk matters in this time-varying sense, then any effect of higher volatility on within-firm wage inequality should be reflected in differential pass-through of positive and negative sales-growth realizations at the jt level to workers. To test this, $Shock_{jt}$ is split into a positive sales-growth-realization and a separate, negative sales-growth-realization component. As both events cannot by definition occur within a given firm-year cluster, the following regression specification would identify differential pass-through of positive and negative sales-growth realizations to workers *within firms but across years*. This is reflected in the following regression specification for the intensive margin and for switchers:

$$\begin{aligned} wage_{ijt} &= \beta_1 Decile_{it-1} \times Sales^{+}_{jt} + \beta_2 Decile_{it-1} \times Sales^{-}_{jt} + \beta_3 Decile_{it-1} \\ &\quad + \beta_4 X_{ijt} + \mu_{jt} + \theta_{ij} + \epsilon_{ijt}, \end{aligned} \tag{4}$$

where $wage_{ijt}$ is worker i's annualized wage at firm j in year t, $Decile_{it-1}$ is a vector for each decile according to the position of her wage in the distribution of firm j or elsewhere in year $t-1$, $Sales^{+}_{jt}$ is the absolute amount of a positive sales-growth realization of firm j in year t, $Sales^{-}_{jt}$ is the absolute amount of a negative sales-growth realization of firm j in year t, X_{ijt} are other controls, such as an indicator for part-time employment, μ_{jt} denotes firm-time fixed effects, and θ_{ij} are worker-firm fixed effects.

Similarly, for the extensive margin of new hires:

$$inequality_{jt} = \beta_1 Sales^{+}_{jt} + \beta_2 Sales^{-}_{jt} + \beta_3 X_{jt} + \mu_j + \eta_t + \epsilon_{jt} \tag{5}$$

and

$$\begin{aligned} wage_{ijt} &= \beta_1 Skill_{it-1} \times Sales^{+}_{jt} + \beta_2 Skill_{it-1} \times Sales^{-}_{jt} + \beta_3 Skill_{it-1} \\ &\quad + \beta_4 X_{ijt} + \mu_{jt}(+\theta_i) + \epsilon_{ijt}, \end{aligned} \tag{6}$$

where – as argued above – one can also add worker fixed effects θ_i for workers that switch jobs at least once.

3.3 Potential empirical challenges:

Final individual wages are equal to the sum of their fixed and variable components, where the latter is a function not only of the variable-wage rate but also of worker effort, which may vary in different dimensions, both qualitatively and quantitatively (e.g., hours worked). One of our empirical challenges is that we do not necessarily observe these wage components. However, it is reasonable to assume that variable wages exist only for certain parts of the wage distribution, and most certainly for top earners. Furthermore, for individuals i with a previous work history leading up to year t in a given occupation (or job role) k at firm j, we can ascertain whether their historical wage variability suggests the existence of variable components of their total wages. We denote this indicator variable $Variable_{ijt-1}$.

To simplify matters, we focus on workers that do not switch firms, so that $Variable_{ijt-1} = Variable_{it-1}$. Then, we can augment (1) as follows to test for any differential impact of variable wages:

$$\begin{aligned} wage_{ijt} &= \beta_1 Decile_{it-1} \times Shock_{jt} \times Variable_{it-1} + \beta_2 Decile_{it-1} \times Shock_{jt} \\ &\quad + \beta_3 Shock_{jt} \times Variable_{it-1} + \beta_4 Decile_{it-1} \times Variable_{it-1} \\ &\quad + \beta_5 Decile_{it-1} + \beta_6 Variable_{it-1} + \beta_7 X_{ijt} + \mu_{jt} + \theta_i + \epsilon_{ijt}. \end{aligned} \tag{7}$$

If β_1 is zero, then we deem the additional effect of variable wages and workers' work responses to their incentives on within-firm wage inequality to be low. Note that in (7), the stand-alone

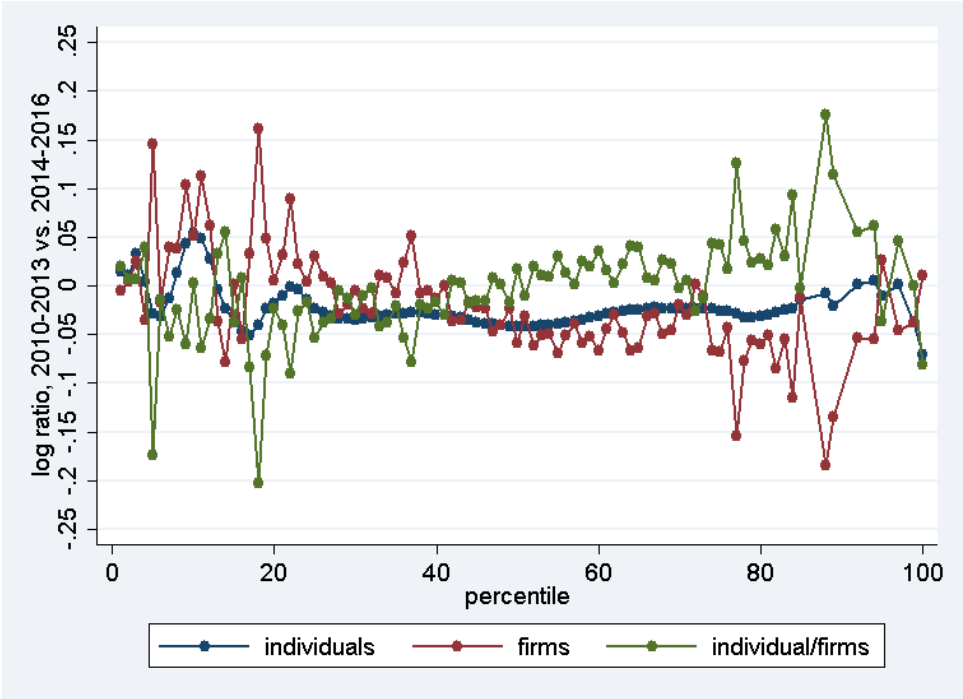
wage effect of workers' incentives following a firm-level shock is captured by β_3.

4 Data

For the first time, this paper combines three major datasets covering bank balance sheets, firm credit relationships, and worker-level for Germany. Specifically, we merge administrative linked employer-employee (IAB-LIAB) data with proprietary syndicated loans (Dealscan) and executive compensation (BoardEx) data.

5 Results

Based on the above econometric decompositions, the following graphs plot overall wage growth (blue dots) decomposed into its between-firms component (red dots), and the within-firms component (green dots) for the difference between treated and control firms. More results are pending approval from the German statistical agency (IAB).



6 Discussion

TBC.

7 Conclusion

Our preliminary results indicate that monetary policy has substantial effects on the distribution of wages within and between employers. We find significant increases in credit supply of affected banks, leading to increased within-firm inequality. The results are driven by executives and highest-paid employees reaping relatively greater benefits from positive firm-level credit supply shocks. At the same time, low-paying firms increase their average pay in response to easier credit access.

References

Fagereng, A., L. Guiso, and L. Pistaferri (2017): "Portfolio Choices, Firm Shocks and Uninsurable Wage Risk," *Review of Economic Studies*, forthcoming.

Guiso, L., L. Pistaferri, and F. Schivardi (2013): "Credit within the Firm," *Review of Economic Studies*, 80(1), 211–247.

Heider, F., F. Saidi, and G. Schepens (2017): "Life Below Zero: Bank Lending Under Negative Policy Rates," *ECB Working Paper*.