Essays in Macroeconomics and Labor Economics

Inauguraldissertation
zur Erlangung des akademischen Grades
eines Doktors der Wirtschaftswissenschaften
der Universität Mannheim

im Frühjahrs-/Sommersemester 2021 vorgelegt von

Moritz Drechsel-Grau

Abteilungssprecher:	Prof. Volker Nocke, Ph.D.
Vorsitzender der Disputation:	Prof. Antonio Ciccone, Ph.D.
Referent:	Prof. Tom Krebs, Ph.D.
Koreferentin:	Prof. Michèle Tertilt, Ph.D.

Tag der Disputation:	07. Juli 2021

Eidesstattliche Erklärung

Hiermit erkläre ich, dass ich die vorliegende Dissertation selbstständig angefertigt und die benutzten Hilfsmittel vollständig und deutlich angegeben habe.

Stuttgart, am 15. Juli 2021

Moritz Drechsel-Grau

Curriculum Vitae

2015–2021	University of Mannheim *Doctoral Program in Economics*
2015–2018	University of Mannheim *Master of Science in Economics (with distinction)*
2010–2014	University of Tübingen *Bachelor of Science in International Economics and European Studies*

Acknowledgments

I am deeply indebted to my advisors Tom Krebs, Michèle Tertilt, Sebastian Siegloch and Andreas Peichl. I am thankful for their guidance and support over the past six years, and in particular heading into the job market. I also thank Klaus Adam for his support and for letting me join his weekly doctoral colloquium where we could present our ideas in a safe and stimulating environment.

My research has benefited from useful feedback by Antonio Ciccone, Ulrich Wagner, Sebastian Findeisen, Andreas Gulyas, Antoine Camous, Katja Kaufmann, Andrey Alexandrov, Lukas Henkel, Timo Reinelt, Benjamin Moll, Dan Greenwald, Amir Sufi, Benjamin Larin, Florian Zimmermann, Rosa Ferrer, and many other visiting researchers who were encouraging and showed interest in my work.

I am grateful for generous financial support from the Graduate School of Economics and Social Sciences (GESS) at the University of Mannheim and the German Research Foundation (through Project A04 of CRC TR 224) that allowed me to focus on my research. I am extremely appreciative of the staff at the University of Mannheim and the GESS, in particular Thomas Behles for his prompt and competent technical support, and Corinna Jann-Grahovac for her help with all the paperwork as well as her kind and caring nature. Chapter 3 of this dissertation would not have been possible without many inspiring exchanges with the anonymous manufacturing firm and their willingness to share their personnel data for scientific purposes. I also thank the state of Baden-Württemberg for providing computational resources through the bwHPC.

I am glad that I could share my PhD experience with many fellow doctoral students. Thank you Tomasz, Yasemin, Andrey, Enrico, Nils, Claudia for passionate discussions over lunch, random laughs in the office and many battles on the tennis court.

I am especially grateful to have met Fabian and Felix who are not only amazing co-authors but have become great friends. They made the struggles of the first year, the excitement over new ideas and the efforts to put them into action both bearable and fun. Thank you for your sharp minds and your optimism. Working together with you has been the most rewarding part of my doctoral studies.

Most importantly, I am thankful to my friends and family for being a source of welcome distractions and for supporting me and my little family when I tried to prepare for the job market in the midst of a pandemic. In particular, I thank my parents Moritz and Zita for always having my back, my brother Edgar for being my closest friend, my wonderful wife Johanna for putting up with me during rough times and celebrating with me when things are going well, and my joyous little son Josef for coming into my life (and my office) and teaching me to put things into perspective. Thank you for your love and support.

Contents

Introduction

This dissertation studies questions in the fields of macroeconomics and labor economics. It consists of three self-contained chapters. While each chapter uses different methods and data to study distinct questions, they all reflect my desire to deepen our understanding of the causes and consequences of socioeconomic inequalities, as well as the policies used to mitigate them.

Over the past decades, many advanced economies such as the US or Germany have experienced a sharp rise in wage and earnings inequality. Improving wage and income growth for low-skill workers is clearly among the most pressing tasks that policy makers face. One of the most controversial but also wide-spread policy tools to raise wages at the bottom and reduce wage inequality is the minimum wage. Currently, many countries are discussing proposals to substantially increase the legal wage floor. However, policy makers lack a comprehensive analysis of the macroeconomic and distributional implications of raising the minimum wage beyond observed levels. The first chapter of my dissertation takes a first step towards filling this gap by using a rich search-and-matching model of the labor market as a laboratory to analyze counterfactually high minimum wage levels.

Rising inequality is not only relevant for distributional concerns and anti-poverty policy. The distribution of income and wealth can also affect aggregate economic variables. In the second chapter, which is joint work with Fabian Greimel, I study whether rising income inequality played a roll in the dramatic increase of US household debt between 1980 and 2007. We ask whether non-rich households increased their demand for housing and mortgages in an attempt to keep up with the housing of the richer Joneses. In order to quantify this mechanism, we enrich a macroeconomic model with the old idea that people's well-being and economic decisions are influenced by the social benchmark set by the rich.

While the gender wage gap has declined considerably, convergence has slowed down and substantial gender inequality persists. Gender pay gaps continue to receive significant attention as policy makers ponder the usefulness of gender quotas and firms roll out programs training their employees to become more aware of gender-related biases. The third chapter, co-authored with Felix Holub, adds to our understanding of the causes of gender pay inequality. We offer the first analysis of the importance of bosses for gender differences in pay and performance ratings. We also ask whether the over-representation of men in management positions is a structural disadvantage for women because male bosses pay and evaluate female employees less favorably.

In the following, I briefly summarize each chapter.

Employment, Output and Welfare Effects of Minimum Wages

The first chapter of my dissertation uses a rich search-and-matching model of the labor market as a laboratory to analyze counterfactual minimum wage levels. In the model, the effect of the minimum wage on employment is ambiguous since firms' vacancy posting and workers' job search decisions are affected in opposite directions. On the one hand, firms will lower their vacancy creation as the minimum wage reduces profits. On the other hand, the minimum wage increases workers' search effort as higher wages make finding a job more attractive. In addition to the employment effect, minimum wages also affect output by changing the *composition* of jobs along two dimensions. First, raising the minimum wage increases average productivity because profits and thus vacancy posting decline more strongly for low-productivity firms. Second, raising the minimum wage increases the average employment level because low-productivity jobs are disproportionately also low-hours jobs.

I estimate the model using German administrative and survey data, and show that it provides a good approximation to the joint distribution of wages, firm productivity and employment levels in Germany prior to the introduction of the minimum wage. More importantly, I find that the model can capture the muted employment response, as well as the reallocation effects in terms of productivity and employment levels found by empirical research on the German introduction of a federal minimum wage in 2015. This reform had a so-called Kaitz index of 47% meaning that the minimum wage of 8.5€ was about 47% of the full-time median wage just before the reform.

I then simulate the short- and long-run effects of increasing the minimum wage. In the long-run, total employment, i.e. the number of jobs, does not fall below the baseline level (without a minimum wage) until the minimum wage reaches a Kaitz index of 70% (13€). This is because higher search effort and lower vacancy posting roughly offset each other. In addition, raising the minimum wage reallocates workers towards full-time jobs and high-productivity firms. Total hours worked and output peak at Kaitz indices of 73% (13.5€) and 78% (14.4€) respectively. However, analyzing the entire transition path reveals that policy makers face an important inter-temporal trade-off as large minimum wage hikes lead to substantial job destruction, unemployment and recessions in the short-run. Finally, not all workers benefit equally from higher minimum wages. For women, who often rely on low-hours jobs, the disutility from working longer hours outweighs the utility of higher incomes. Moreover, high minimum wages force low-skill workers into long-term unemployment.

Falling Behind: Has Rising Inequality Fueled the American Debt Boom?

This chapter investigates whether rising income inequality and *Keeping up with the richer Joneses* fueled the US mortgage boom through an increase in the *demand for housing*. We begin by documenting novel aspects of the US mortgage and housing boom which call for such a demand-side mechanism to complement supply-side drivers of household debt. First, we document that (non-rich) household debt grew substantially more in US states where top incomes grew faster. Second, we show that higher state-level top incomes drive up mortgage debt but do not affect non-mortgage debt suggesting that housing plays a key role in the transmission from rising top incomes to rising household debt. Third, we find that house prices grew faster in states with a stronger increase in top incomes pointing to an important role for housing demand.

We then assess the aggregate consequences of rising income inequality in the presence of social comparisons—both analytically and quantitatively. We incorporate social comparisons into a heterogeneous agent model of the macro economy. In the model, households not only care about their own consumption and housing, but also about how their house compares to the benchmark set by the rich. When top incomes rise and the rich upgrade their houses, the non-rich substitute status-enhancing housing for status-neutral consumption to keep up with the rich. These houses are mortgage-financed, causing a boom in debt-to-income ratios across the entire income distribution, and an increase in house prices.

In a stylized version without idiosyncratic income risk, we can show analytically how this status externality affects aggregate debt depending on who cares about whom in the network of social comparisons. In the empirically relevant case where households care about the rich, the debt-to-income ratio of the non-rich is increasing in top incomes.

We then calibrate the full model with idiosyncratic income risk in order to quantify the contribution of this mechanism to the observed increase in mortgages and house prices between 1980 and 2007. We find that in the presence of *Keeping up with the richer Joneses* (KURJ), the rise in income inequality can explain half of the observed 120%-increase in the mortgage-to-income ratio and two thirds of the observed 60%-increase in house prices between 1980 and 2007. The model also accounts for about half of the observed 65%-increase in the house-value-to-income ratio between 1980 and 2007.

Finally, we compare the effects of our demand-side mechanism to those of the *Global Saving Glut*, i.e. the surge in the foreign net debt position of the US from about 0% of GDP in 1980 to about 50% of GDP in 2007. In our model, this increase in the supply of credit generates a similar debt boom through lowering the real interest rate by about 40%. However, the Global Saving Glut increases house prices by only 2% and the ratio of house values to income by only 4%. Both mechanisms together can explain three quarters of the observed 120%-increase in the mortgage-to-income ratio.

Gender Gaps and the Role of Bosses

This chapter investigates the contribution of managers to gender gaps and analyzes whether the over-representation of men in management positions puts women at a disadvantage. To that end, we bring in unique personnel data provided by one of the largest European manufacturing firms. These data not only contain detailed information on job characteristics, demographics, compensation and performance evaluations, they also allow us to trace out the entire firm hierarchy and identify every employee's coworkers, superiors, and subordinates.

We first decompose the raw gender gaps in base salaries, bonus payments and the probability of receiving a high performance evaluation. For all outcomes, we find that the sorting of male and female workers to different managers increases the gender gap (in favor of men). That is, conditional on job and worker characteristics, men work for managers who pay and evaluate their employees better.

We then analyze adjusted gender gaps—controlling not only for job characteristics but also for manager fixed effects. Consistent with prior research on gender gaps, we find that adjusted pay gaps are positive, meaning that men earn more than observationally equivalent women. Notably, the gap is significantly larger for bonus pay. However, for performance evaluations, the adjusted gender gap is negative suggesting that the pay gap cannot be rationalized by a performance gap. We find that women receive better ratings than observationally equivalent men working the same job. This implies that, conditional on performance evaluations, the adjusted bonus pay gap is even larger.

In the final part of the chapter, we ask whether gender gaps are larger in teams led by male managers. Using a difference-in-difference framework that controls for unobserved worker and manager heterogeneity, we find that this is indeed the case for bonus gaps. This is driven by the fact that performance ratings are more favorable to men if handed out by a male manager. We present suggestive evidence that the relevance of manager gender for pay gaps is driven by (unconscious) discrimination rather than same-gender complementarities in productivity. However, independent of the root cause of these differences in evaluations by manager gender, the findings imply that a higher number of male managers increases gender gaps and thus constitutes a structural disadvantage for women.

Chapter 1

Employment, Output and Welfare Effects of Minimum Wages

1 Introduction

The minimum wage is one of the most frequently used labor market policies in developed countries. In the benchmark model of fully competitive labor markets, wages equal marginal productivity and a binding minimum wage always reduces employment, output and welfare. However, a large body of empirical research has found only very muted employment effects for observed minimum wages ranging between 30 and 60% of the full-time median wage (Kaitz index[1]) (Dube, 2019).[2] In addition, recent evidence shows that minimum wages not only increase earnings but also improve the quality of jobs by reallocating workers towards high-productivity firms and jobs with higher employment levels (Dustmann et al., 2020). Against this backdrop, many countries are discussing proposals to substantially increase the minimum wage, but policy makers lack a comprehensive analysis of the macroeconomic and distributional implications of raising the minimum wage beyond observed levels.[3]

This paper takes a first step towards filling this gap. Specifically, I use a rich search-and-matching model in order to analyze how the minimum wage affects employment, output and welfare. I first estimate the model using German administrative and survey data from 2014 – the year before Germany introduced a federal minimum wage that affected more than ten percent of jobs. Second, I evaluate the macroeconomic and distributional

[1]The Kaitz index is the ratio of the minimum to the full-time median wage multiplied by 100%.

[2]The majority of high-income OECD – including for example the US, Canada, Japan, South Korea or Australia and 21 out of 27 European countries countries – has a minimum wage in place. The Kaitz index varies between 30 and 60 percent. This variation is also present within the US where state-level minimum wages vary between the federal minimum of 7.25 USD and 14 USD.

[3]For example, US Democrats have proposed federal minimum wage of 15 USD (Kaitz index $\approx$ 75%). In Germany, there is a discussion about raising the minimum wage to 12 EUR ($\approx$ 62%). The Polish government plans a 73% increase over and the UK government officials plan to raise the minimum wage to 67% of the median over. The Italian government plans to introduce a minimum wage.

implications of the German minimum wage reform and show that the model is consistent with recent reduced-form evidence on the reform's short-run employment and productivity effects (e.g. Dustmann et al., 2020). Finally, I use the estimated model to quantify the short-run and long-run effects of a hypothetical reform that raises the minimum wages well beyond the current level in Germany.

The analysis is based on a search-and-matching model of the labor market with substantial worker and firm heterogeneity, differences in employment levels, and a progressive tax-and-transfer system. In the model, the effect of the minimum wage on employment is ambiguous since firms' vacancy posting and workers' job search decisions are affected in opposite directions (Flinn, 2006; Acemoglu, 2001). On the one hand, firms will lower their vacancy creation as the minimum wage cuts into match profits. On the other hand, the minimum wage increases wages, earnings and thus the surplus of finding a job, which leads workers to exert more search effort. The net effect on employment is therefore a quantitative question.[4]

In addition to the employment effect, minimum wages also affect output by changing the *composition* of jobs along two dimensions. First, raising the minimum wage increases average productivity because profits and thus vacancy posting decline more strongly for low-productivity firms (Eckstein and Wolpin, 1990; Acemoglu, 2001). Second, raising the minimum wage increases the average employment level, as was recently documented for example by Dustmann et al. (2020). In particular, my model allows for three different employment levels, which I call full-time, part-time and marginal jobs. Differences in employment levels are particularly important as most tax- and transfer systems in developed countries subsidize low-earnings jobs. As a result, low-hours jobs are concentrated in the bottom part of the wage distribution and become relatively less profitable in the presence of a binding minimum wage.[5] The shift towards full-time jobs is amplified by the fact that workers' incentives to search for full-time jobs increase in the hourly wage.

The analysis proceeds in three steps. In a first step, I estimate the model using German administrative linked employer-employee as well as survey data from 2014, i.e. the last year where the economy was not distorted by a federal minimum wage. I show that the estimated model is able to match the joint distribution of wages, firm productivity and employment levels. This is important as it determines the scope for reallocation and thus output effects of increasing the minimum wage. The model also matches the distribution of labor market states across demographics which allows for an analysis of heterogeneous welfare effects.

[4] Acemoglu (2001) briefly discusses this possibility, but does not analyze which channel dominates quantitatively. Bagger and Lentz (2018) allow for endogenous search effort in their analysis of firm-worker sorting. Outside the minimum wage literature, the idea that the surplus of employment affects workers' search effort and employment is standard in the literature on unemployment benefits (e.g. Meyer, 1990; Chetty, 2008; Schmieder et al., 2012; Marinescu and Skandalis, forthcoming).

[5] In Germany in 2014, full-time jobs accounted for only one third of the jobs affected by the initial minimum wage which affected more than ten percent of all jobs.

In the second step, I assess the macroeconomic implications of the introduction of a federal minimum wage in Germany in 2015. I find that the introduction of a minimum wage of 8.5 EUR (Kaitz index of 47%) had negligible employment effects, but led to an increase in average hours worked and firm productivity of 1.4% and 0.6% respectively. The model predicts that this change in the composition of jobs increased output by 0.4% over the first five years and will increase output by almost 0.5% in the long-run. However, I also find that the German tax- and transfer-system prevents consumption growth from keeping up with earnings growth. Higher earnings reduce the level of transfer payments workers receive. Together with higher disutility of longer working hours, this implies that the welfare gains of the reform are negligible. Nevertheless, workers are now less reliant on government transfers to top up their earnings. Importantly, the model's short-run predictions of a null-effect on total employment, a shift from marginal to part-time and full-time jobs, and an increase in average firm productivity are qualitatively and quantitatively consistent with the short-run effects documented by recent reduced-form studies (e.g. Garloff, 2016; vom Berge et al., 2016; Caliendo et al., 2017; Dustmann et al., 2020). The fact that the model is consistent with the reduced-form evidence on a large and observed minimum wage reform lends credibility to the following analysis of counterfactual minimum wage levels.

In the third and most important step, I analyze how raising the minimum wage well beyond observable levels affects employment, output and welfare. Importantly, I analyze not only the new stationary equilibrium but the entire transition path. Focusing on the long-run effects, I find that total employment, i.e. the number of jobs, slightly increases in the minimum wage up until a Kaitz index of 60% (11 EUR) as higher search effort outweighs lower vacancy posting.[6] As the minimum wage increases further, the reduction in vacancies dominates and total employment starts to fall. I further find that raising the minimum wage can substantially increase total output as the share of low-hours and low-productivity jobs monotonically decreases in the minimum wage. Total hours worked are maximized at a Kaitz index of 73% (13.5 EUR). At the long-run output maximum at a Kaitz index of 78% (14.4 EUR), total output is about 3.6% above the baseline level. Average firm productivity and total hours worked are 3.3% and 4.0% above the baseline level respectively. Quantitatively, the increase in average output per job more than offsets the lower number of low-skill jobs whose contribution to total output is relatively small.[7]

In addition to the steady-state analysis, this paper goes beyond the existing literature by analyzing the entire transition path. The results show that short- and long-run effects

[6] While there is no evidence for Germany, Cengiz et al. (2019) find no disemployment effects for past US minimum wage reforms with Kaitz of up to 60%.

[7] I assume that there are no skill-complementarities in production. This assumption is supported by the findings of Cengiz et al. (2019) who demonstrate that the minimum wage elasticity for higher-skilled employment should be very small with a neoclassical production function and plausible parameter values for the elasticity of substitution between low- and high-skill workers. This is mainly driven by the small output share of minimum wage workers. In addition, they find no evidence for labor-labor substitution.

differ significantly. Specifically, a sudden increase in the minimum wage will cause an initial drop in employment even if employment hardly changes in the long-run. The larger the increase in the minimum wage, the more workers initially lose their job because it has become unprofitable for the firm. Search frictions imply that it takes time for the unemployment rate to drop again. This is amplified by the fact that firms now post substantially fewer vacancies. For minimum wages above 60% of the median wage, output declines on impact. At the long-run output optimum, the unemployment rate initially more than doubles and, on average, is about 60% (45%) higher over the first two (five) years after the minimum wage hike. As a result, the economy goes through a recession of almost two years.

Finally, I show that the minimum wage does not benefit all workers equally. Women, who tend to prefer jobs with fewer weekly hours, experience increasing disutility from work as firms offer fewer vacancies for low-hours jobs. This disutility outweighs the utility gains from higher consumption.[8] In addition, low-skill workers become non-employable and are stuck in long-term unemployment as firms will no longer hire them at the minimum wage.

In sum, this paper makes three contributions. First, I incorporate endogenous job search effort, differences in employment levels, and a progressive tax-transfer system into a search-matching model with worker and firm heterogeneity and show that the estimated model matches the joint distribution of wages, firm productivity and employment levels. Second, I use the estimated model to assess the macroeconomic and distributional implications of the introduction of a federal minimum wage in Germany in 2015 and show that the model is consistent with the available, empirical evidence. Third and most importantly, I provide a comprehensive analysis of the short- and long-run impact on employment, output and welfare of raising the minimum wage beyond observable levels.

Related Literature. My paper speaks to several strands of the literature. Most importantly, my paper adds to the large literature investigating the effects of minimum wages in labor markets with search frictions. Some early contributions assume that contact rates are exogenously given and not affected by the minimum wage in wage posting models (Burdett and Mortensen, 1998; Bontemps et al., 1999; van den Berg and Ridder, 1998). Both Eckstein and Wolpin (1990) and Acemoglu (2001) allow for endogenous vacancy creation and show theoretically that a minimum wage induces a trade-off between the total number of jobs and their average productivity. Flinn (2006) estimates a stylized search-matching model with endogenous contact rates in which the employment effect of minimum wages need not be negative even though firms can adjust vacancy posting. Engbom and Moser (2018) estimate a wage-posting model with worker and firm heterogeneity as well as endogenous vacancy creation in order to quantify the contribution of an increase

[8] Note that I interpret this disutility as a rather general proxy capturing not only the utility of leisure but also outside constraints such as childcare obligations.

in the minimum wage to the decline of wage inequality in Brazil. In simultaneous and independent work, Blömer et al. (2020) estimate the wage posting model by Bontemps et al. (1999) to analyze minimum wage effects on steady state full-time employment in Germany.[9] I contribute to this literature by quantifying employment effects when both vacancy posting and search effort are optimally chosen by firms and workers.[10] In addition, my paper is the first to analyze how minimum wages affect output when jobs differ not only by firm productivity but also employment level.[11] My model also differs by allowing for a progressive tax- and transfer system that subsidizes low-earnings jobs, as is the case in most developed countries. This is important for our understanding of reallocation effects as it shapes the joint distribution of employment levels and wages. Finally, this is the first paper to analyze transition dynamics of minimum wage hikes and show that policy makers have to weigh long-run output and welfare gains of higher minimum wages against severe short-term unemployment.

I further contribute to the literature evaluating past minimum wage reforms which mostly consists of reduced-form papers. Harasztosi and Lindner (2019) analyze who pays for the minimum wage in Hungary. Portugal and Cardoso (2006) and Dube et al. (2016) show that minimum wages reduce employment flows. Cengiz et al. (2019) provide an extensive analysis of employment effects of past minimum wage reforms in the US. The short-run effects of the German minimum wage reform of 2015 is analyzed most prominently by Dustmann et al. (2020) as well as e.g. Garloff (2016), Caliendo et al. (2017), Holtemöller and Pohle (2017), and Burauel et al. (2020). This paper's structural approach is able to add a macroeconomic perspective by analyzing output and welfare effects. In addition, the model with endogenous search effort can rationalize why reduced-form studies have not found significant disemployment effects even for high levels of the minimum wage (e.g. Cengiz et al., 2019; Dustmann et al., 2020).

Finally, by including endogenous search effort, my paper is also related to the literature on employment effects of other labor market policies that target the surplus of employment. The large literature on unemployment benefits has worked to understand how benefits or

[9]Apart from my aforementioned general contributions, my paper differs along a number of dimensions. First, Blömer et al. (2020) only focus on full-time employment which accounts for only one third of the jobs affected by the initial minimum wage of 8.5 EUR and less than half of the jobs between 8.5 and 12.5 EUR in 2014. I show that taking non-full-time work into account is important for understanding the reallocation effects of the minimum wage. Second, they do not discuss transition dynamics which I show to be of great importance when assessing a minimum wage hike. Thirdly and most importantly, they focus solely on the unemployment rate while my paper provides a joint analysis of employment, output and welfare effects of increasing the minimum wage.

[10]While Acemoglu (2001) theoretically shows that endogenous search effort can mute disemployment effects of minimum wages, he does not quantify the contribution of this channel in an estimated model. Bagger and Lentz (2018) use endogenous search effort to explain sorting of workers and firms. Krebs and Scheffel (2016) use a search matching model with endogenous search to evaluate the German Hartz reforms.

[11]A recent paper by Doppelt (2019) shows theoretically and using reduced form evidence that higher minimum wages lead workers to work longer hours. However, he does not embed the mechanism in a quantitative model to analyze output effects.

benefit duration affect employment by influencing workers' incentives to exert search effort and find a job (e.g. Ljungqvist and Sargent, 1998; Chetty, 2008; Krebs and Scheffel, 2013; Schmieder et al., 2016). There is also a literature in macroeconomics analyzing unemployment insurance policies over the business cycle in search-matching models (e.g. Mortensen and Pissarides, 1994; Krause and Uhlig, 2012; Hagedorn et al., 2019; Mitman and Rabinovich, 2019). While these papers study how the surplus of employment evolves when unemployment benefits change, I analyze how minimum wages affect employment because the value of employment is affected by the minimum wage.[12] My paper thus suggests that unemployment benefits and minimum wages interact and should potentially be set jointly.

Outline. The remainder of the paper is structured as follows. Section 2 presents the equilibrium search-matching model. Section 3 describes the parameterization, identification and estimation of the model and evaluates the model fit. Section 4 analyzes the introduction of the German minimum wage. Section 5 analyzes counterfactually high minimum wages. Finally, section 5 discusses the results and concludes.

2 Model

I study an economy where a unit mass of workers meet a mass m_f of firms in a labor market with search frictions. Time is discrete and both workers and firms are infinitely-lived. Workers differ by human capital and demographics, and firms differ by productivity. Both worker and firm heterogeneity is exogenous and time-invariant.

2.1 Workers

Workers differ by gender and family status. In particular, I distinguish between the following five sociodemographic groups indexed by j: married men, single men, single women with and without kids, and married women (see Table 1).[13] Let P_j denote the population share of group j. A worker's sociodemographic type determines her preferences over employment levels as well as her tax-and transfer schedule.[14]

Workers further differ by their time-invariant human capital (skill) h. The gender-specific distribution function of human capital is $\Phi^{g(j)}$ where g is the gender of group j. I assume that the labor market is segmented with respect to workers' skill levels such that there is a continuum of independent labor markets – one for each level of h (van den Berg and Ridder, 1998; Engbom and Moser, 2018).

[12] In a recent paper by Hartung et al. (2018), the value of unemployment not only affects job finding rates but also separation rates as it leads workers to accept lower wages in return for greater job stability.

[13] As men with and without children are similar with respect to all targeted moments, I only distinguish between single and married men. The same holds for married women.

[14] Whenever possible, I will drop the subscript j for worker types to improve readability.

Table 1: Sociodemographic Types

	$P_j = \Pr(j)$	$\Pr\big(g(j)\big)$	$\Pr\big(j \vert g(j)\big)$
Sociodemographics			
Men, Single	0.214	0.514	0.416
Men, Married	0.300	0.514	0.584
Women, Single, No Kids	0.168	0.486	0.346
Women, Single, Kids	0.046	0.486	0.095
Women, Married	0.272	0.486	0.560

Notes: The share of each sociodemographic group conditional on gender $g(j)$ is computed from the SOEP and then multiplied by the respective gender share in the SIAB data. Source: SOEP, SIAB, own calculations.

A type-j worker with human capital h can be employed, $s = e$, short-term unemployed, $s = su$ or long-term unemployed, $s = lu$. There are three *employment levels* which I label full-time ($x = ft$), part-time ($x = pt$) and marginal employment ($x = mj$). In addition, jobs differ with respect to the employer's productivity p which will be described below. While short-term unemployed workers receive unemployment insurance proportional to their previous earnings, all long-term unemployed workers receive the same unemployment benefits, i.e. a subsistence minimum. In sum, for each skill level h there is a continuum of idiosyncratic states for employed and short-term unemployed workers and a single state for long-term unemployment. The state space of a type-j worker with human capital h is

$$\mathcal{S} = \Big\{\big\{(s, x, p) \mid s \in \{e, su\}, x \in \{ft, pt, mj\}, p \geq 1\big\}, lu\Big\}$$

In the following I denote by σ one point in the state-space of a worker ($\mathcal{S}$) and F the distribution of endogenous states (given j and h).

When a worker with human capital h works a type-x job at a firm with productivity p, the match output is $f(h, x, p) = e_x a_x hp$ for $x \in \{ft, pt, mj\}$. The parameters e_x denote the hours worked in full-time, part-time and marginal jobs respectively.[15] The parameters $a_x > 0$ allow for constant productivity differences between full-time, part-time and marginal jobs. Workers earn a fixed and exogenous share $r \in (0, 1)$ of the match

[15] I normalize hours worked in full-time employment to one.

output.[16] In the presence of a minimum wage $\bar{w}$, the hourly wage is

$$w(h,x,p) = \max\{rf(h,x,p), \bar{w}\} \tag{1}$$

Gross earnings and net earnings are given by

$$\begin{aligned} \tilde{y}(h,x,p) &= e_x w(h,x,p) \\ y^j(h,x,p) &= \tilde{y}(h,x,p) - T^j(\tilde{y}(h,x,p)) \end{aligned} \tag{2}$$

where $T^j(\tilde{y})$ is a tax function that depends on the worker's sociodemographics.[17]

Short-term unemployed workers receive a share b of their previous net earnings up to a maximum amount of B_{max} (unemployment insurance). Long-term unemployed workers receive subsistence benefits B_{min} independent of their skill level or previous earnings. Short-term unemployment insurance is capped from below by B_{min}. Employed workers are also eligible for unemployment benefits to top up their net earnings or unemployment insurance. In doing so, a share τ_{top} of net earnings will be deducted from B_{min}. Finally, married workers receive non-labor income y^j_{free} which is always deducted from B_{min}.[18] Hence, subsistence benefits for type-j workers may not exceed $B^j_{min} = \max\{B_{min} - y^j_{free}, 0\}$.

As there is no savings device, consumption c equals net income. A type-j worker with skill h faces the following consumption schedule

$$c^j(h,\sigma) = \begin{cases} y^j(h,x,p) + \max\left\{B^j_{min} - \tau_{top} y^j(h,x,p), 0\right\} + y^j_{free} & \text{if } s = e \\ by^j(h,x,p) + \max\{B^j_{min} - by^j(h,x,p), 0\} + y^j_{free} & \text{if } s = su \\ B^j_{min} + y^j_{free} & \text{if } s = lu \end{cases} \tag{3}$$

where $\sigma \in \mathcal{S}$ denotes one state in the worker's state space.

[16] There are a number of reasons for not using a more involved wage setting mechanism such as Nash bargaining (Cahuc et al., 2006) or wage posting Burdett and Mortensen (1998). First, not having to solve for a wage-posting schedule or bargained wage keeps the estimation of the model feasible as the combination of endogenous worker search effort, and multiple worker types and employment levels makes the computation of the equilibrium time-consuming. Second, match-level wage determination in search-matching models remains a black box and little is known about the validity of the wage-posting or bargaining assumptions. While certainly too simple, the assumption of an exogenous piece rate ensures that (i) I match the aggregate labor share and (ii) the results are not driven by a poorly-understood mechanism. Third, recent evidence by Jäger et al. (2020) shows that – even for previously unemployed workers – wages are insulated from the value of non-employment. Fourth, a recent paper by Di Addario et al. (2020) finds that a core prediction of the sequential auction model (Postel-Vinay and Robin, 2002; Bagger et al., 2014) is not supported by Italian social security data. In particular, the productivity of the firm where the worker is poached/hired from has almost no effect on the wage at the destination firm. Fifth, wage posting implies substantial wage spillover which have not been found by Cengiz et al. (2019) and Dustmann et al. (2020). Sixth, a fixed piece-rate could be motivated by Nash bargaining over the match output instead of the match surplus.

[17] I refer to taxes as the sum of income taxes and social security contributions.

[18] The type-specific and exogenous non-labor income y^j_{free} represents a share of the partner's income for married workers. Singles do not receive such non-labor income.

Workers exert costly search effort ℓ to find (better) jobs in their skill-segment of the labor market. A worker in employment state s meets a vacancy with probability

$$\lambda_\sigma(\ell|h) = \phi^\sigma \ell \Lambda(\theta_h) \tag{4}$$

where labor market tightness θ_h is taken as given and ϕ^σ is a search efficiency parameter. I will assume that search efficiency differs by employment level and between short- and long-term unemployed $(\phi^{su}, \phi^{lu}, \phi^{e_x})$. Importantly, not every meeting has to result in a match as search cannot be directed towards certain employment levels or high-productivity firms, and workers may decline lower-valued offers.

The mass of search-weighted workers of type-j is denoted by $S^j(h)$ and the mass of all search-weighted workers in skill segment h is

$$S(h) = \sum_j \underbrace{P_j \int_\sigma \phi^\sigma \ell(\sigma|j,h) \mathrm{d}F(\sigma|j,h)}_{S^j(h)} \tag{5}$$

where $\ell(\cdot|j,h)$ and $F(\cdot|j,h)$ represent the optimal search effort and stationary distribution functions for type-j workers in skill segment h.

Workers' utility depends on consumption, the employment level and job search:

$$u^j(\ell|h,\sigma) = \tilde{u}(c^j(h,\sigma)) - d(\ell) + \nu^j\big(x(\sigma)\big) \tag{6}$$

Thereby, $\tilde{u}(c)$ is a concave flow utility function of consumption, $d(\ell)$ is a convex search cost function and $\nu^j(x(\sigma))$ captures the (dis-)utility of different employment levels relative to nonemployment. The latter may depend on workers' sociodemographics j. Single women with kids may for example have a strong preference for part-time or marginal jobs.[19] Heterogeneity in $\nu^j(x)$ will allow the model to match the joint distribution of employment levels and sociodemographics.

2.2 Firms

There is a mass m_f of risk-neutral firms with heterogeneous productivity $p \sim \Gamma$. Firms employ workers of all skill levels h at all employment levels x. I assume that firms operate a linear production technology such that total output of a firm with productivity p is the

[19] I emphasize that these "preference" parameters not only capture the tastes for leisure, but also exogenous constraints such as childcare obligations. As I do not explicitly model policies affecting child care constraints, using such a proxy is justified even though the parameter is not policy-invariant outside the model.

sum of the match outputs

$$\sum_x \int_{\underline{h}}^{\bar{h}} f(h,x,p)L(h,x,p)\mathrm{d}h$$

where $L(h,x,p)$ is the firm's mass of employees with skill h and demographics j working a type-x job. This implies that there are no complementarities between low- and high-skill workers.[20]

Firms attract workers for type-x jobs in skill segment h by posting vacancies $v(h,x)$ at a convex cost $\kappa_x(h,v)$. As hiring a worker does not affect future recruitment, firms will not reject workers of a particular demographic type even if different workers are more or likely to switch employers than others. Denote by $N(h,x)$ the mass of type-x vacancies in skill segment h and the total number of vacancies as $N(h) = \sum_x N(h,x)$. In addition, let $\Psi(h)$ denote the distribution of employment levels and productivities among all vacancies in skill segment h. Firms' vacancy posting response to a binding minimum wage can affect both the $N(h)$ and $\Psi(h)$. The former impacts labor market tightness, job finding probabilities and the total number of jobs. The latter will determine the composition of jobs and thus the average productivity and employment level.

2.3 Labor Market

Recall that labor markets are segmented by worker skill h and workers cannot direct search towards a certain employment level or towards high-productivity firms. Hence, the total mass of search and vacancies in a skill-segment are matched by the matching function

$$M(h) = N(h)^{\xi} S(h)^{1-\xi} \tag{7}$$

where ξ is the elasticity of matches with respect to the mass of posted vacancies. Labor market tightness is defined as

$$\theta(h) = \frac{N(h)}{S(h)} \tag{8}$$

and the aggregate contact rates for a unit of search and a vacancy are $\Lambda(\theta) = \theta^{\xi}$ and $\Pi(\theta) = \theta^{\xi-1}$ respectively.

Employment relationships are terminated for three mutually exclusive reasons. First, workers may voluntarily change firms and/or employment levels as a result of on-the-job

[20]This assumption is rather standard in papers studying frictional labor markets (e.g. Bagger et al., 2014; Bagger and Lentz, 2018). The assumption is also supported by the findings of Cengiz et al. (2019) who demonstrate that the minimum wage elasticity for higher-skilled employment should be very small with a neoclassical production function and plausible parameter values for the elasticity of substitution between low- and high-skill workers.

search. In equilibrium, firms with low productivity will be more likely to experience this event.

Second, workers may be hit by a so-called Godfather shock which forces them to switch to a different job that is randomly drawn from the distribution of vacancies. This is important to account for the substantial share of job-to-job transitions that are accompanied by a wage cut and cannot be explained by on-the-job search (Jolivet et al., 2006). The Godfather shock arrives with probability $\pi_{e|e_x}(h) = \psi_x \Lambda(\theta)$ and captures involuntary and unintended job-to-job transitions unrelated to workers' search effort. These may be the result of firms' outplacement programs, workers' search effort after an advance-notice layoff or family-related events that force workers to move and look for a new job immediately.

Third, matches can be destroyed such that the worker transitions into short-term unemployment. This happens with probability $\pi_{su|e_x}$ and if a minimum wage hike makes the match unprofitable for the firm.

2.4 Worker Problem

Workers choose search effort ℓ and reject or accept job offers in order to maximize discounted lifetime utility. Labor market tightness and the distribution of vacancies are taken as given.

The value of long-term unemployment for a type-j worker with human capital h solves the following Bellman equation:

$$\begin{aligned} V_{lu}^j(h) = \max_{\ell} \Big\{ & u^j\big(\ell, h, lu\big) + \beta \lambda_{lu}(\ell|h) \mathbb{E}_{(x,p)} \Big[\max\big\{ V_e^j(h,x,p), V_{lu}^j(h) \big\} \big| h \Big] \\ & + \beta\big(1 - \lambda_{lu}(\ell|h)\big) V_{lu}^j(h) \Big\} \end{aligned} \tag{9}$$

Search effort ℓ is associated with lower flow utility but a higher probability of meeting a firm. Upon on meeting a firm offering a (x, p) job, the worker accepts the job if and only if the value of the employment relationship, $V_e^j(h, x, p)$, exceeds the value of remaining long-term unemployed. The max-operator in the continuation value captures this acceptance decision. The expectation is taken with respect to the distribution of vacancies in the worker's skill segment. With probability $1 - \lambda_{lu}(\ell|h)$, the worker does not meet a firm and remains long-term unemployed.

The value of short-term unemployment when the previous job was of type x at a type-p firm is

$$\begin{aligned} V_{su}^j(h,x,p) = \max_{\ell} \Big\{ & u^j\big(\ell|h, (su,x,p)\big) + \beta \pi_{lu|su} V_{lu}^j(h) \\ & + \beta \lambda_{su}(\ell|h) \mathbb{E}_{(x',p')} \Big[\max\big\{ V_e^j(h,x',p'), V_{su}^j(h,x,p) \big\} \big| h \Big] \\ & + \beta\big(1 - \pi_{lu|su} - \lambda_{su}(\ell|h)\big) V_{su}^j(h,x,p) \Big\} \end{aligned} \tag{10}$$

The only difference to long-term unemployment is that the worker transitions from short- to long-term unemployment with exogenous probability $\pi_{lu|su}$.

The value of a worker employed at a type-p firm on a type-x job is

$$\begin{aligned} V_e^j(h,x,p) = \max_{\ell} \Big\{ & u^j\big(\ell|h,(e,x,p)\big) + \beta\pi_{su|e_x} V_{su}^j(h,x,p) \\ & + \beta\lambda_{e_x}(\ell|h)\mathbb{E}_{(x',p')}\Big[\max\big\{V_e^j(h,x',p'),V_e^j(h,x,p)\big\}\big|h\Big] \\ & + \beta\pi_{e|e_x}(h)\mathbb{E}_{(x',p')}\Big[V_e^j(h,x',p')\big|h\Big] \\ & + \beta\big(1-\pi_{su|e_x}-\lambda_e(\ell|h)-\pi_{e|e_x}(h)\big)V_e^j(h,x,p)\Big\} \end{aligned} \tag{11}$$

Employed workers become short-term unemployed with probability $\pi_{su|e_x}$, receive a job offer that they can decline through on-the-job search with probability $\lambda_{e_x}(\ell|h)$ and are involuntarily reallocated to a different job with probability $\pi_{e|e_x}(h)$.

All workers may have an incentive to search for a (better) job. The first order condition determining optimal search effort is given by

$$\frac{\mathrm{d}d^j(\ell)}{\mathrm{d}\ell} = \beta\frac{\partial\lambda_\sigma(\ell|h)}{\partial\ell}\Bigg(\underbrace{\mathbb{E}_{(x,p)}\Big[\max\big\{V_e^j(h,x,p),V^j(h,\sigma)\big\}\big|h\Big]-V^j(h,\sigma)}_{\text{expected surplus of meeting a firm}}\Bigg) \tag{12}$$

For a worker in state σ, the job finding probability is the result of optimal search effort $\ell(\sigma)$ as well as the worker's acceptance decision

$$\pi^j(\ell|h,\sigma) = \lambda_\sigma\big(\ell|h)\mathbb{E}_{(x,p)}\Big[\mathbb{1}\{V_e^j(h,x,p)>V^j(h,\sigma)\}\big|h\Big] \tag{13}$$

2.5 Firm Problem

Firms maximize expected discounted profits taking as given labor market tightness, the distribution of vacancies and the distribution of workers' search effort. As total production is additive in h and x, the firm faces a sequence of independent optimization problems – one for each (h,x)-segment. Each period, firms post vacancies which may result in an employment relationship starting in the subsequent period. Unfilled vacancies are not carried over to the next period but have to be re-posted. Additive production combined with the fact that the cost of posting vacancies is independent of the current workforce further implies that the firm's optimal amount of vacancies is independent of the current workforce. For the same reasons, firms will not reject workers of a particular demographic type.

A type-x employment relationship with a type-j employee may be dissolved either due to exogenous job destruction, a Godfather shock or on-the-job search with probability:

$$\delta^j(h,x,p) = \pi_{su|e_x} + \pi_{e|e_x}(h) + \pi^j(\ell(\sigma)|h) \tag{14}$$

The probability of filling a vacancy is equal to the aggregate contact rate times the probability that the contacted worker accepts the offer:

$$\eta(h,x,p) = \Pi(\theta_h)\frac{S(h,x,p)}{S(h)} \tag{15}$$

Thereby, $S(h)$ is the total search-weighted mass of workers in skill segment h and $S(h,x,p)$ is the mass of search-weighted workers in segment h willing to accept a type-x job at a firm with productivity p:

$$S(h,x,p) = \sum_j S^j(h,x,p) \tag{16}$$

$$S^j(h,x,p) = P_j \int_\sigma \phi_\sigma \ell(\sigma|j,h)\mathbb{1}\{V_e^j(h,x,p) > V^j(h,\sigma)\}\mathrm{d}F(\sigma|j,h) \tag{17}$$

Let $(1-r^+)$ be the firm's profit share of the match output. If the minimum wage is binding for a (h,x,p)-job, $(1-r^+)$ is lower than the baseline profit share, $(1-r)$. Given r^+, the value $W^j(h,x,p)$ of a type-x employment relationship with a worker of type j in segment (h,x) for a firm with productivity p is given by

$$\begin{aligned} W^j(h,x,p) &= \underbrace{(1-r^+)f(h,x,p)}_{\text{flow profit}} + \beta^f\big(1-\delta^j(h,x,p)\big)W^j(h,x,p) \\ &= \frac{(1-r^+)f(h,x,p)}{1-\beta^f\big(1-\delta^j(h,x,p)\big)} \end{aligned} \tag{18}$$

where β^f is the firms' discount factor. When posting a vacancy, the firm has to take the expectation over worker types as they differ in their on-the-job search effort which affects the separation probability and expected value of a match. The ex-ante expected value of filling a vacancy is thus

$$\begin{aligned} \mathbb{E}\big[W(h,x,p)\big] &= \sum_j \frac{S^j(h,x,p)}{S(h,x,p)}W^j(h,x,p) \\ &= (1-r^+)f(h,x,p)\underbrace{\sum_j \frac{S^j(h,x,p)}{S(h,x,p)}\frac{1}{1-\beta^f\big(1-\delta^j(h,x,p)\big)}}_{\text{discounted expected match duration}} \end{aligned} \tag{19}$$

Knowing the expected value of an employment relationship, the optimal number of vacancies has to satisfy

$$\kappa'(v,h,x) = \beta^f\eta(h,x,p)\mathbb{E}\big[W(h,x,p)\big] \tag{20}$$

Optimal vacancy posting then requires firms to post vacancies until the marginal cost of posting another vacancy is equal to the discounted expected value of an employment relationship weighted by the probability of filling the vacancy.

2.6 Equilibrium

A stationary equilibrium consists of value functions, $V_{lu}^{j}(h)$, $V_{su}^{j}(h,x,p)$, $V_{e}^{j}(h,x,p)$, search effort policy functions, $\ell^{j}(h,\sigma)$, vacancy posting policy functions, $v(h,x,p)$, labor market tightness, $\theta(h)$, distribution of vacancies, $\Psi(h,x,p)$, and a distribution of workers across states, $F^{j}(h,\sigma)$, that satisfy the following conditions. First, given labor market tightness and the distribution of vacancies, the value and search effort policy functions solve the workers' problem. Second, given labor market tightness, the distribution of vacancies, workers' search policies and the distribution of workers across states, firms' vacancy posting policy functions solve the firms' optimality conditions. Third, the distribution of workers across states is stationary. That is, given the economy starts at this distribution and given the policy functions and labor market tightness, the distribution of workers across states will not change.

3 Estimation

In this section, I first describe the pre-set parameters and parameterize workers' flow utility and skill distributions, firms' productivity distribution and vacancy posting cost function and the tax schedule (section 3.1). Second, I discuss which moments I target in the method of simulated moments in order to identify the jointly estimated parameters (section 3.2). Third, I evaluate the estimation results and model fit (section 3.4).

3.1 Parameterization and Pre-Set Parameters

One period in the model corresponds to one quarter. I set the quarterly discount factor of both workers and firms equal to $\beta = 0.98$ and choose the minimum wage of EUR 8.5 as the numéraire.

The employment level for full-time employment, e_{ft} is normalized to one and e_{pt} and e_{mj} are set to match the ratio of average weekly hours of part-time and marginal workers relative to full-time employed workers reported by Dustmann et al. (2020) who have access to hours worked in the German security data. This yields $e_{pt} = 0.615$ and $e_{mj} = 0.223$.

I set $r_{ft} = r_{pt} = 0.62$ which approximately match the aggregate labor share in Germany between 2010 and 2014. The labor share for marginal jobs r_{mj} is estimated and allowed to be lower in order to match the joint distribution of wages and employment levels. As marginal jobs constitute a tiny share of the aggregate wage bill, this does not affect the labor share significantly.

The German transfer system distinguishes between short- and long-term unemployment. During the first year of unemployment, workers are paid a fixed fraction $b = 0.6$ of their previous earnings (ALG I), but not less than the subsistence minimum B_{min}. With a constant net replacement rate for short-term unemployed workers, benefits differ by previous earnings. Long-term unemployed workers receive the subsistence minimum B_{min} independent of their previous earnings (ALG II). I set the policy parameter B_{min} to 800 EUR which corresponds to about 55% of of full-time monthly earnings at the minimum wage of 8.5 EUR. For employed workers, 80% of their net earnings is deducted from the amount of subsistence benefits they are eligible to receive on top of their earnings ($\tau_{top} = 0.8$). Hence, all workers with monthly net earnings of at least 1,000 EUR are not eligible for top-up transfers. Workers with net earnings below this threshold are eligible for subsistence transfers if they do not receive non-labor income y^j_{free} from their spouse. Using SOEP data that allow me to link spouses, I calculate average net earnings of the spouses of the married men and women in my sample. I then assign half of that amount to the spouse as non-labor income. On average, married women have roughly EUR 894 and married men EUR 409 in non-labor income from their spouses' net earnings. In the model, non-labor income is deducted from subsistence benefit eligibility. With $B_{min} = 800$, this implies that married women are not eligible for subsistence benefits and married men receive at most half of total subsistence benefits. Singles are assumed to have no non-labor income and are hence eligible for the full amount of subsistence benefits.

Gross earnings are subject to taxation. Note that I refer to the sum of taxes and social security contributions simply as taxes. I assume that workers pay a constant marginal tax rate τ^j on earnings above an exemption level D^j.

$$y_{net} = \min\{y_{gross}, D^j\} + (1 - \tau^j) \max\{0, y_{gross} - D^j\} \tag{21}$$

and estimate the parameters on SOEP data for gross and net earnings for the years 2013 and 2014 separately for different socioeconomic types. Figure 1 shows that the estimated average tax function provides a good fit to the binned data.

I assume that firm productivity $p \geq 1$ is drawn from a Log Gamma distribution with shape and scale parameters α and θ. Productivity differences across job types are governed by $a_{pt}, a_{mj} \in (0, 1]$ with a_{ft} normalized to one. Human capital is drawn from a gender-specific left-truncated Log Normal distribution defined by μ^g_h and σ^g_h. The truncation bound h_{min} is chosen such that the lowest possible wage – resulting from a match between the least productive firm ($p_{min} = 1$) and lowest skilled worker generates a wage of 4 EUR, i.e. $rh_{min}p_{min}a_{mj} = 4$. Data from the SOEP as well as the German Survey of Earnings Structure show that there are virtually no jobs with an hourly wage below 4 EUR (Minimum Wage Commission, 2018).

FIGURE 1: Fit of Estimated Tax Functions

Notes: This figure shows estimated average tax functions as well as the mean average tax rate in various gross earnings bins. The spikes show the range between the 10^{th} and 90^{th} percentile of average tax rates in those bins. The average tax function is $T(y) = (1 - \tau^j)\max\{0, y - D^j\}/y$.

Workers' utility depends on consumption, job search and the employment level in the following way:

$$u^j(\ell|h,\sigma) = \frac{c^j(h,\sigma)^{1-\gamma_c}}{1-\gamma_c} - \ell^\zeta + h^\epsilon \sum_x \gamma_x^j \mathbb{1}\{x(\sigma) = x\} \tag{22}$$

where $\zeta^j > 1$ and γ_x^j are constants that capture the (dis-)like for the different employment levels (relative to nonemployment) for type-j workers. The state-specific constants will allow the model to match the distribution over employment levels for each sociodemographic group. The state-constants are scaled by h^ϵ where $\epsilon > 0$ implies that the absolute importance of the state-(dis-)utilities grows with human capital. The value of ϵ may helps to match the joint distribution of wages and employment levels.[21]

[21]For example, if flow utility of consumption is linear $\gamma_c = 0$, $\gamma_{pt} > \gamma_{ft}$ and $\epsilon = 0$, the surplus of part-time work over full-time work will be larger smaller for high-skill workers compared to low-skill workers resulting in relatively more part-time jobs in the lower skill segments.

Finally, I assume that the cost of posting v vacancies for type-x jobs in skill segment h is given by

$$\kappa(v,h,x) = \underbrace{e_x \kappa_1}_{\equiv \kappa_{1x}} v^{\kappa_2^x} f(h)^{1-\kappa_2^x} \tag{23}$$

where f is the density of workers' human capital and e_x is the employment level.[22] The convexity of the cost function may depend on the job type. I scale the cost of posting vacancies by the density of human capital due to the assumption of segmented labor markets. This implies that optimal vacancy creation satisfies

$$v(h,p,x) = \left(\frac{(1-r^+) f(x,h,p) A(h,p,x)}{\kappa_{1x} \kappa_2^x} \right)^{\frac{1}{\kappa_2^x - 1}} f(h) \tag{24}$$

where $A(h,x,p)$ is a term depending on the hiring probability and the discounted expected match duration. The elasticity of vacancy creation with respect to the profit share is $1/(\kappa_2^x - 1)$.

3.2 Estimation Strategy

The remaining structural parameters will be estimated using the simulated method of moments to match important aspects of the German labor market in 2014. I estimate the model using a two-step multiple-restart procedure similar to the TikTak-estimation method proposed by Arnoud et al. (2019). In the first stage, I search a compact parameter space by evaluating the objective function at about three million quasi-random Sobol points. I then select the best three thousand points as starting points for local minimizations and pick the local minimizer with the lowest local minimum as the global minimizer.

The parameters to be jointly estimated are the gender-specific skill distribution parameters (α^g, θ^g), the firm productivity distribution parameters (μ_p, σ_p), the sociodemographic-specific preference parameters (γ_x^j), the type-independent preference parameters $(\gamma_c, \zeta, \varepsilon)$, the search efficiency parameters $(\phi^{su}, \phi^{lu}, \phi^{e_x})$, the vacancy cost parameters (κ_1, κ_2^x), the mass of firms (m_f), the probability of becoming long-term unemployed $(\pi_{lu|su})$, and the labor share of marginal jobs (r_m).

To inform these parameters, I target (a) the joint distribution of labor market states and sociodemographics, (b) average and sociodemographic-specific job finding rates out of unemployment, (c) the average elasticity of job finding probabilities with respect to unemployment insurance for short-term unemployed workers, (d) job-to-job transition probabilities conditional on employment level, (e) selected wage quantiles conditional on gender and employment level, (f) the distribution of gender and employment levels in

[22] This functional form is similar to those used in Shephard (2017) and Engbom and Moser (2018).

selected wage groups, (g) selected quantile ratios of the gender-specific distributions of worker fixed effects of full-time workers, (h) selected quantile ratios of the distribution of full-time clustered firm fixed effects weighted by the number workers in each employment level, (i) the standard deviation of the log of full-time firm size, and (j) the aggregate job vacancy rate. While all of the parameters are jointly identified by all moments, I will provide intuition for the selection of moments.

In the absence of a minimum wage, the wage equation in my model is very simple. As in Abowd et al. (1999) (henceforth AKM), the wage w of a full-time worker employed at firm with productivity p is log-additive in her skill h and the firm's productivity

$$\log(w) = \log(r) + \log(h) + \log(p) \tag{25}$$

where r is the exogenous piece-rate. I estimate the empirical distribution of worker and firm-class fixed effects using a clustered AKM approach (Bonhomme et al., 2019). In particular, I first cluster firms based on their wage distributions and use firm-class fixed effects instead of firm fixed effects. See Appendix B for details.

To inform the parameters of the skill and productivity distributions, I target selected quantile ratios of the distribution of worker (by gender) and firm fixed effects for full-time workers as well as selected quantile ratios of the distribution of full-time firm fixed effects weighted by the number of part-time and marginal jobs.

Apart from the fixed effects distributions, I target selected quantiles of the gender-specific wage distributions and the overall wage distributions of full-time, part-time and marginal workers. Explicitly targeting the wage distribution is important as the model needs to be able to replicate the pre-reform distribution of wages and employment levels as well as possible.

The search efficiency parameters are closely related to the average job finding probability of short- and long-term unemployment as well as the probability of job-to-job transitions conditional on the current employment level.

The (dis-)utility parameters γ^j_{ft}, γ^j_{pt} and γ^j_{mj} drive heterogeneity in employment status across sociodemographics. The curvature-parameter ζ in the disutilty of job search affects the elasticity of job search with respect to the surplus of employment. Based on the quasi-experimental literature on the UI-elasticity of job finding probabilities I target an average elasticity of 0.5 across all workers (e.g. Chetty, 2008; Schmieder et al., 2012).

The scale parameter κ_1 affects the overall labor market tightness by making vacancies more or less costly and is thus related to the job vacancy rate. The curvature parameters κ_2^x affect the share of type-x jobs across skill-segments and hence across the wage distribution. Increasing κ_2^{mj} relative to κ_2^{ft} will lead to more type-x vacancies in low skill segments as type-x vacancy posting becomes more inelastic with respect to the expected value of vacancy which in turn tends to increase in h. Moreover, decreasing κ_2^{ft} will make it easier for more productive firms to grow large relative to unproductive firms such that the standard deviation of the log of full-time firm size increases. The curvature parameters are thus informed by both the share of part-time and marginal jobs across the wage distribution as well as the standard deviation of the log of full-time firm size.

3.3 Data

The main data source is a 2% sample of administrative social security records of German workers (SIAB) from 2011 to 2014. The SIAB is a linked employer-employee data set containing information on daily earnings and employment levels (full-time, part-time and mini-job). Sociodemographic characteristics (apart from gender and age) are only available for nonemployed workers. I thus complement it with survey data from the German Socioeconomic Panel (SOEP) which contains annual information on more than 15 thousand workers. For firm-level moments I use administrative data from the Establishment History Panel and the Job Vacancy Survey of the Institute for Employment Research (IAB). I focus on prime-aged workers aged 25 to 60.

3.4 Estimation Results

The model parameters are reported in Table 2 and 3. The skill distribution of men has a higher mean but lower standard deviation than that of women. Figure 2 show the distributions of human capital and firm productivity.

FIGURE 2: Estimated Human Capital and Firm Productivity Distributions

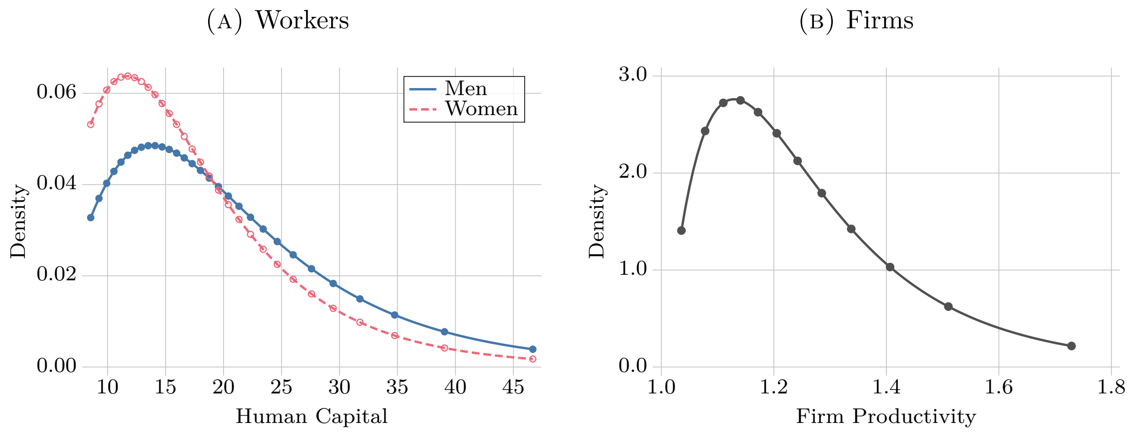

Notes: This figure shows the density of the estimated human capital distributions of workers (by gender) and firm productivity distribution. All distributions are truncated log normal distributions. The markers refer to the grid points used to discretize the distribution.

Table 2 shows that, apart from married men, workers receive utility from working fewer hours as $\gamma_f^j < \gamma_p^j < \gamma_m^j$. All women have a higher preference for part-time and marginal jobs. Single women with kids receive the highest disutility from working full-time. The

Table 2: Worker Parameters

Name	Description	Value	Source
All Workers			
β	Discount factor	0.980	–
γ_c	CRRA parameter	0.727	estimated
ζ_2	Search disutility (convexity)	2.056	estimated
ϵ	Relation btw. h and state utilities	0.173	estimated
Skill Distribution of Men			
μ	Mean of $\log(h)$	2.920	estimated
σ	Std. dev. of $\log(h)$	0.542	estimated
Skill Distribution of Women			
μ	Mean of $\log(h)$	2.725	estimated
σ	Std. dev. of $\log(h)$	0.517	estimated
Men, Single			
γ^j_{ft}	State utility of $s = ft$	-0.070	estimated
γ^j_{pt}	State utility of $s = pt$	-0.117	estimated
γ^j_{mj}	State utility of $s = mj$	0.484	estimated
Men, Married			
γ^j_{ft}	State utility of $s = ft$	0.384	estimated
γ^j_{pt}	State utility of $s = pt$	0.130	estimated
γ^j_{mj}	State utility of $s = mj$	0.480	estimated
Women, Single, No Kids			
γ^j_{ft}	State utility of $s = ft$	0.007	estimated
γ^j_{pt}	State utility of $s = pt$	0.226	estimated
γ^j_{mj}	State utility of $s = mj$	0.857	estimated
Women, Single, Kids			
γ^j_{ft}	State utility of $s = ft$	-0.501	estimated
γ^j_{pt}	State utility of $s = pt$	0.531	estimated
γ^j_{mj}	State utility of $s = mj$	0.896	estimated
Women, Married			
γ^j_{ft}	State utility of $s = ft$	-0.210	estimated
γ^j_{pt}	State utility of $s = pt$	0.984	estimated
γ^j_{mj}	State utility of $s = mj$	1.962	estimated

convexity of search cost is close to two. The positive value for ϵ implies that the state (dis-)utilities are scaled up in higher skill segments.

Table 3 shows the firm and labor market parameters. The within-firm relative productivity of part-time and marginal jobs is estimated to be 1.05 and 0.91 respectively.

The vacancy posting cost function for full- and part-time jobs is not very convex as $\kappa_{2f} = 1.75$, $\kappa_{2p} = 1.53$ and $\kappa_{2m} = 2.09$ are not substantially greater than two.[23]

The top bars in each of the panels of Figure 3 show that the model is able to capture the overall distribution of labor market states and job finding rates.[24] In the estimated model (data), 7.5% (6.4%) of workers are unemployed with 51.4% (51.8%) of them in long-term

[23]For Brazil, Engbom and Moser (2018) estimate a value of 1.45. Shephard (2017) assumes a quadratic vacancy posting cost function in the UK.

[24]See table A.3 for the values underlying Figure 3.

Table 3: Firm, Labor Market and Policy Parameters

Name	Description	Value	Source
Firms			
m	Mass of firms	0.025	estimated
α	Scale of $log(p)$	2.269	estimated
θ	Shape of $log(p)$	0.106	estimated
α_{ft}	Relative productivity ($x = ft$)	1.00	normalized
α_{pt}	Relative productivity ($x = pt$)	1.05	estimated
α_{mj}	Relative productivity ($x = mj$)	0.91	estimated
κ_1^{ft}	Vacancy posting cost (weight), $x = ft$	100.0	estimated
$\kappa_1^{pt}/\kappa_1^{ft}$	Relative vacancy posting cost, $x = pt$	0.850	estimated
$\kappa_1^{mj}/\kappa_1^{ft}$	Relative vacancy posting cost, $x = mj$	0.791	estimated
κ_2^{ft}	Vacancy posting cost (convexity), $x = ft$	1.750	estimated
κ_2^{pt}	Vacancy posting cost (convexity), $x = pt$	1.534	estimated
κ_2^{mj}	Vacancy posting cost (convexity), $x = mj$	2.087	estimated
Labor Market			
ξ	Vacancy-elasticity of matches	0.3	literature
$\bar{r}_F$	Wage rate ($x = ft$)	0.605	estimated
$\bar{r}_{pt}$	Wage rate ($x = pt$)	0.605	estimated
$\bar{r}_x$	Wage rate ($x = mj$)	0.548	estimated
e_{ft}	Hours ($x = ft$)	1.0	normalized
e_{pt}	Hours ($x = pt$)	0.615	SOEP
e_{mj}	Hours ($x = mj$)	0.223	SOEP
$\pi_{su\|e_{ft}}$	Transition from e_{ft} to su	0.010	SIAB
$\pi_{su\|e_{pt}}$	Transition from e_{pt} to su	0.019	SIAB
$\pi_{su\|e_{mj}}$	Transition from e_{mj} to su	0.030	SIAB
$\pi_{lu\|su}$	Transition from su to lu	0.075	estimated
ϕ_{su}	Search efficiency, $s = su$	0.337	estimated
ϕ_{lu}/ϕ_{su}	Relative search efficiency, $s = lu$	0.384	estimated
ϕ_{ft}/ϕ_{su}	Relative search efficiency, $s = e_{ft}$	1.147	estimated
ϕ_{pt}/ϕ_{su}	Relative search efficiency, $s = e_{pt}$	0.911	estimated
ϕ_{mj}/ϕ_{su}	Relative search efficiency, $s = e_{mj}$	0.834	estimated
ψ_{ft}	Godfather shock, $x = ft$	0.017	SIAB
ψ_{pt}	Godfather shock, $x = pt$	0.022	SIAB
ψ_{mj}	Godfather shock, $x = mj$	0.050	SIAB

unemployment. Among the employed workers, 9.0% (9.6%) have a marginal job, 27.4% (24.0%) work part-time and 63.6% (66.3%) have a full-time job. The job finding rate out of short-term unemployment is 28.5% (29.6%) and considerably lower for long-term unemployed workers with 7.0% (6.7%). The difference in job-finding rates reflects the fact that search is estimated to be substantially less efficient in generating matches with firms ($\phi_{lu} < \phi_{su}$). In addition, long-term unemployed workers have lower human capital and thus lower incentives to search for jobs compared to short-term unemployed workers.

The estimated model is also able to capture most of the heterogeneity across sociodemographic groups. Compared to men, a much larger share of women and in particular single women with kids and married women work in part-time or marginal jobs. While the model can replicate the observed heterogeneity in employment levels, the unemployment rate of single men and especially single women with kids and married women is less than perfectly matched.

Figure 4 and table A.6 show the distribution wages over selected wage bins. The overall fit (panel A) is remarkably good given the limited flexibility imposed by the parametric skill and productivity distributions and the fact that there are no skill-dependent parameters.[25] Only 2.4% (1.8% in the data) of all jobs pay a wage below 6.5 Euro, 8.5% (9.8%) of wages are above 6.5 Euro but below 8.5 Euro, 22.1% (18.8%) of wages are between 8.5 and 12.5 Euro, 33.6% (34.6%) are between 12.5 and 20 Euro and 33.4% (35.0%) of wages exceed 20 Euro. The model is also able to capture gender-specific heterogeneity as a larger share of women find themselves in the lower wage bins. Similar to the data, 14.1% (16.5%) of women are affected by the initial minimum wage, only 7.8% (6.7%) of men earn less than 8.5 Euro per hour. However, the right tail of the wage distribution of men is too short while that of women is too long relative to the data.

The differences in the job-type-specific wage distribution (panels B to D) are also replicated by the model. Full-time jobs pay substantially higher wages than part-time jobs which in turn pay higher wages than marginal jobs. Hence, minimum wages will cut deeper into the wage distribution of part-time and marginal jobs compared to full-time jobs. In particular, the initial minimum wage affects 45.8% (53.9%) of marginal jobs, 10.8% (12.1%) of part-time jobs but only 5.8% (5.5%) of full-time jobs. The most important difference between model and data is that the distribution of wages for marginal jobs is too dispersed. There are too many jobs paying a wage below 6.5 Euro or above 12.5 Euro and too few jobs in the range between 6.5 and 12.5 Euro. In addition, too few full-time jobs pay wages between 8.5 and 12.5 Euro. This will affect how the distribution of job types is affected by the minimum wage. Figure 5 shows the share of full-time, part-time, marginal jobs and men in each of these wage bins. Marginal jobs are over-represented in the lowest wage bin. In addition, part-time jobs are over-represented in the wage bins around the initial minimum wage of 8.5 Euro as there are not enough full-time jobs in this range. While these differences between model and data need to be kept in mind, the model delivers a good fit to the joint distribution of wages and job types which is a complicated object.

The distribution of worker and firm fixed effects for full-time jobs is shown in Figure 6. Figure 7 shows the distribution of full-time firm fixed effects among part-time and marginal jobs. In particular, panels C and D show the employment weighted variation in firm productivity among part-time and marginal jobs which the model is able to match quite closely. Panels E shows the percent difference between the q^{th} quantile of the firm productivity distribution weighted by part-time employment and the corresponding quantile of the firm productivity distribution weighted by full-time employment. Both in the data and the model, firm productivity is just slightly lower among full-time workers (about

[25] Engbom and Moser (2018) estimate a set of labor market parameters for each skill segment.

Table 4: Model Fit – Clustered AKM Wage Decomposition

	Total	Workers	Firms	Sorting	
	$var(\ln w)$	$var(\ln h)$	$var(\ln p)$	$2cov(\ln h, \ln p)$	$corr(\ln h, \ln p)$
Value					
Data	0.219	0.119	0.028	0.072	0.624
Model	0.213	0.175	0.016	0.022	0.215
Share					
Data	–	54.4 %	12.8 %	32.9 %	–
Model	–	82.2 %	7.3 %	10.5 %	–

Notes: This table shows the variance decomposition of log wages into a worker component, a firm component and their covariance. The worker and firm-class fixed effects are estimated using the years 2011 to 2014 and 25 firm classes. The column "Total" refers to the total *explained* variance, i.e. the total variance minus the residual variance. In the data, the residual variance accounts for only 3% of the total variation. In the model, there is no distinction between explained and total variance. Data: SIAB, own calculations.

5%). Using marginal workers as weights instead of full-time workers, the firm productivity distribution shifts downward by around 20% in the data but by significantly more in the model (panel F). Hence, marginal workers in the model work at firms that pay too low full-time wages compared to the data.[26] Table 4 shows the variance decomposition of full-time wages. Worker heterogeneity contributes 83.9% (54.4%), firm heterogeneity 6.1% (12.8%) and sorting 10.0% (32.9%) to the overall variation in full-time wages. The correlation between worker and firm fixed effects is 0.221 in the model and 0.624 in the data.[27] The fact that the model cannot capture this large degree of positive sorting implies that I underestimate the overall variation in full-time wages by about four log points. The fact that the model cannot capture this large degree of positive sorting implies that I underestimate the overall variation in full-time wages by about four log points.

[26] See Table A.7 and Table A.8 for details.

[27] The correlation of 0.624 in the data is rather high. Using the same methodology, Bonhomme et al. (2019) find a correlation of 0.5 for Sweden. In order to match the observed correlation of worker and firm fixed effects, one may extend the model to make the probability of job destruction dependent on worker skill and firm productivity (true in the data).

Figure 3: Model Fit – Employment Moments

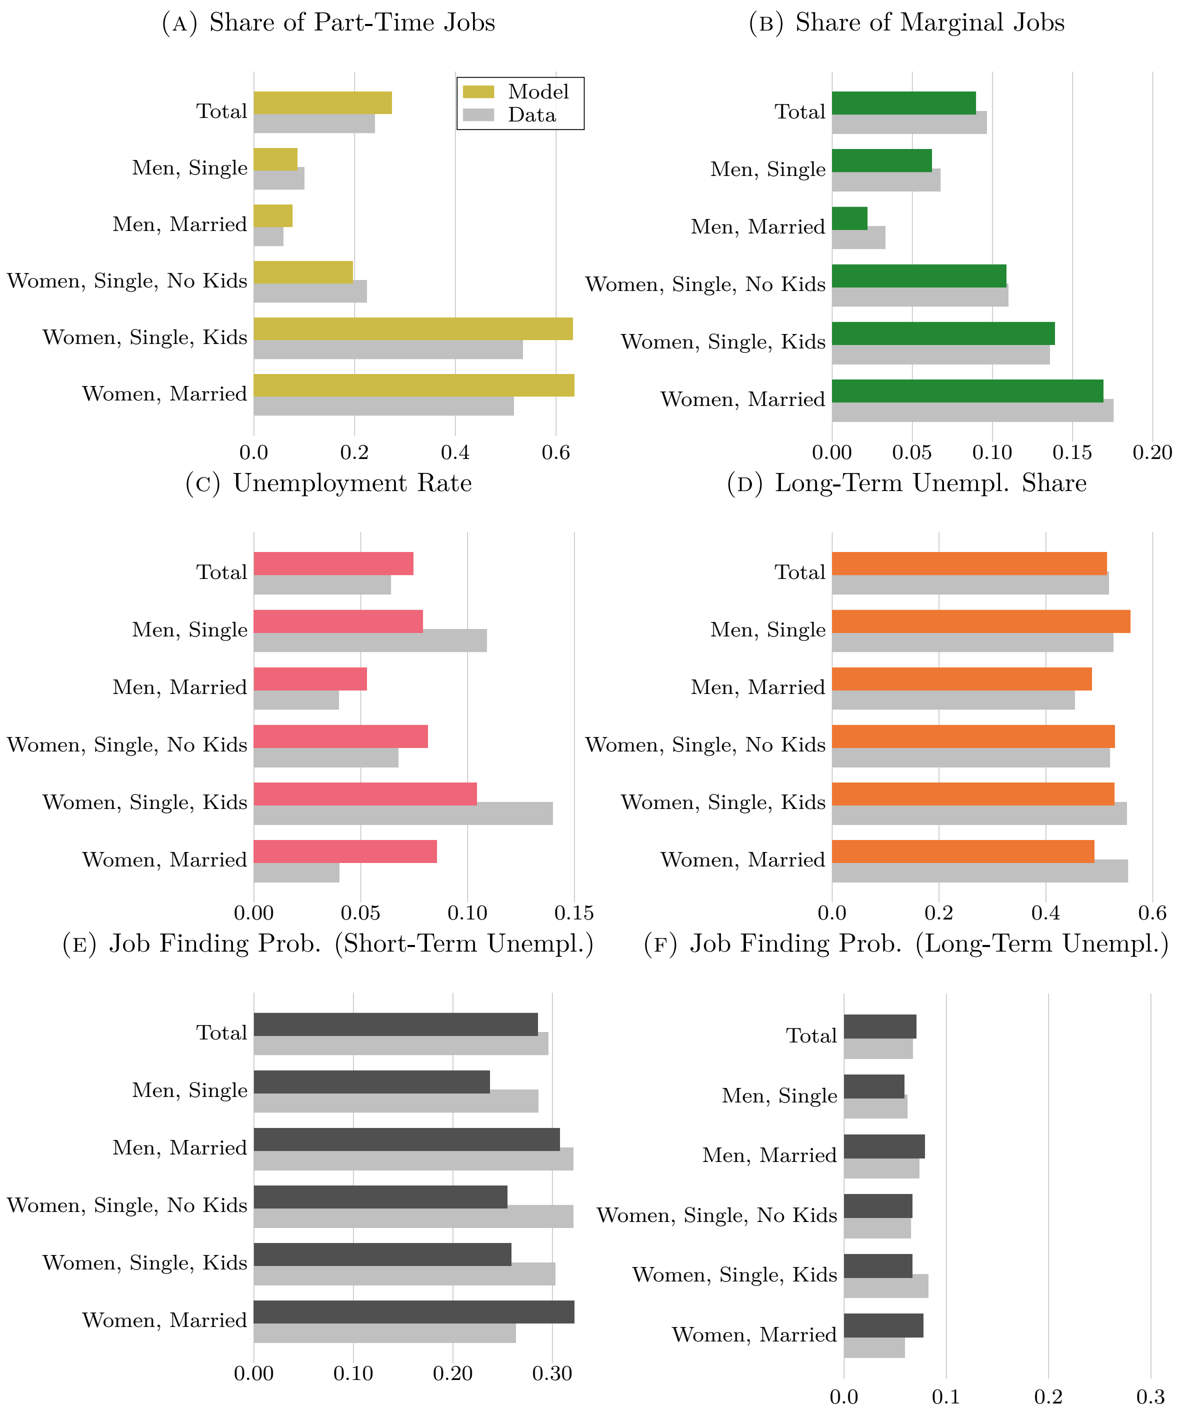

Notes: This figure shows labor market moments targeted in the estimation for the full population (Total) and within the sociodemographic groups. Subfigures 1 and 2 show the probability of working a part-time and marginal job conditional on being employed. Subfigure 3 shows the unemployment rate and subfigure 4 the share of long-term unemployed workers conditional on being unemployed. Figures 5 and 6 show the job finding probabilities for short- and long-term unemployed workers. Data: SIAB, SOEP.

FIGURE 4: Model Fit – Wage Groups by Job Types and Gender

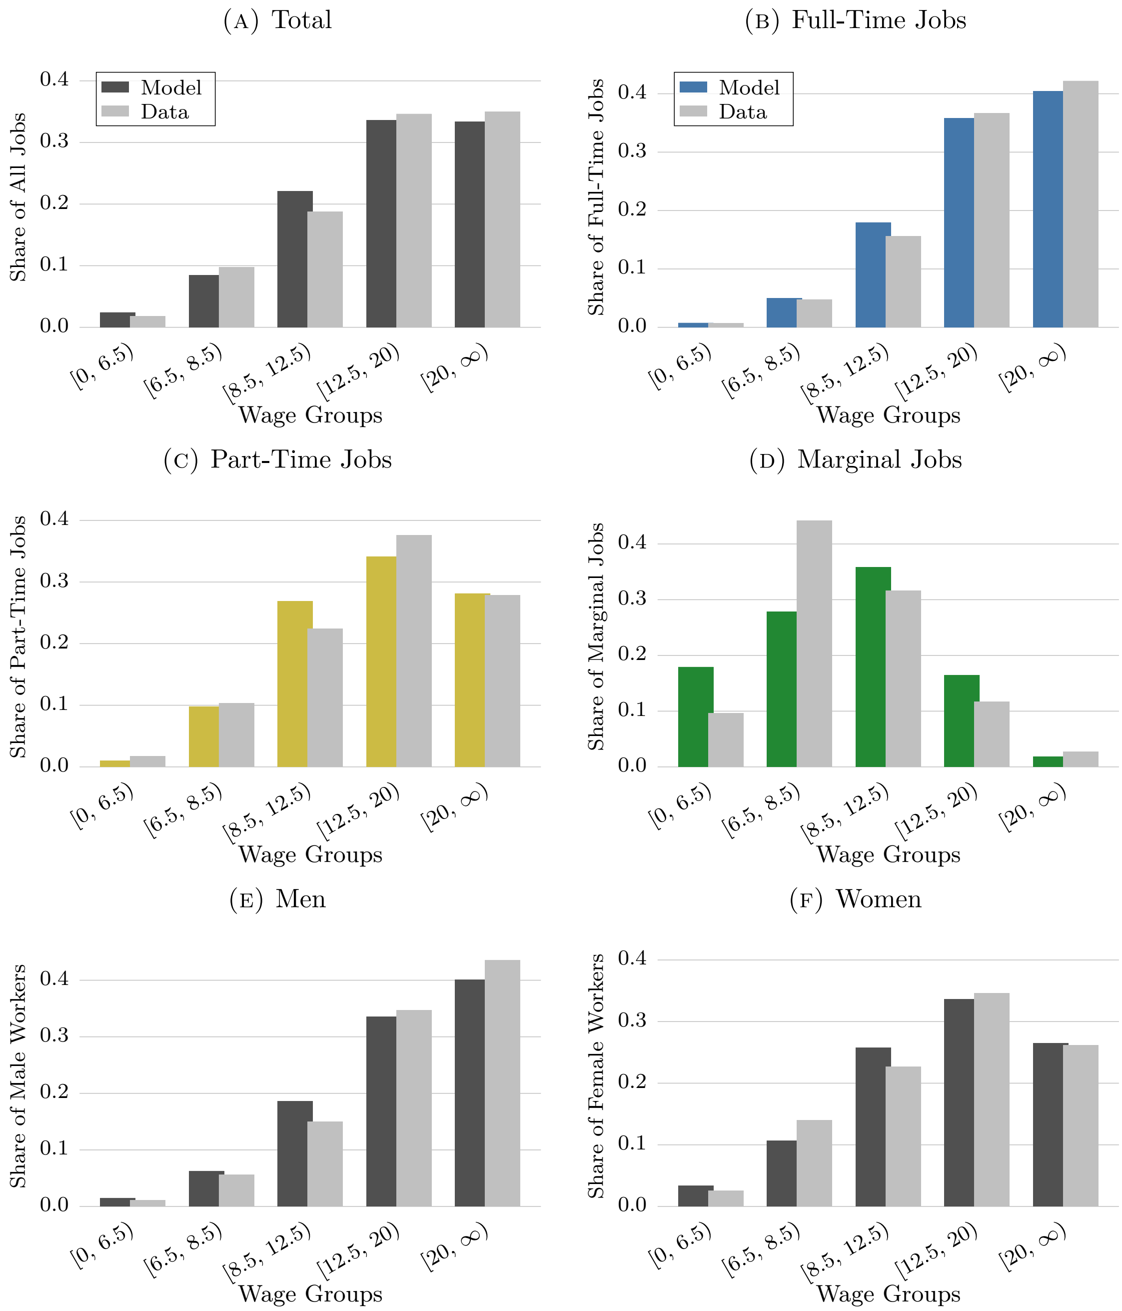

Notes: This figure shows the distribution of jobs over four wage groups for all workers and separately for full-time, part-time, marginal job, male and female workers in the model and data. Data: SIAB, SOEP.

FIGURE 5: Model Fit – Job Types and Gender By Wage Groups

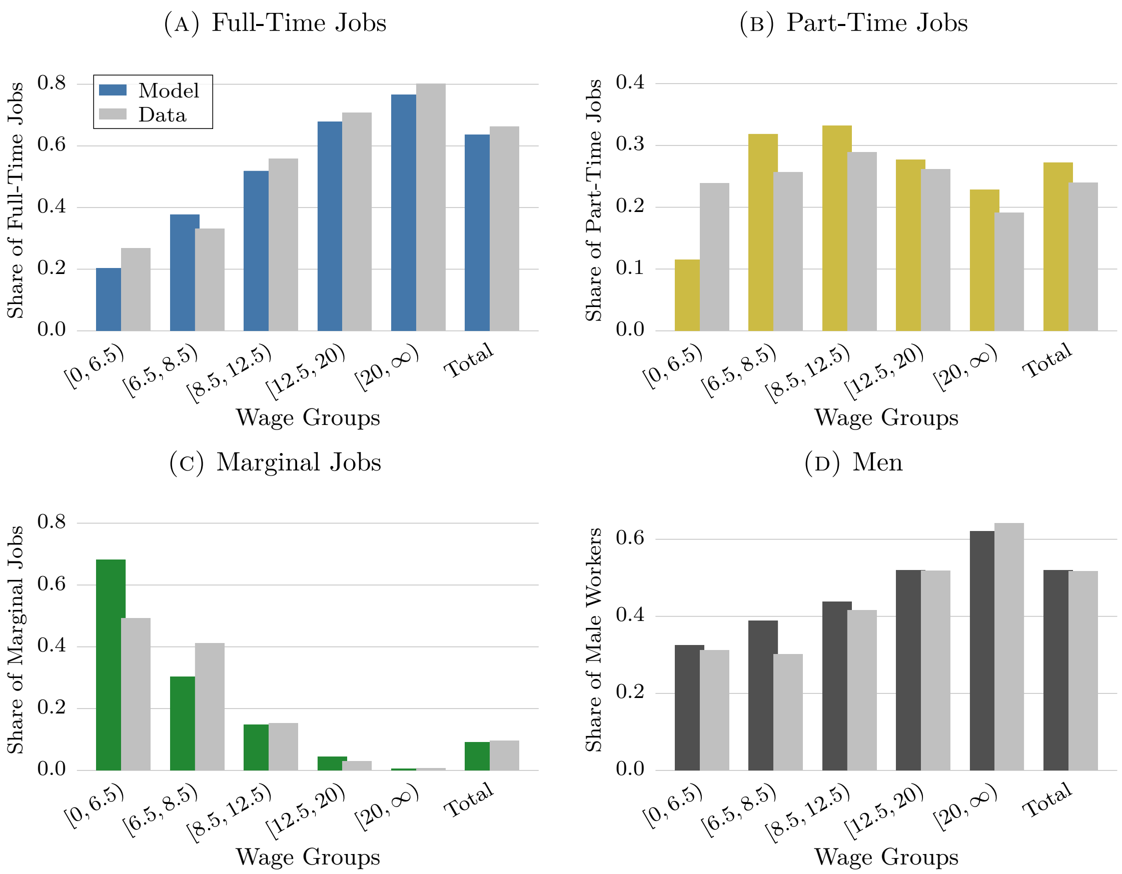

Notes: This figure shows the share of full-time, part-time and marginal jobs as well as the share of men within various bins of the wage distribution in the model and data. Data: SIAB, SOEP.

FIGURE 6: Model Fit – Clusterd AKM Fixed Effects

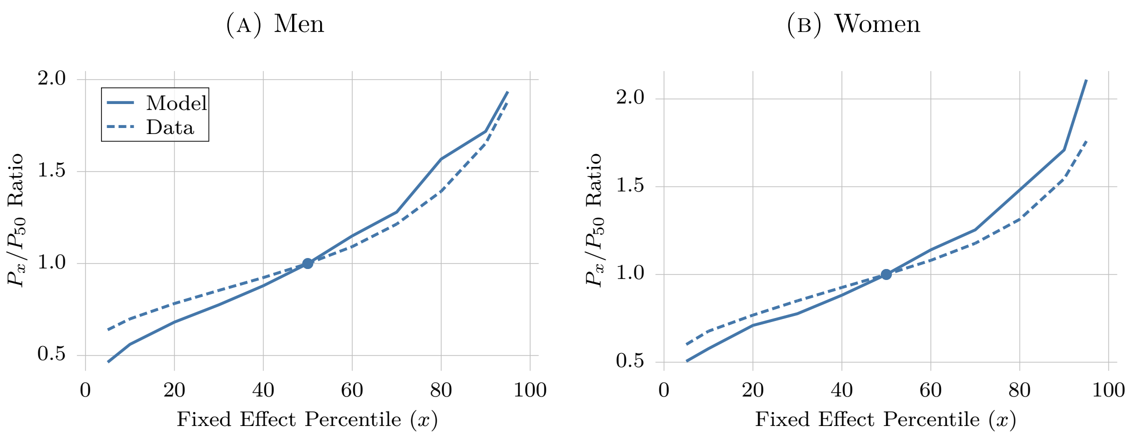

Notes: This figure shows the ratios of selected percentiles to the median of the distributions clustered AKM worker fixed effects for men and women. See appendix B for details. Data: SIAB.

FIGURE 7: Model Fit – Firm Fixed Effect

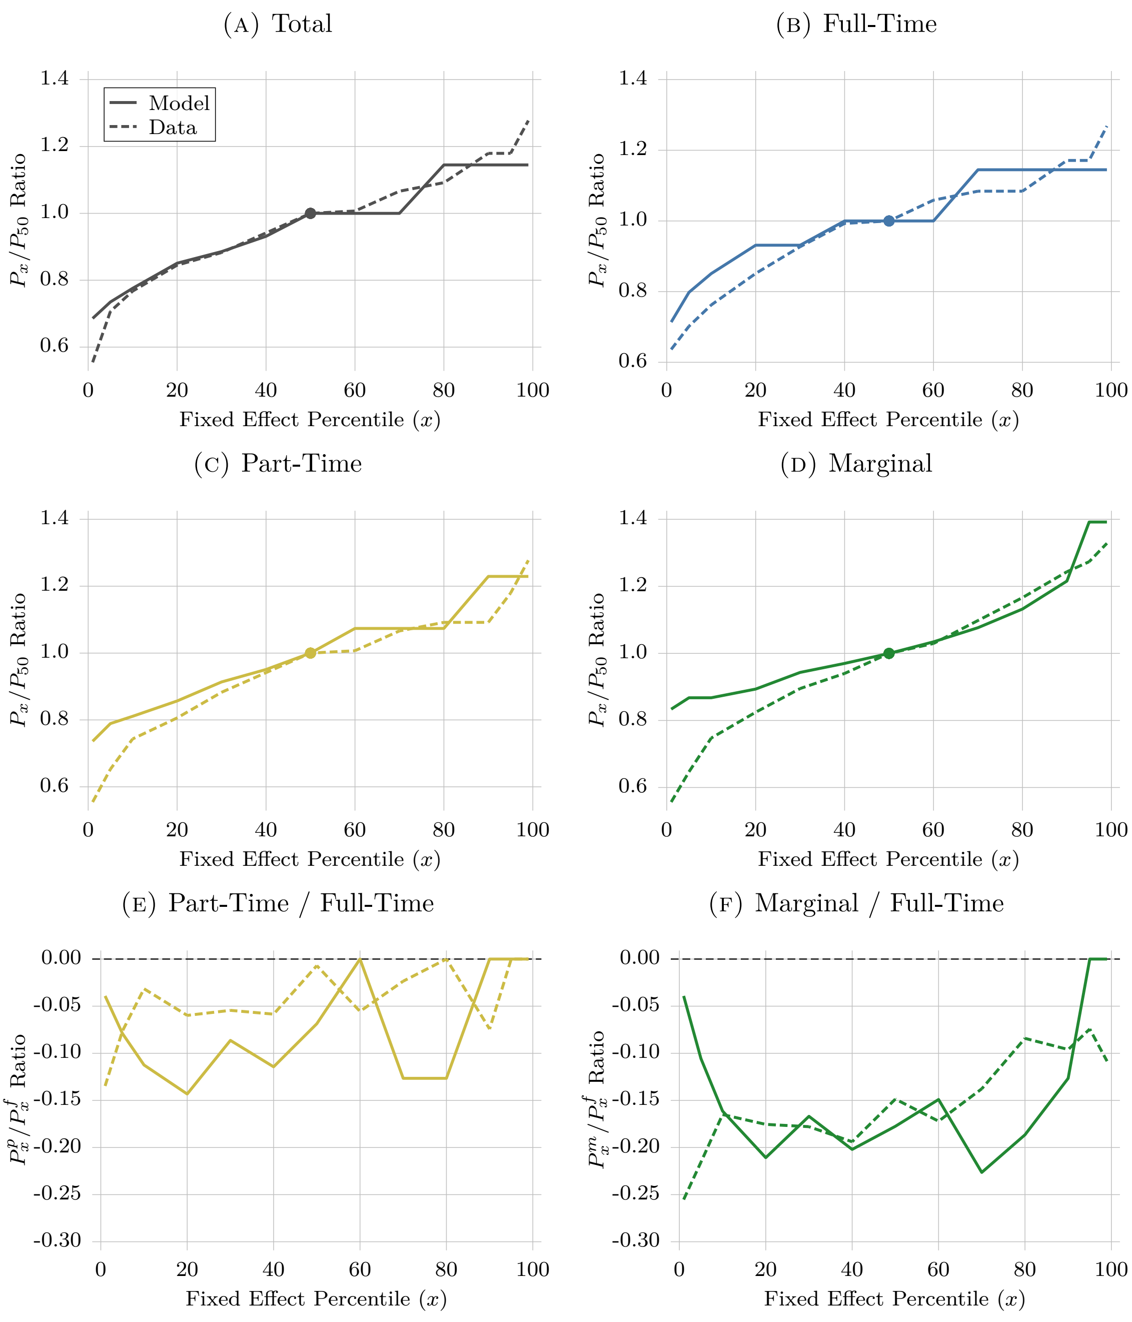

Notes: This figure shows the distribution of (clustered) firm fixed effects estimated using clustered AKM on full-time jobs. In panels A, B, C and D, all jobs, only full-time, only part-time jobs and only marginal jobs are used as weights respectively. Panels E and F show how the distributions change when weighting by part-time and marginal jobs instead of full-time jobs. Data: SIAB.

4 The German Minimum Wage Reform of 2015

In 2015, the German government introduced a federal minimum wage of 8.5 EUR (Kaitz index of 47%) that cut deep into the wage distribution affecting more than 10% of all jobs. In this section, I use the estimated model to analyze how the initial federal minimum wage affected employment, productivity and output. First, I compare the pre- and post-reform steady states and highlight the mechanisms at play (4.1). Second, I analyze the transitional dynamics (4.4).

4.1 Steady State Comparison

Table 5 compares the steady state without a minimum wage to the steady state with a minimum wage of 8.5 EUR (column 2). The difference between the new and the old stationary equilibrium is shown in column 3.

Relative to the steady state without a minimum wage, the unemployment rate is slightly lower (by 0.035 percentage points) in the steady state with a minimum wage of 8.5 EUR. The small change in the number of jobs masks heterogeneity across employment levels. In particular, while the share of marginal jobs among all jobs drops from 9.14% to 7.94%, the share of part-time and full-time jobs increases by 0.81 and 0.39 percentage points respectively.

The slight decrease in the unemployment rate occurs despite a small drop in the average job finding rate out of unemployment, $\Pr(e|u)$, by 0.124 percentage points (0.07%). the effect of the lower job-finding rate is counterbalanced by a decrease in the average job destruction probability of about 0.02 percentage points (1.2%). Note that this is a direct result of the reallocation away from unstable marginal jobs towards more stable part-time and full-time jobs.[28]

Average wages in the new stationary equilibrium are up by about 2.1%. Part of this increase is driven by reallocation to more productive firms. In other words, workers now work at firms where they would have received 0.5% higher wages in the absence of a minimum wage. While over two thirds of the increase in productivity reflects reallocation to more productive firms (higher p), part of the increase in productivity ($a_x p$) is a direct result of the shift away from relatively unproductive marginal jobs. Note that this is broadly consistent with the evidence reported by Dustmann et al. (2020) who show that about 25% of the wage increase of employed workers can be attributed to the reallocation channel.

Average gross earnings increase by more than wages (+3.5%) reflecting the shift towards jobs with longer hours (+1.4%). Taxes and transfers result in a 2.8% increase in

[28] Recall that the job destruction probability of any given employment relationship is exogenous and thus not affected by the minimum wage. Endogenous job destruction only occurs if a job becomes unprofitable due to a prohibitively high minimum wage. In the new steady states, however, these jobs are not created in the first place. The drop in the average job destruction rate is hence a pure composition effect.

Table 5: Minimum Wage Effects – General Equilibrium

	(1)	(2)	(3)
	Baseline ($\bar{w} = 0$)	New Equilibrium ($\bar{w} = 8.5$)	
	Value	Value	Change
Labor Market States			
Unemployment Rate	7.44%	7.40%	-0.035
Long-Term Share	51.17%	51.34%	0.170
Full-Time Share	63.60%	63.99%	0.394
Part-Time Share	27.26%	28.07%	0.812
Marginal Share	9.14%	7.94%	-1.206
Transition Probabilities			
$\Pr(e\|u)$	17.54%	17.42%	-0.124
$\Pr(su\|e)$	1.41%	1.39%	-0.017
Wages, Earnings & Incomes			
Log Wages	2.776	2.796	0.021
Log Productivity	0.382	0.388	0.005
Log Hours	3.389	3.403	0.014
Log Earnings	7.631	7.665	0.035
Log Net Earnings	7.279	7.308	0.028
Log Income	7.583	7.590	0.008
Macro Aggregates			
Log Output	8.305	8.310	0.005
Log Transfers	4.554	4.493	-0.060
Log Labor Taxes	6.719	6.728	0.009

Notes: This table shows the long-run general equilibrium effects of the introduction of a federal minimum wage of 8.5 EUR relative to the baseline equilibrium without a minimum wage (first column). Changes refer to the absolute difference to the baseline outcome (e.g. percentage points or log points).

average earnings and a 0.8% increase in incomes. The relatively weak increase in incomes follows from the fact that many low-skill workers top up their earnings with unemployment benefits. Reallocation to better firms and longer hours leads total output to grow by 0.5%. While the tax-and-transfer scheme mutes the increase in incomes, total transfer payments decrease by 6.0%. In addition, the government's revenues from labor taxation increase by 0.9% as average earnings grow and the unemployment rate falls slightly.

In sum, the introduction of the minimum wage moved the economy into an equilibrium with higher productivity, output and employment. While the unemployment rate decreases only slightly, employment weighted by hours worked increases markedly as the share of part-time and full-time jobs rises. While the tax- and especially the transfer-system prevents incomes from growing more strongly, workers are less reliant on government transfers to top up their earnings. Combined with the fact that higher average earnings raise tax revenues, the reform improved the government's budget position.

Figure 8: Mechanisms - Partial vs. General Equilibrium

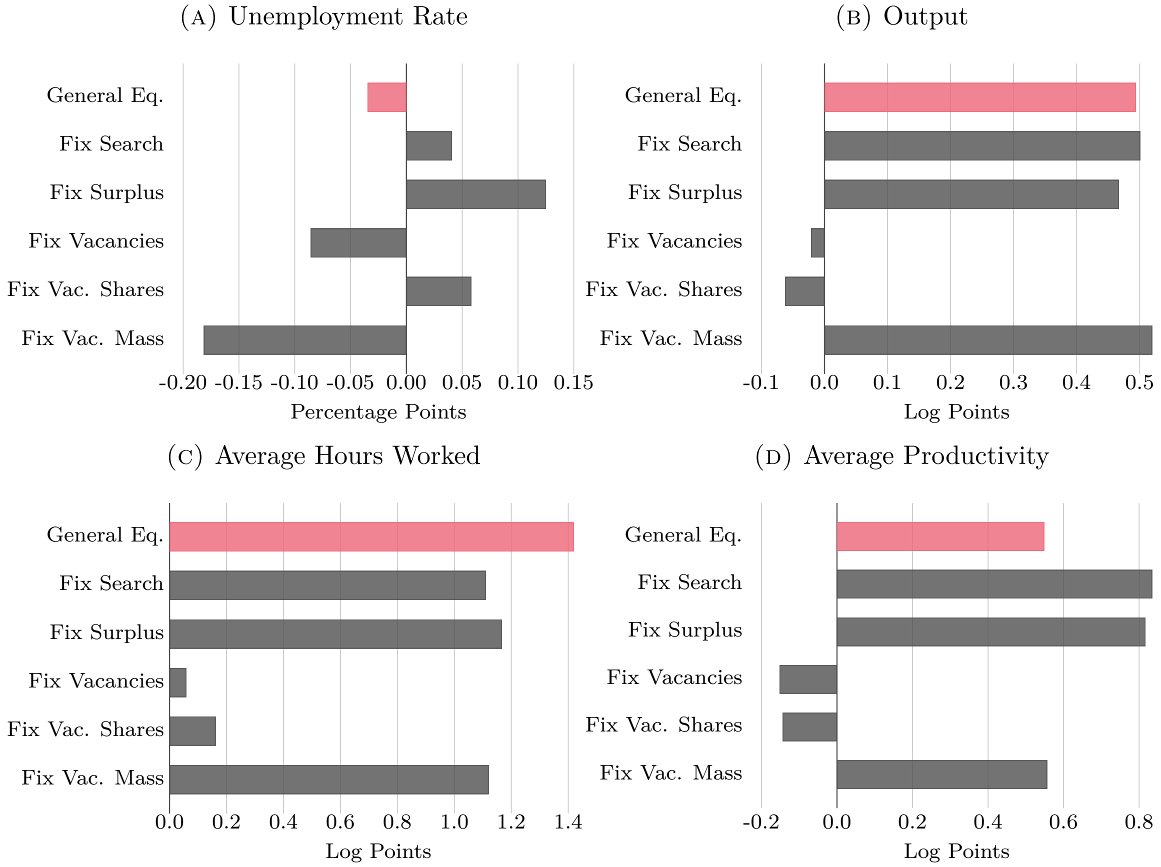

Notes: This figure shows the effects of the minimum wage of 8.5 EUR for different partial equilibrium scenarios that result from shutting down different margins of adjustment. If a partial equilibrium differs from the general equilibrium scenario, the respective channel is important for generating the general equilibrium effect.

4.2 Mechanisms

I now study the importance of the different mechanisms that feed into the general equilibrium effects. To that end, I shut down different margins of adjustment one at a time. Table 6 shows the partial equilibrium effects of fixing workers' search effort, workers' surplus of successful search, firms' vacancy posting, the vacancy shares, and the mass of vacancies to the baseline levels in columns 3 through 7. Columns 1 and 2 report the baseline levels and general equilibrium effects from Table 5. Figure 8 visualizes the general and partial equilibrium effects for unemployment, output, hours worked and firm productivity.

Panel A of Figure 8 shows that eliminating workers' job-finding surplus or search effort pushes the unemployment rate up while shutting down firms' vacancy posting pushes it down. When workers' surplus of successful search is held fixed, the unemployment rate increases by about 0.13 percentage points instead of the decrease by 0.035 percentage points in general equilibrium. The effect of fixing search effort is smaller than that of fixing the surplus as search effort is negatively affected by the drop in labor market tightness.

TABLE 6: Minimum Wage Effects – Mechanisms

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Baseline	GE	Fix Search	Fix Vacancies	Fix Surplus	Fix Vac. Shares	Fix Vac. Mass
	Value	Change	Change	Change	Change	Change	Change
Unemployment							
Unemployment Rate	7.44%	-0.035	0.041	-0.086	0.125	0.058	-0.181
$\Pr(e\|u)$	17.54%	-0.124	-0.263	0.212	-0.481	-0.174	0.299
$\Pr(su\|e)$	1.41%	-0.017	-0.013	-0.001	-0.014	-0.002	-0.014
Employment Level							
Log Hours	3.389	0.014	0.011	0.001	0.012	0.002	0.011
Full-Time Share	63.60%	0.394	0.191	0.023	0.239	0.101	0.370
Part-Time Share	27.26%	0.812	0.808	0.023	0.792	0.009	0.553
Marginal Share	9.14%	-1.206	-0.999	-0.047	-1.032	-0.111	-0.923
Wages, Earnings & Incomes							
Log Wages	2.776	0.021	0.022	0.018	0.022	0.019	0.020
Log Productivity	0.382	0.005	0.008	-0.002	0.008	-0.001	0.006
Log Earnings	7.631	0.035	0.033	0.018	0.034	0.020	0.031
Log Net Earnings	7.279	0.028	0.027	0.015	0.028	0.017	0.025
Log Income	7.583	0.008	0.008	0.004	0.008	0.005	0.007
Macro Aggregates							
Log Output	8.305	0.005	0.005	-0.000	0.005	-0.001	0.005
Log Transfers	4.554	-0.060	-0.048	-0.034	-0.042	-0.026	-0.066
Log Labor Taxes	6.719	0.009	0.008	0.005	0.008	0.005	0.009

Notes: This table shows the long-run effect of the introduction of a federal minimum wage of 8.5 EUR relative to the baseline equilibrium without a minimum wage (first column). The second columns shows the new general equilibrium. In the remaining columns, different margins of adjustment are switched off one at a time. In column 3, workers' search effort is held fixed. In column 4, firms' vacancy posting is held fixed. In column 5, workers can adjust their search effort, but the expected surplus of meeting a firm is held constant (search effort will still react to changes in labor market tightness). In column 6, the hours-and productivity shares of posted vacancies are held fixed, but the total mass of vacancies is allowed to adjust. In column 7, the total mass of vacancies is held fixed, but the hours- and productivity distribution of vacancies is allowed to adjust. In all scenarios, unprofitable vacancies may not be posted. Changes refer to the absolute difference to the baseline outcome (e.g. percentage points or log points).

When firms' vacancy posting policies are held fix, the unemployment rate decreases by 0.9 percentage points. Taking a closer look at the role of vacancy posting, we see that there are two opposing effects. On the one hand, the total mass of vacancies is reduced which drives up unemployment (via lower job finding rates). On the other hand, the change in the composition of posted vacancies away from unstable low-hours jobs lowers unemployment (by reducing the average job destruction rate). Besides this effect on average job destruction rates, the change in the hours-distribution of vacancies raises searchers' expected disutility from longer working hours and thus dampens the increase in the surplus of successful search and hence search effort and job finding rates. The reduction in overall vacancy posting, however, dominates such that the net effect of endogenous vacancy posting drives up the unemployment rate.

The increase in average hours worked (Panel C) and firm productivity (Panel D) is driven mainly by firms' vacancy posting and in particular by the change in the composition of vacancies. In general, firms create fewer vacancies for jobs that are (strongly) affected by the minimum wage. As the minimum wage affects low-hours and low-productivity jobs relatively often, the reduction in vacancies is not symmetric across employment levels. Conditional on meeting a firm, the probability of being offered a low-hours or low-productivity job declines. Figure 9 shows how the productivity distribution of vacancies changes.

For total output (Panel B), I find that the reallocation effect is much more important than changes in the number of jobs. Only fixing the mass of vacancies – which has a relatively big effect on the unemployment rate – does not lead to a lower output effect. The distribution of vacancies across employment levels and firm productivity drives the positive output effect in general equilibrium.

4.3 Heterogeneity Across Sociodemographics

The different sociodemographic groups in the model and the data are differently affected by the minimum wage. Figure 10 shows that women are significantly more likely to earn less than 8.5 EUR per hour. I now analyze how the effects of the reform vary across sociodemographic characteristics in the new stationary equilibrium.

Panel B of Figure 10 displays the percentage point changes in the distribution of labor market states (full-time, part-time, marginal and total employment). All bars sum to zero. While the reallocation pattern away from marginals towards part-time and full-time jobs is the same qualitatively, there is substantial variation in magnitude. Men and single women without kids move to both part-time and full-time jobs. In contrast, the share of married women and single women with kids working full-time jobs hardly increases because of the high disutility of working full-time for this group. As a result, their unemployment rate increases slightly while total unemployment drops.

Figure 9: Productivity Distribution of Vacancies

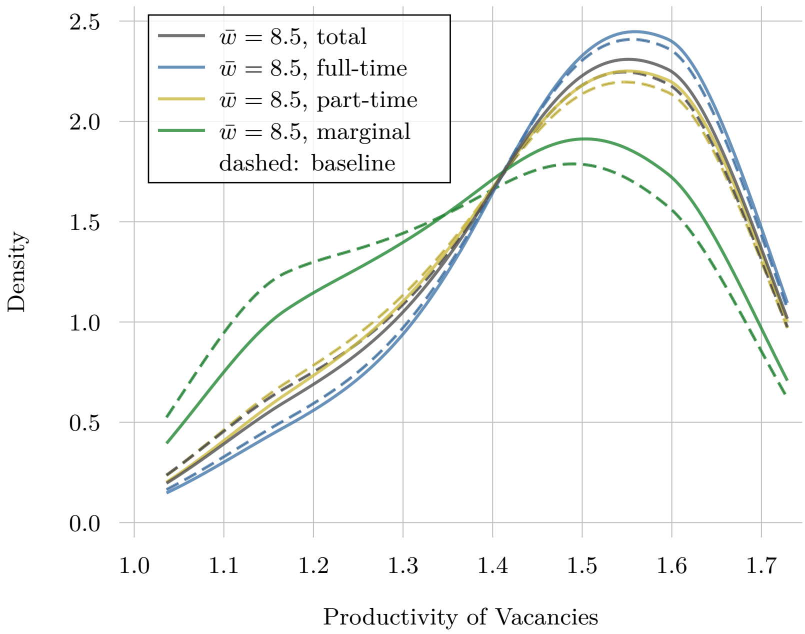

Notes: This figure shows how the productivity distribution of vacancies offered by firms changes in response to the introduction of the minimum wage of 8.5 EUR. I exclude skill segments in which none of the minimum wages considered is binding for any job, i.e. where all wages in the baseline equilibrium are above 14.5 EUR.

Panel C shows how lifetime utility, income and earnings change relative to the baseline equilibrium. Although earnings increase substantially, income growth is much weaker due to the fact that many low-wage workers top up their earnings with government transfers and thus loose the majority of the earnings increase. Perhaps surprisingly, lifetime utility remains almost unchanged and is slightly negative for women. This is because the small increase in income (consumption) is counteracted by lower state utility as workers now work longer hours. Especially those workers who have a strong preference for or rely on marginal jobs with low working hours experience utility losses from the reallocation towards part- and full-time jobs.

To see this more clearly, panel D decomposes the average change in flow utility (closely correlated with lifetime utility) into the components of the utility function. While utility from consumption, $u(c)$, increases, hours-related utility, $\nu(s)$, decreases.[29]

4.4 Transitional Dynamics

In the presence of search frictions, the process of worker reallocation takes time. Workers whose jobs survive the introduction of the minimum wage will gradually transition to more productive firms or jobs with longer hours. More importantly, the minimum wage will

[29] Disutility from search plays almost no role.

make some jobs unprofitable. These workers become unemployed and finding a (better) job takes time. While worker reallocation pushes up output in the long-run, the short-run effects of introducing or raising the minimum wage may be significantly less desirable. It is thus paramount to study the transitional dynamics triggered by the minimum wage reform.[30]

Figure 11 shows how the economy reacts to the reform. Panel A shows that there is indeed a drop in total employment as some jobs become unprofitable. It takes about five years until the employment response turns positive. The matgnitude of the initial layoff shock, however, is very small (0.052 percentage points). It takes roughly ten years for the shift towards part-time and full-time jobs to unfold. This shift directly maps into a reduction of the average job destruction rate (panel B). The job finding rate out of unemployment exhibits rather weak transitional dynamics.

Panel C shows the evolution of average wages. Wages jump up immediately and increase only slightly over the following years. However, the decomposition of wage growth changes over time. Initially wage growth is almost entirely driven by lower profit margins and thus a higher average labor share. As time progresses, workers reallocate to more productive firms. Hence the profit margin recovers and workers' higher wages are increasingly the result of working for more productive firms. We also see that wage growth is not driven by selection of relatively high-skill workers into employment.

Panel D shows that while both average wages and earnings increase immediately after the reform, earnings continue to grow substantially over the following ten to fifteen years. This is driven by the increase in average hours worked. As employment is essentially constant, the increase in average earnings translates into an increase in the government's revenue from labor income taxation. Similarly, low-skill workers receive lower transfers payments to top up their growing earnings (panel F).

As the initial dip in employment is negligible, total output increases monotonically following the reform (panel D). The minimum wage is low enough to prevent an initial dip due to the destruction of unprofitable matches. After five years, total output is already 0.39 log points above the pre-reform level. Taking into account that it takes a substantial amount of time until the new stationary equilibrium is reached, it is more informative to compare the net present value of output with and without the minimum wage. The minimum wage raises the net present value of output by 0.38 log points – about 77% of the log point difference in output between the two steady states.[31]

4.5 Comparison with Reduced-Form Evidence

I now briefly discuss how the model predictions line up with the available reduced-form evidence on the initial introduction of the minimum wage which can be seen as an inde-

[30] Appendix C documents how the transition path is computed.

[31] For this computation, I use the model discount factor to compute the net present values.

pendent test of the model. There are several studies documenting the short-run effects of the 2015 minimum wage reform using individual or regional variation in the bite of the minimum wage (e.g. Garloff, 2016; Caliendo et al., 2017; Holtemöller and Pohle, 2017; Burauel et al., 2020; Dustmann et al., 2020). The results of these studies boil down to the following points.

First, both hourly wages increased significantly and consistent with near full compliance from 2014 to 2016. Earnings grew by more than wages suggesting that work hours increased (Dustmann et al., 2020).[32]

Second, none of the afore-mentioned studies find significant adverse effect on overall employment. However, the minimum wage induced a shift from marginal jobs towards part-time and full-time jobs (Garloff, 2016; Holtemöller and Pohle, 2017). Caliendo et al. (2017) estimate that approximately 2.4% of marginal jobs were lost due to the minimum wage in the first year of the reform. vom Berge et al. (2016) document that the number of marginal workers dropped by about 2% and 4% from December 2014 to January and September 2015 respectively.[33] The model predicts that about 2.7% of marginal jobs were lost on impact and about 4.4% in the fall of 2015. In addition and consistent with the model, turnover rates decreased as both job finding and separation rates were reduced (Bossler and Gerner, 2016).

Third, there is also evidence that the minimum wage reallocated workers to larger, more productive firms (Dustmann et al., 2020). Quantitatively, reallocation to more productive firms seems to have happened slightly quicker in the data than in the model. (Dustmann et al., 2020) attribute about one quarter of the increase in wages from 2014 to 2016 to reallocation to better firms. In the model, about 15% of the wage gain comes from reallocation. This is likely due to the fact that job destruction in the model does not depend on human capital or firm productivity. In the data, the probability of job destruction decreases in both. The broad pattern, however, is consistent with this empirical finding. After five years and in the new steady state, 25% of the wage gain is driven by productivity gains.

In sum, the estimated model captures all these effects qualitatively and does a good job of replicating them quantitatively. The fact that the model not only matches well the labor market moments in the pre-reform period, but is also broadly consistent with the rich reduced-form evidence on the minimum wage reform lends credibility to the optimal policy analysis in section 5.

[32] Recall that I assume full compliance throughout the paper. There is some evidence of non-compliance in the first year after the introduction of the minimum wage (Burauel et al., 2020). However, the issue of non-compliance seems to have been rather transitory. In addition, earnings increased more strongly than wages over the first two years suggesting that, if anything, hours increased.

[33] Caliendo et al. (2017) include all marginal workers while vom Berge et al. (2016) only include those workers where the marginal job is the main job. Note that I use the same definition as vom Berge et al. (2016).

Figure 10: Heterogeneous Effects by Sociodemographics

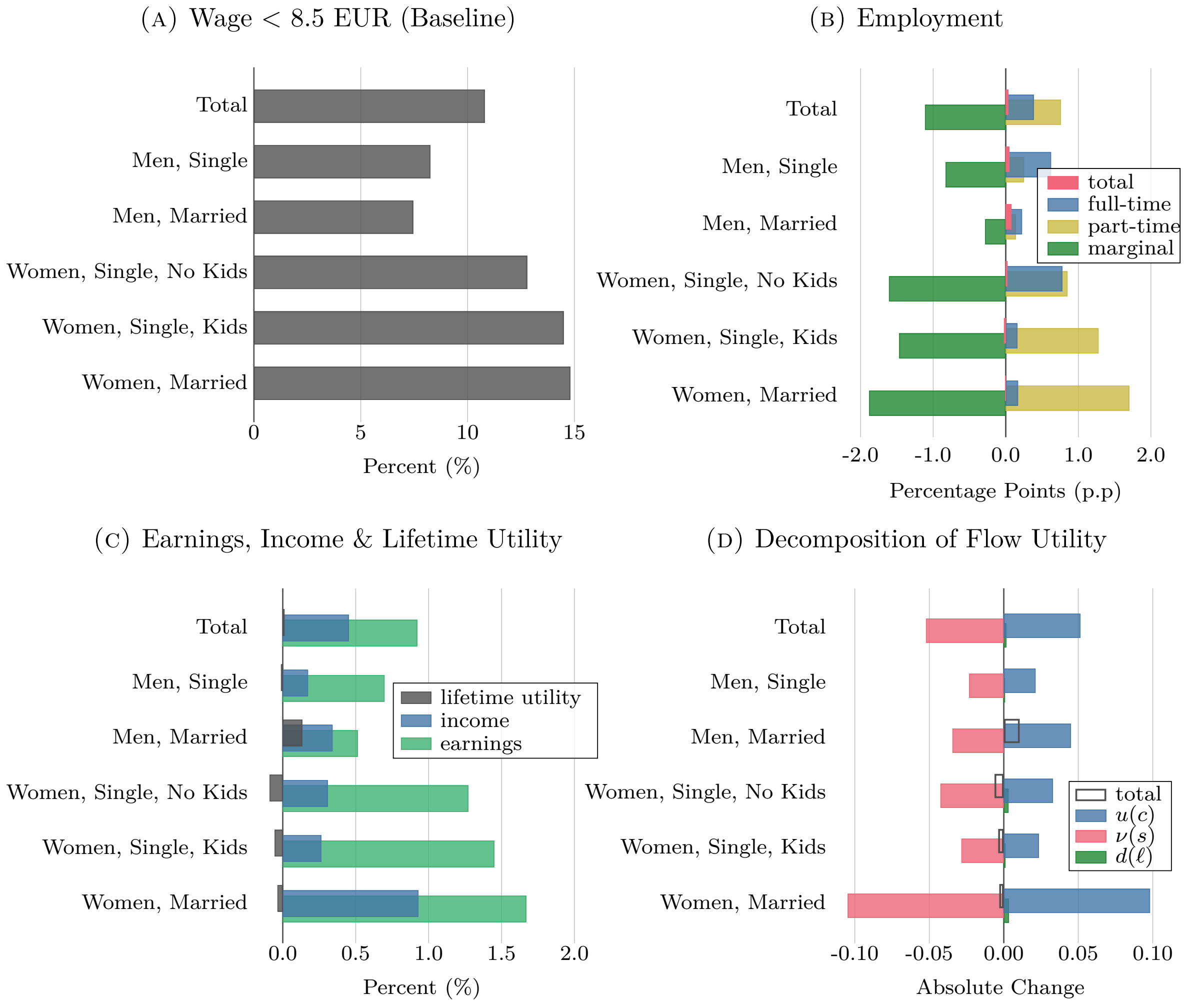

Notes: This figure shows how the effects of the minimum wage of 8.5 EUR vary across sociodemographic groups. Panel A shows how many employed workers are affected by the minimum wage, panel B shows how the distribution of labor market states changes (the bars sum to zero). Panel C shows the relative change in average earnings, income/consumption and lifetime utility. Panel D decomposes the average change in flow utility into its components (see equation 22).

FIGURE 11: Dynamic Effects of the Initial Minimum Wage

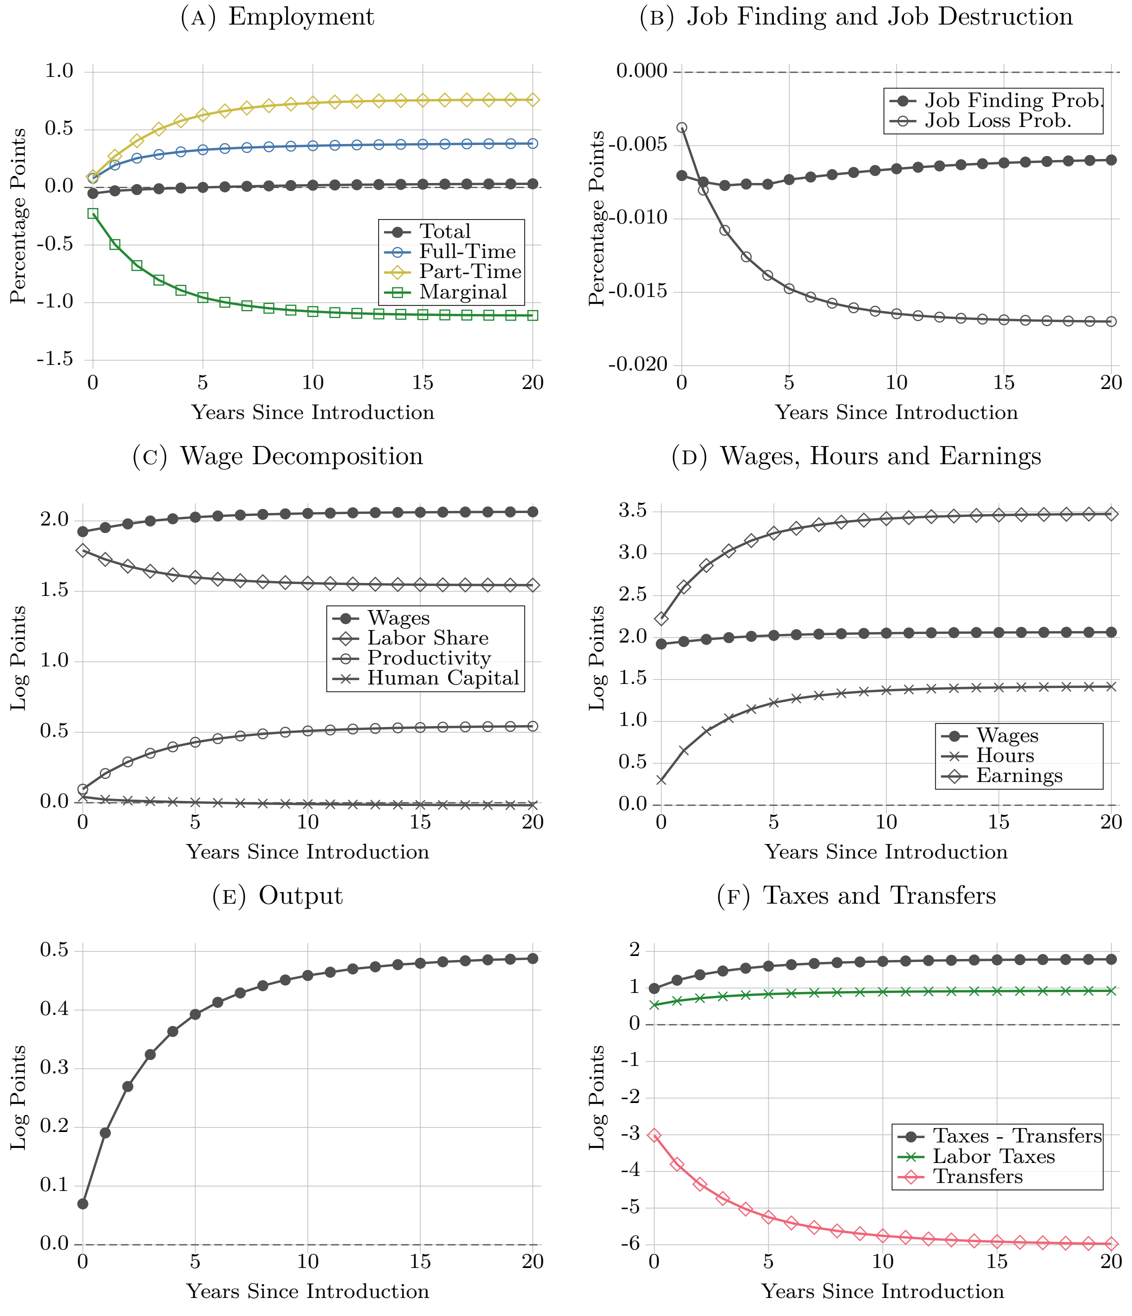

Notes: This figure shows the predicted changes in employment by job type (panel A), average job finding and job destruction probabilities (panel B), the components of wage growth (panel C), averages wages, hours and earnings (panel D), total output (panel E) and total taxes and transfers (panel F) following the introduction of a minimum wage of 8.5 EUR.

5 Counterfactuals: Increasing the Minimum Wage

In this section, I use the structural model to analyze how increasing the minimum will impact employment and output and welfare. First, I will analyze the long-run employment and output effects (section 5.1). Second, I will analyze the entire transition path and compare short- and long-run effects (section 5.2). Third, I will discuss heterogeneity in welfare effects of minimum wages (section 5.3).

5.1 Long-Run Effects

I first take a long-run perspective by comparing the stationary equilibrium that emerges for different minimum wages to the baseline equilibrium. Figure 12 shows steady-state employment, hours worked, output and average productivity as a function of the minimum wage. Panel A shows that total employment, i.e. the share of employed workers, is a non-monotonic function of the minimum wage. Employment is maximized at a minimum wage of 11.0 EUR (Kaitz = 60%). Employment does not drop below the baseline level for minimum wages below 13.0 EUR (Kaitz = 70%). Quantitatively, the positive employment effect of moderate minimum wages is very small ($\leq$ 0.2 p.p.) while the decline in employment for high minimum wages is quite steep.

Panel B shows that the minimum wage not only affects the extensive margin of employment but also the average employment level. While the number of employed workers increases only slightly for moderate minimum wage hikes, total hours worked do increase significantly. Importantly, the hours-maximizing minimum wage of 13.5 EUR (Kaitz = 73%) is considerably higher than the employment-maximizing minimum wage. Hence, even though employment starts to decline, total hours continue to increase because the average employment level increases. At the hours-optimum, total hours worked are 4.3% above the baseline level.

This increase in total hours worked implies that the output-maximizing wage is significantly higher than the employment-maximizing one and that the minimum wage can lead to long-run output growth. Indeed, Panel C shows that output increases considerably in the minimum wage up until 14.4 EUR (Kaitz = 78%). At the optimum, total output is 3.7% higher compared to the baseline without a minimum wage. Beyond that point, output starts to decline as total employment drops sharply. Panel D shows that *average* firm productivity increases monotonically in the minimum wage. This second margin of reallocation explains why the output-optimum is above the hours-optimum. At the output-maximizing minimum wage, jobs are on average 4.1% more productive.

In order to understand what drives these minimum wage effects in Figure 12, I fix (a) workers' search policies, (b) the expected surplus of meeting a firm, (c) the productivity-hours distribution of vacancies, and (d) the mass of vacancies at the corresponding baseline levels. Figure 13 illustrates the results in each of these partial equilibrium scenarios.

Figure 12: Long-Run Minimum Wage Effects

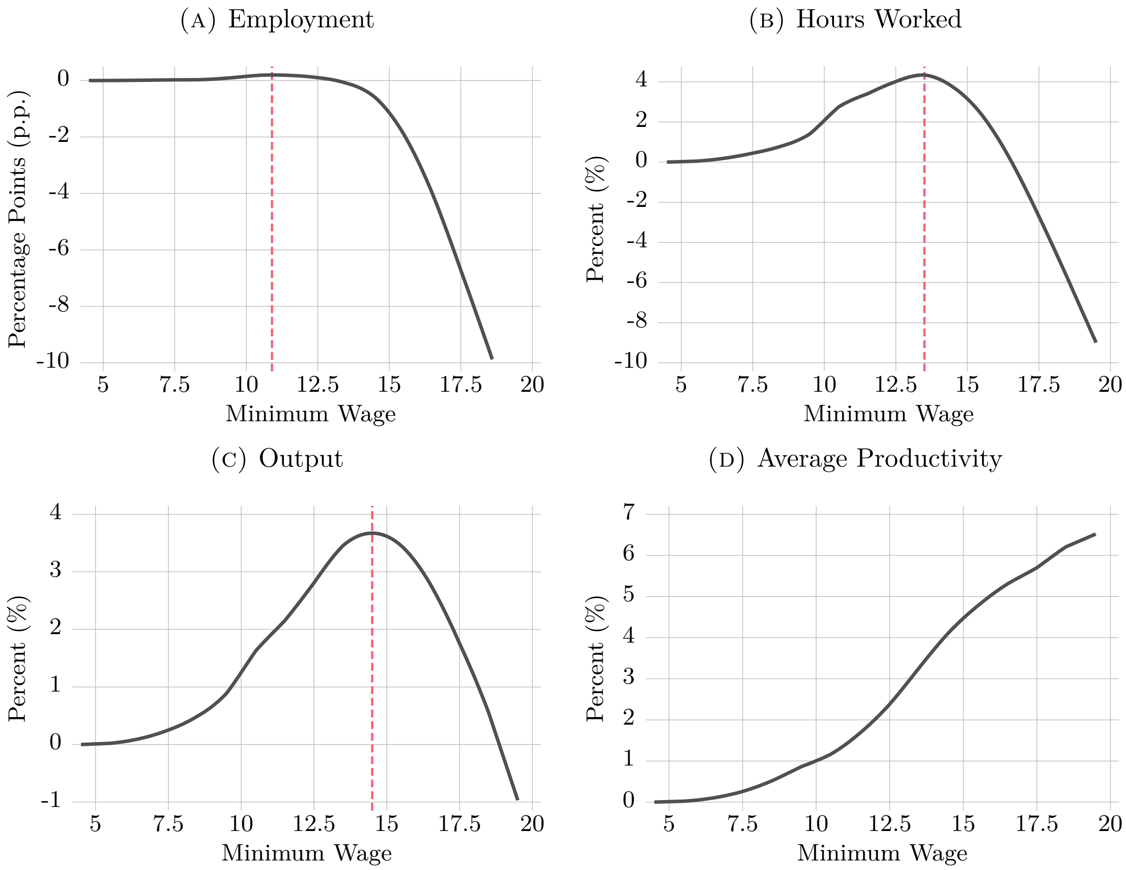

Notes: This figure shows the predicted long-run minimum wage effects on employment, hours worked, output and average productivity in panels A through D respectively. The red dashed lines in panels A, B and C indicate the maximum.

Panel A shows that endogenous search effort and, in particular, the increase in the expected surplus of meeting a firm is responsible for the lack of disemployment effect for moderate minimum wages in general equilibrium. If workers' are not allowed to adjust their search behavior, employment decreases monotonically as the minimum wage exceeds 7.5 EUR. If workers are allowed to re-optimize, but the surplus of contacting a firm is held at its baseline level, the drop in employment is even more pronounced. This is because workers now reduce their search effort as firms post fewer vacancies and the aggregate contact rate drops (search effectiveness). In contrast, when the number of vacancies is held constant and workers' can adjust their search effort, total employment increases significantly in the minimum wage. The negligible employment effects for moderate minimum wages is therefore not due to a muted reduction of firms' vacancy posting, but rather the net effect of two off-setting forces. In other words, changes in the demand and supply of labor largely offset each other for moderate minimum wage hikes. As the minimum wage approaches

Figure 13: Long-Run Minimum Wage Effects – Mechanisms

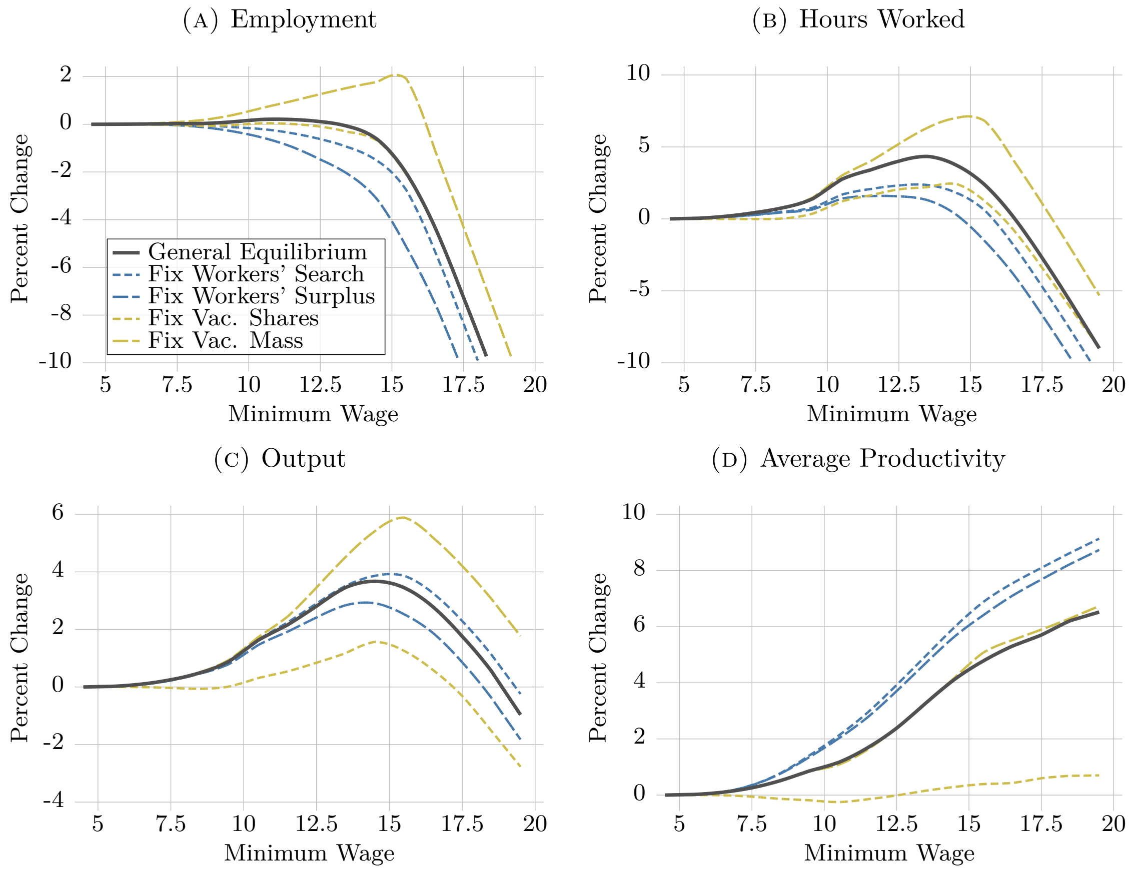

Notes: This figure shows the predicted relative changes in total employment, total hours worked, total output and average productivity as a function of the minimum wage for different scenarios. The blue short-dashed line shows a scenario where workers' search effort is held fixed at the pre-reform levels. The blue long-dashed line refers to a scenario that allows search effort to adjust, but keeps the surplus of meeting a firm constant. The yellow short-dashed line displays a scenario where the hours-and productivity distribution of vacancies is held fixed while the total mass of vacancies is allowed to adjust. The yellow long-dashed line refers to a scenario where the distribution of vacancies is flexible but the mass of vacancies is fixed.

15 EUR, however, broad non-employability of low-skill workers kicks in and employment falls.[34]

Changes in the productivity and employment level distribution of vacancies have only a small impact on total employment, they do impact the response of both total hours worked and average productivity (Panels B and D). It is clear that the profitability of low-productivity jobs declines relative to that of high-productivity jobs for a given minimum wage. Optimal behavior by firms implies that asymmetric declines in profitability lead to asymmetric vacancy reductions. This can be seen in Panel A of Figure 14 which plots the equilibrium productivity distribution of all vacancies for selected minimum wages. Clearly, the minimum wage shifts this distribution to the right. Conditional on meeting a firm, average firm productivity thus increases and workers move to more productive firms.

[34]Note that in the scenario where the mass of vacancies is held fixed, I exclude non-employable vacancies.

Figure 14: Productivity Distribution of Vacancies

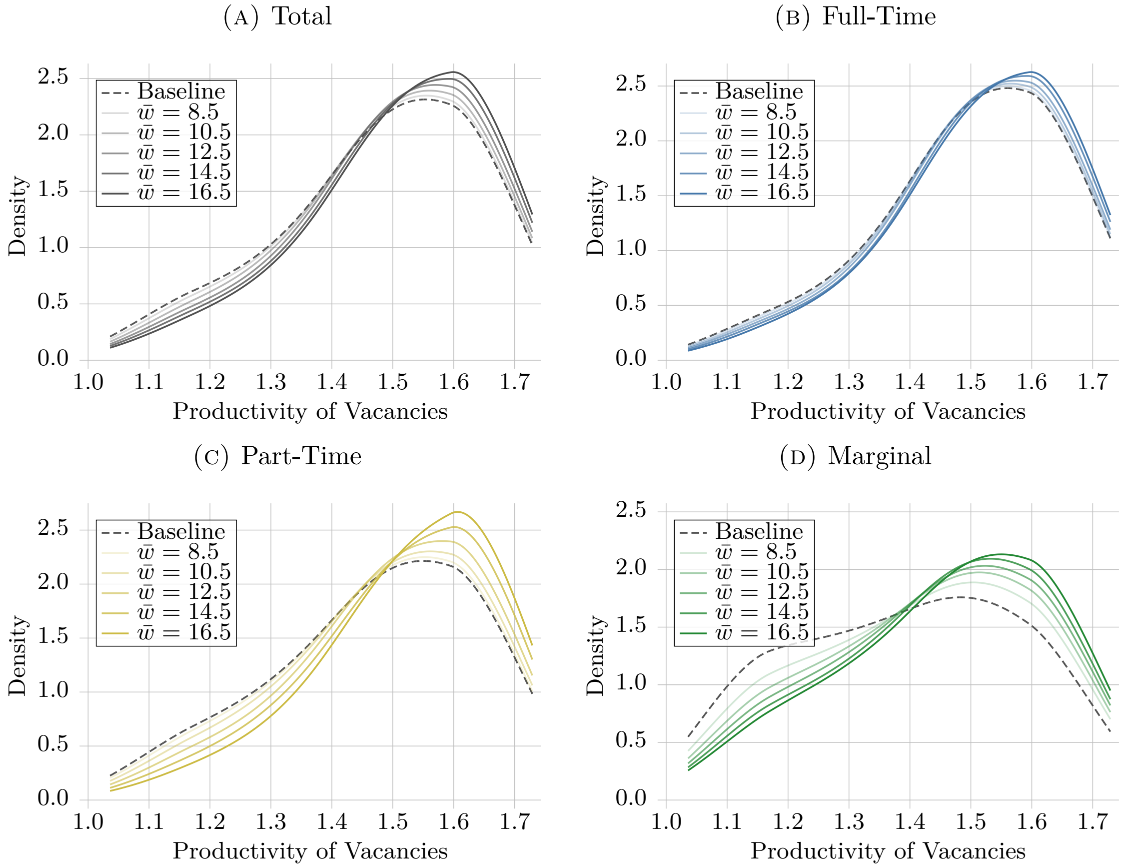

Notes: This figure shows how the productivity distribution of vacancies offered by firms changes with the minimum wage. Panel A shows the distribution for all vacancies. Panels B, C and D show the productivity distribution of full-time, part-time and mini-job vacancies respectively. I exclude skill segments in which none of the minimum wages considered is binding for any job, i.e. where all wages in the baseline equilibrium are above 16.5 EUR.

Panel D shows that fixing the composition of vacancies kills the positive productivity effect of minimum wages.

Panels B, C and D of Figure 14 further show that full-time vacancies are relatively more productive than vacancies for part-time and especially marginal jobs. Hence, the response of firms' vacancy posting is not only asymmetric in terms of productivity but also employment levels. Therefore, fixing the composition of vacancies also mutes the effect on total hours worked as there is less reallocation towards high-hours jobs.

While firms' vacancy posting decisions drive up total hours worked and average productivity in response to a minimum wage hike, the optimal response of workers increases hours worked but decreases average firm productivity. On the one hand, a binding minimum wage increases the relative value of high-hours jobs and thus increases the incentives of marginal and part-time workers to engage in on-the-job search for a job with a higher employment level. This is because, for a given increase in the hourly wage, earnings and therefore consumption growth is higher for jobs with higher employment levels.[35] Therefore, fixing workers' search effort or their surplus of meeting a firm at the baseline levels reduces the positive hours effect of minimum wages by about 50%. Conversely, a binding minimum wage reduces the surplus of working for a high productivity firm as the minimum wage eliminates or reduces productivity-related wage differentials. This reduces the incentives for on-the-job search and the probability that a worker at a low-productivity firm will accept a job offer from a high-productivity firm (with the same employment level). Hence, fixing this adjustment mechanism amplifies the positive productivity effect of increasing the minimum wage.

Finally, Panel C shows how total output evolves in these partial equilibrium scenarios. Allowing firms' to adjust their vacancy posting decisions reduces output because the total mass of vacancies and thus employment drops but increases output because the hours- and productivity distribution of vacancies shifts toward more productive and full-time jobs. Allowing workers to adjust their search effort has only small effects on total output as it has a positive effect on employment and hours worked but a negative effect on average productivity.

In sum, the steady state analysis shows that increasing the minimum creates a trade-off between employment and output. Policy makers can use the minimum wage to improve the average productivity and employment level of jobs and thereby average output per job. However, the model predicts that, for minimum wages beyond a Kaitz index of 60%, improved job composition has to be traded off against total employment.

5.2 Transition Dynamics: Long-Run Gain vs. Short-Run Pain

The steady state comparisons show that reallocation is crucial in order to understand the effects of increasing the minimum wage. In a world of search frictions, reallocation will take time and can be quite painful. Figure 15 shows how many jobs in the baseline equilibrium will become unprofitable for different minimum wages. The higher the minimum wage, the more jobs will be destroyed following the minimum wage hike. While initial job destruction is not important for minimum wages below 10 EUR, it is increasingly important for higher minimum wages. At the long-run output maximum of 14.4 EUR, for example, over 10% of all jobs are destroyed initially.

These workers become unemployed and have to find a new (better) job which takes time and effort. Taking search frictions seriously thus requires one to analyze the entire transition path following a minimum wage hike.[36] Figure 16 shows how the unemploy-

[35] Note that the minimum wage does not change the disutility of working long hours. Hence, whatever the initial relative value of full-time jobs, it will increase in the minimum wage. This is consistent with the theoretical and empirical results presented by Doppelt (2019) who analyzes a stylized model where workers can choose the number of hours worked.

[36] I allow for a notice period of one quarter. See Appendix C for details on how to compute the transition path.

FIGURE 15: Initial Job Destruction

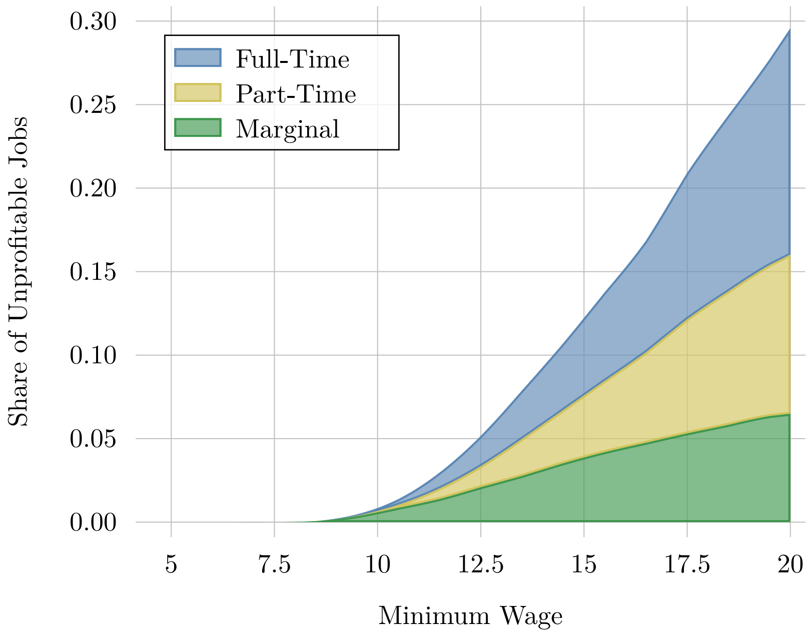

Notes: This figure shows what share of jobs in the baseline equilibrium becomes unprofitable for different minimum wages. The different areas decompose the total share into different employment levels.

ment rate and output evolve following minimum wage hikes of different magnitudes. In particular, the black line corresponds to the observed minimum wage reform (8.5 EUR) and the green, yellow and red lines correspond to the long-run employment maximum (11 EUR), the highest minimum wage without long-run disemployment effects (13.0 EUR), and the long-run output-maximizing minimum wage (14.4 EUR) respectively.

As expected, we see significant spikes in the unemployment rate at the time the minimum wage is imposed (Panel A). At the long-run employment-maximizing minimum wage of 11 EUR, the unemployment rate increases by about 1.5 percentage points on impact (increase of 20%). While a minimum wage of 13.0 EUR does not lead to disemployment effects in the long-run, it does so in the short- and medium run as the economy takes about 10 to 15 years to converge to the new stationary equilibrium. At the long-run output-maximizing minimum wage of 14.4 EUR, the unemployment rate more than doubles following the reform and takes three years to fall below 10%. After five years, the unemployment rate is still 24% above the baseline and 15% above the new long-run unemployment rate.

Panel B shows that output gains also take time to materialize. For minimum wages above 11 EUR, the minimum wage hike forces the economy into a recession. On the way to the long-run output maximum, output falls below its baseline level for almost two years. Nevertheless, output gains take less time to kick in than it takes the unemployment rate to drop. This is because high short-run unemployment is mostly driven by workers at the

FIGURE 16: Minimum Wage Effects Along the Transition Path

(A) Unemployment (B) Output

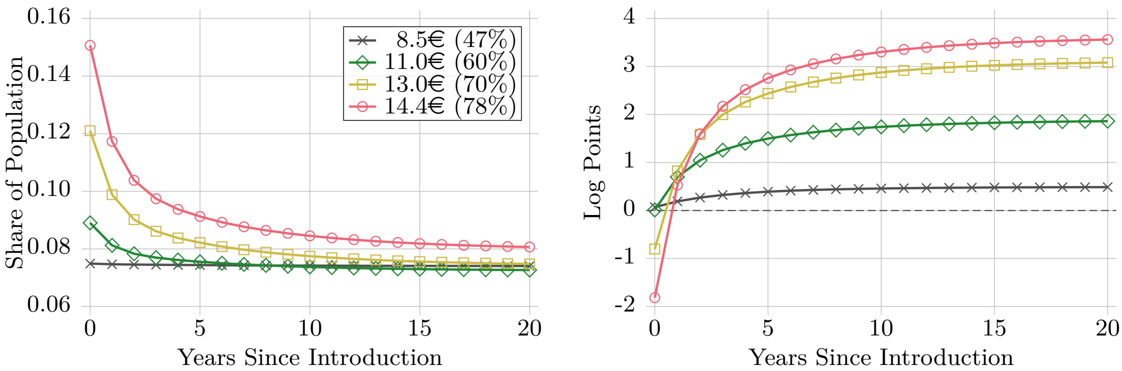

Notes: This figure shows how the unemployment rate and output evolve following minimum wage hikes of different magnitude – always starting at the baseline equilibrium without a minimum wage. I assume that the minimum wage hike is announced one quarter before it becomes binding.

bottom end of the skill distribution and the contribution of these workers to total output is relatively small. The trade-off between output and employment thus becomes magnified in the short- and medium run.

In order to formalize how policy makers' planning horizon affects their assessment of higher minimum wages, Figure 17 shows how the average discounted unemployment rate (panel A) and the net present value of output (panel B) evolve as a function of the minimum wage and for different time horizons T. Panel A shows the average discounted unemployment rate between $t = 0$ and $t = T$:

$$\frac{1}{\sum_{t=0}^{T} \beta^t} \sum_{t=0}^{T} \beta^t u_t(\bar{w}) \tag{26}$$

where $u_t(\bar{w})$ is the unemployment rate t periods after the minimum wage is raised from zero to $\bar{w}$ and the workers' discount factor ($\beta = 0.98$) is used. Panel B shows the log difference in the net present value of output relative to the baseline without a minimum wage for time horizon T:

$$\log\left(\sum_{t=0}^{T} \beta^t Y_t(\bar{w})\right) - \log\left(\sum_{t=0}^{T} \beta^t Y_t(0)\right) \tag{27}$$

where $Y_t(\bar{w})$ is total output in period t after a minimum wage of $\bar{w}$ was introduced. The lighter the line, the longer the time horizon T. The darkest line corresponds to $t = 0$ and the red dashed line corresponds to $T = \infty$. The lines in between show the change in output and the unemployment rate 1, 2, 5, 10 and 20 years after the introduction of the minimum wage.

FIGURE 17: Discounted Output and Employment Effects

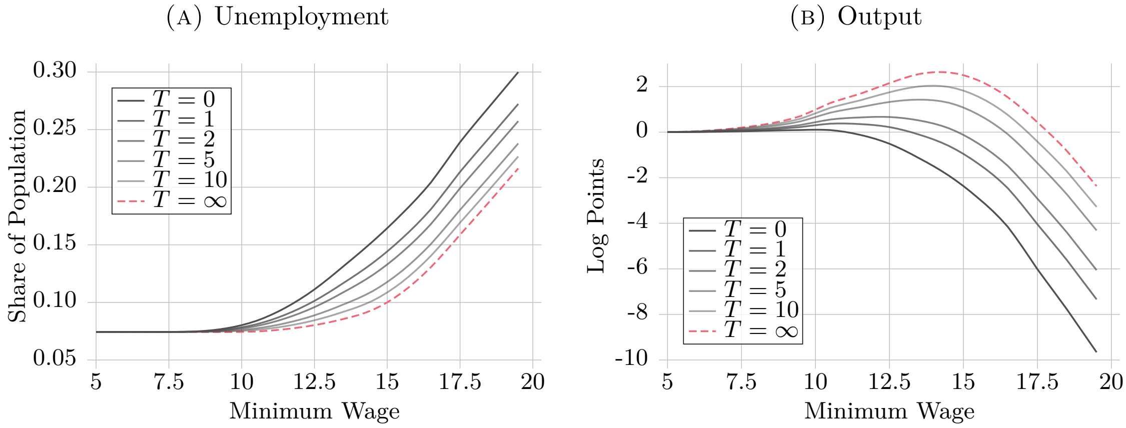

Notes: This figure shows and the average discounted unemployment rate (panel A) and the net present value of output (panel B) of different minimum wage policies for different time horizons. Lighter lines correspond to longer time horizons. The red dashed line corresponds to an infinite time horizon.

Panel A of Figure 17 shows that there is no binding minimum wage that decreases the average discounted unemployment rate – even for an infinite time horizon. The long-run reduction in the unemployment rate is not big enough to outweigh higher short-run unemployment rates. The long-run net present value of output is maximized at a minimum wage of 14.2 EUR and 2.68 log points above the net present value of output without a minimum wage. Adopting a five and two year horizon, the output maximizing minimum wage drops to 13.5 EUR and 12.2 EUR respectively with smaller but still significant discounted output gains of 1.43 and 0.67 log points.

5.3 Who Benefits from High Minimum Wages?

As a final step, I use the model to analyze how increasing the minimum wage affects lifetime utility of workers. Panel A of Figure 18 shows that average lifetime utility in the population increases up until a minimum wage of over 16 EUR. However, at the per capita optimum, a significant share of low-skill workers experiences large welfare losses. For the bottom 5, 10 or 20 percent of the human capital distribution, average lifetime utility peaks between 13 and 14 EUR and declines sharply to the right of the optimum. This is because high minimum wages make low-skill workers unemployable and forces them into long-term unemployment. While low-skill workers are the ones who benefit the most from increasing the minimum wage, they also suffer the most if the minimum wage is set so high that they become unemployable.

Panel B shows how lifetime utility changes for different sociodemographic groups. As with the initial minimum wage, welfare gains are not distributed equally. While average lifetime utility of men and single women without kids grows strongly with higher minimum wage levels, single women with kids and married women do not benefit from the

Figure 18: Heterogeneous Welfare Effects

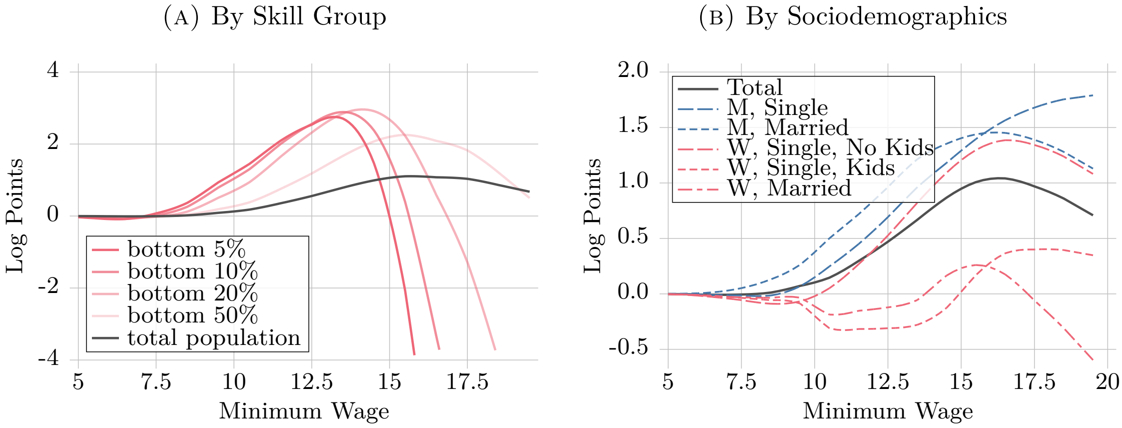

Notes: This figure shows how the minimum wage changes average lifetime utility of different sub-populations. Panel A distinguishes between different parts of the human capital distribution and Panel B presents the effects by sociodemographic characteristics.

reallocation effects in terms of their lifetime utility. The latter actually experience small welfare losses for moderately high minimum wages between 10 and 14 EUR. Reallocation away from low-hours jobs comes at a disutility cost of longer working hours. This disutility is estimated to be substantially larger for single women with kids and married women reflecting the large share of non-full-time jobs among these workers in the baseline equilibrium. Time constraints due to childcare obligations that feature into this disutility thus interact with the minimum wage.

In the end, the minimum wage remains a crude policy tool that is not targeted at certain skill levels or sociodemographics. Hence, different sub-populations do not benefit equally from higher minimum wages even though aggregate welfare – measured by the average lifetime utility across all workers – increases in the minimum wage.

5.4 Different Baseline: Increasing the Minimum Wage Starting at 8.5 Euro

Thus far, the analysis was conducted using the economy without a minimum wage as the baseline. While this is useful to understand the short- and long-run effects, most countries – including Germany today – already have some positive minimum wage in place. I now assume that the economy has already converged to the stationary equilibrium with a minimum wage of 8.5 EUR – the initial level set in 2015. Figure 16 suggests that this is a reasonable assumption as both output and employment are close to the new stationary equilibrium five years after the reform.

Analogous to Figure 17, Figure 19 shows how the net present value of output and the average unemployment rate evolve at different time horizons as a function of the minimum wage using the 8.5 EUR scenario as the point of departure. However, the

Figure 19: Discounted Output and Employment Effects with Baseline at 8.5 EUR

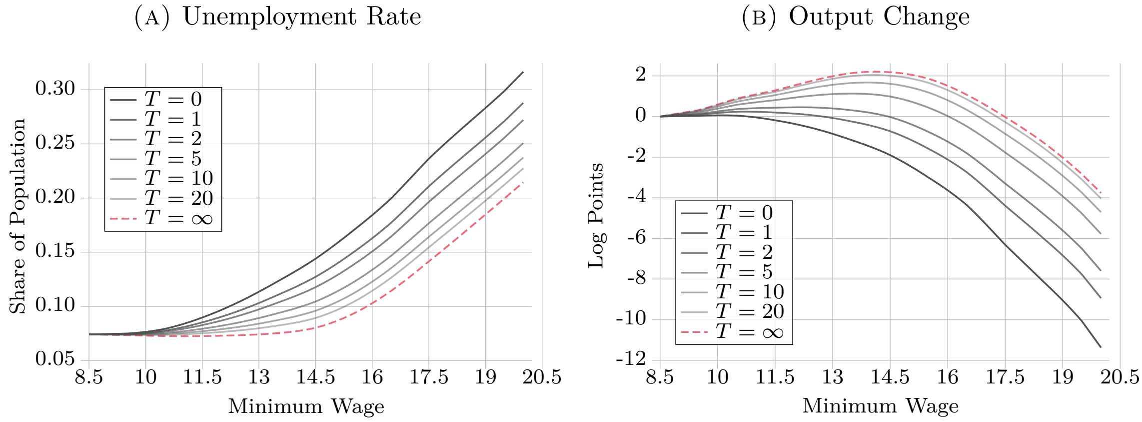

Notes: This figure shows and the average discounted unemployment rate (panel A) and the net present value of output (panel B) of different minimum wage policies for different time horizons starting from an equilibrium with a minimum wage of 8.5 EUR. Lighter lines correspond to longer time horizons. The red dashed line corresponds to an infinite time horizon.

results change only marginally. The reason for this is that a minimum wage of 8.5 EUR is still relatively low and does not alter the starting point enough to affect the transition quantitatively. Hence, policy makers still face the same trade-offs when thinking about raising the minimum wage with a positive minimum wage already in place.

6 Discussion & Conclusion

The main goal of this paper is to construct a rich quantitative search-matching model that is consistent with recent reduced-form evidence on employment and reallocation effects of observed minimum wages (Cengiz et al., 2019; Dustmann et al., 2020) and can be used to analyze how output, employment and welfare react to increasing the minimum wage. This analysis is motivated by recent proposals to increase minimum wages in developed countries and the lack of quantitative structural models that can inform policy makers. While more research is clearly necessary to better understand how the economy would react to higher minimum wages, this paper takes multiple steps in that direction.

I show that a rich model with two-sided heterogeneity, endogenous search and vacancy posting, a realistic tax- and transfer system and multiple employment levels can not only match important aspects of pre-reform micro data but also replicate the available reduced form evidence on the German minimum wage (Dustmann et al., 2020). In particular, the minimum wage of 8.5 EUR had negligible employment effects while increasing productivity, wages and the average employment level.

I use the estimated and tested model to analyze how increasing the minimum wage affects employment and output in the short- and long-run. The analysis offers at least four important insights. First, total employment increases slightly in the minimum wages up to a Kaitz index of 60% as increasing search effort by workers offsets decreasing vacancy posting by firms. Allowing not only firms but also workers to changes in the respective value of an employment relationship can thus explain why past and current minimum wage hikes have not lead to significant disemployment effects Cengiz et al. (2019). Second, while increasing the minimum wage above 60% of the median wage leads to a decline in the number of jobs, total hours worked continue to increase and peak at a Kaitz index of 73%. This shows that taking into account differences in employment levels and the fact that relatively many low-wage jobs are also low-hours jobs is crucial in order to understand how the minimum wage affects the composition of jobs and hence average output per job. Before the introduction of the minimum wage in Germany, full-time jobs accounted for only one third of all jobs affected by the minimum wage of 8.5 EUR (Kaitz index of 47%). Moreover, the composition of jobs shifts towards more productive firms as low-productivity firms become relatively less profitable. While most of the literature on minimum wages has focused on (dis-)employment effects of minimum wages, these results suggest that analyzing job composition and output effects is equally important. In the model, a Kaitz index of 78% maximizes steady state output in the economy. Third, I show that it is paramount to take search frictions seriously and analyze the transition dynamics following a minimum wage hike. Since large minimum wage hikes destroy jobs on impact, unemployment rates spike. Search frictions and lower vacancy posting make finding a (better) job very time-consuming. Long run output gains thus go hand-in-hand with substantial short- and medium-run unemployment rates and even recessions of up to two years. Fourth, not all workers benefit from higher minimum wages. Low-skill workers become non-employable and are stuck in long-term unemployment. Moreover, many women who rely on marginal jobs that are replaced by part-time and full-time jobs experience substantially smaller welfare gains compared to men. The results thus suggest that providing adequate child care opportunities is important to allow all workers to benefit from higher minimum wages.

Against the backdrop of these results, a number of avenues for future research seem particularly fruitful. First, my paper – as well as the literature as a whole – abstracts from real-world features that may become more important for large minimum wages. First, I suspect that, at some point, high minimum wages will lead firms in the tradeable sector to consider moving their workforce to countries with lower wage floors. Empirical analyses of past reforms point to somewhat stronger disemployment effects in the tradeable sector (e.g. Cengiz et al., 2019). As wages in the tradeable sector are relatively high even without minimum wages, this distinction has been quantitatively unimportant. Analyzing to what extent and at what point firms decide to relocate to other countries will be important to assess the costs and benefits of higher minimum wages. As for taxation, international cooperation may become important for minimum wage laws.

Second, more work is needed to assess whether the large output effects are robust to changes in the production function. Complementarities between low- and high-skill tasks can reduce demand for high-skill jobs as low-skill jobs become non-profitable, but may also limit the decline in vacancy posting for low-skill jobs as they are required for more important high-skill tasks. In addition, it will be fruitful to investigate how firms' investment decisions are affected by the minimum wage. Will high minimum wages lead firms to replace labor with capital or lead to higher productivity growth.

Third, exploring the effects on endogenous human capital accumulation seems important. On the one hand, higher minimum wages may decrease workers' incentives to invest in their education as wage differentials are reduced. On the other hand, the disappearance of jobs in low-skill segments of the labor market will increase human capital accumulation.

Fourth, the inter-temporal trade-offs associated with higher minimum wages deserve more attention. At what speed should policy makers implement increases in the minimum wage in order to mitigate the short-run losses while still benefiting from productivity- and output-enhancing reallocation effects?

Finally, the effects of the minimum wage interact with other labor market policies such as the design of unemployment insurance or earned income tax credits. As both lower unemployment benefits and higher minimum wages affect workers' surplus of employment, the optimal generosity of the social safety net and the level of the minimum wage should be determined jointly.

Bibliography

Abowd, J., F. Kramarz, and D. Margolis (1999): “High Wage Workers and High Wage Firms,” *Econometrica*, 67, 251–333.

Acemoglu, D. (2001): “Good Jobs versus Bad Jobs,” *Journal of Labor Economics*, 19, 1–21.

Antoni, M., A. Ganzer, P. Vom Berge, et al. (2016): “Sample of Integrated Labour Market Biographies (SIAB) 1975-2014,” FDZ-Datenreport 04/2016, Institute for Employment Research.

Arnoud, A., F. Guvenen, and T. Kleineberg (2019): “Benchmarking Global Optimizers,” Working Paper 26340, National Bureau of Economic Research.

Bagger, J., F. Fontaine, F. Postel-Vinay, and J.-M. Robin (2014): “Tenure, Experience, Human Capital, and Wages: A Tractable Equilibrium Search Model of Wage Dynamics,” *American Economic Review*, 104, 1551–96.

Bagger, J. and R. Lentz (2018): “An Empirical Model of Wage Dispersion with Sorting,” *The Review of Economic Studies*, 86, 153–190.

Blömer, M., N. Guertzgen, L. Pohlan, H. Stichnoth, and G. J. Van den Berg (2020): “Unemployment Effects of the German Minimum Wage in an Equilibrium Job Search Model,” .

Bonhomme, S., T. Lamadon, and E. Manresa (2019): “A Distributional Framework for Matched Employer Employee Data,” *Econometrica*, 87, 699–739.

Bontemps, C., J.-M. Robin, and G. J. Van den Berg (1999): “An Empirical Equilibrium Job Search Model with Search on the Job and Heterogeneous Workers and Firms,” *International Economic Review*, 40, 1039–1074.

Bossler, M. and H.-D. Gerner (2016): “Employment Effects of the New German Minimum Wage: Evidence from Establishment-level Micro Data,” IAB-Discussion Paper 10/2016, Institute for Employment Research.

Brenzel, H., J. Czepek, H. Kiesl, B. Kriechel, A. Kubis, A. Moczall, M. Rebien, C. Roettger, J. Szameitat, A. Warning, et al. (2016): “Revision of the IAB Job Vacancy Survey: Backgrounds, Methods and Results,” *IAB-Forschungsbericht*.

Burauel, P., M. Caliendo, M. M. Grabka, C. Obst, M. Preuss, C. Schröder, and C. Shupe (2020): “The Impact of the German Minimum Wage on Individual Wages and Monthly Earnings,” *Jahrbücher für Nationalökonomie und Statistik*, 240, 201–231.

Burdett, K. and D. T. Mortensen (1998): “Wage Differentials, Employer Size, and Unemployment,” *International Economic Review*, 257–273.

Cahuc, P., F. Postel-Vinay, and J.-M. Robin (2006): “Wage Bargaining with On-the-Job Search: Theory and Evidence,” *Econometrica*, 74, 323–364.

Caliendo, M., A. Fedorets, M. Preuss, C. Schröder, and L. Wittbrodt (2017): “The Short-Run Employment Effects of the German Minimum Wage Reform,” IZA Discussion Papers 11190, Institute for the Study of Labor (IZA).

Card, D., J. Heining, and P. Kline (2013): "Workplace Heterogeneity and the Rise of West German Wage Inequality," *The Quarterly Journal of Economics*, 128, 967.

Cengiz, D., A. Dube, A. Lindner, and B. Zipperer (2019): "The Effect of Minimum Wages on Low-wage Jobs," *The Quarterly Journal of Economics*, 134, 1405–1454.

Chetty, R. (2008): "Moral Hazard versus Liquidity and Optimal Unemployment Insurance," *Journal of Political Economy*, 116, 173–234.

Di Addario, S., P. Kline, R. Saggio, and M. Solvsten (2020): "It Ain't Where You're From, It's Where You're at: Hiring Origins, Firm Heterogeneity, and Wages," Tech. Rep. 104-20, IRLE Working Paper.

Doppelt, R. (2019): "Minimum Wages and Hours of Work," *mimeo*.

Dube, A. (2019): ""Impacts of minimum wages: Review of the international evidence," *Independent report, https://www. gov. uk/government/publications/impacts-of-minimum-wages-review-of-the-international-evidence*.

Dube, A., T. W. Lester, and M. Reich (2016): "Minimum Wage Shocks, Employment Flows, and Labor Market Frictions," *Journal of Labor Economics*, 34, 663–704.

Dustmann, C., A. Lindner, U. Schönberg, M. Umkehrer, and P. vom Berge (2020): "Reallocation Effects of the Minimum Wage," CReAM Discussion Paper Series 2007, Centre for Research and Analysis of Migration (CReAM), Department of Economics, University College London.

Eckstein, Z. and K. I. Wolpin (1990): "Estimating a Market Equilibrium Search Model from Panel Data on Individuals," *Econometrica*, 783–808.

Engbom, N. and C. Moser (2018): "Earnings Inequality and the Minimum Wage: Evidence from Brazil," *Federal Reserve Bank of Minneapolis-Opportunity and Inclusive Growth Institute Working Paper*, 7, 18–50.

Flinn, C. J. (2006): "Minimum Wage Effects on Labor Market Outcomes under Search, Matching, and Endogenous Contact Rates," *Econometrica*, 74, 1013–1062.

Garloff, A. (2016): "Side Effects of the New German Minimum Wage on (Un-)Employment : First Evidence from Regional Data," IAB Discussion Paper 201631, Institut für Arbeitsmarkt- und Berufsforschung (IAB), Nürnberg.

Hagedorn, M., F. Karahan, I. Manovskii, and K. Mitman (2019): "Unemployment Benefits and Unemployment in the Great Recession: The Role of Macro Effects," Working Paper 19499, National Bureau of Economic Research.

Harasztosi, P. and A. Lindner (2019): "Who Pays for the Minimum Wage?" *American Economic Review*, 109, 2693–2727.

Hartung, B., P. Jung, and M. Kuhn (2018): "What Hides behind the German Labor Market Miracle? Unemployment Insurance Reforms and Labor Market Dynamics," CEPR Discussion Paper DP13328, Centre for Economic Policy Research.

Holtemöller, O. and F. Pohle (2017): "Employment Effects of Introducing a Minimum Wage: The Case of Germany," Tech. rep., IWH Discussion Papers.

Jolivet, G., F. Postel-Vinay, and J.-M. Robin (2006): "The Empirical Content of the Job Search Model: Labor Mobility and Wage Distributions in Europe and the US," *European Economic Review*, 50, 877 – 907.

Jäger, S., B. Schoefer, S. Young, and J. Zweimüller (2020): "Wages and the Value of Nonemployment," *The Quarterly Journal of Economics*, qjaa016.

Krause, M. U. and H. Uhlig (2012): "Transitions in the German Labor Market: Structure and Crisis," *Journal of Monetary Economics*, 59, 64 – 79, carnegie-NYU-Rochester Conference Series on Public Policy at New York University on April 15-16, 2011.

Krebs, T. and M. Scheffel (2013): "Macroeconomic Evaluation of Labor Market Reform in Germany," *IMF Economic Review*, 61, 664–701.

——— (2016): "Structural Reform in Germany," Tech. rep., IZA Discussion Papers 9787.

Ljungqvist, L. and T. J. Sargent (1998): "The European Unemployment Dilemma," *Journal of Political Economy*, 106, 514–550.

Marinescu, I. E. and D. Skandalis (forthcoming): "Unemployment insurance and job search behavior," *Quarterly Journal of Economics*.

Meyer, B. D. (1990): "Unemployment Insurance and Unemployment Spells," *Econometrica*, 58, 757–782.

Minimum Wage Commission (2018): "Second Report on the Effects of the Minimum Wage," *Berlin*.

Mitman, K. and S. Rabinovich (2019): "Do Unemployment Benefit Extensions Explain the Emergence of Jobless Recoveries?" .

Mortensen, D. T. and C. A. Pissarides (1994): "Job Creation and Job Destruction in the Theory of Unemployment," *The review of economic studies*, 61, 397–415.

Portugal, P. and A. R. Cardoso (2006): "Disentangling the Minimum Wage Puzzle: An Analysis of Worker Accessions and Separations," *Journal of the European Economic Association*, 4, 988–1013.

Postel-Vinay, F. and J.-M. Robin (2002): "Equilibrium Wage Dispersion with Worker and Employer Heterogeneity," *Econometrica*, 70, 2295–2350.

Schmieder, J. F., T. von Wachter, and S. Bender (2012): "The Long-Term Effects of UI Extensions on Employment," *American Economic Review*, 102, 514–19.

——— (2016): "The Effect of Unemployment Benefits and Nonemployment Durations on Wages," *American Economic Review*, 106, 739–77.

Shephard, A. (2017): "Equilibrium Search and Tax Credit Reform," *International Economic Review*, 58, 1047–1088.

van den Berg, G. J. and G. Ridder (1998): "An Empirical Equilibrium Search Model of the Labor Market," *Econometrica*, 66, 1183–1221.

vom Berge, P., S. Kaimer, S. Copestake, J. Eberle, and T. Haepp (2016): "Arbeitsmarktspiegel: Entwicklungen nach Einführung des Mindestlohns (Ausgabe 2)," Tech. rep., IAB-Forschungsbericht.

A Additional Tables and Figures

Table A.1: Model Fit – Job-to-Job Transitions

	full-time	part-time	mini-job
Job-to-job transition			
Data	0.028	0.034	0.088
Model	0.039	0.046	0.062
Godfather shock			
Data	0.017	0.022	0.065
Model	0.017	0.022	0.050

Notes: This table shows the probability of job-to-job transitions for full-time, part-time and mini-job workers. The top panel shows the probability of any job-to-job transition and the bottom panel shows the probability of being hit by the Godfather shock. Data: SIAB.

Table A.2: Model Fit – Other Moments

	Model	Data
Job Vacancy Rate		
Job Vacancy Rate	0.034	0.025
Job Vacancy Rate (full-time)	0.029	–
Job Vacancy Rate (part-time)	0.040	–
Job Vacancy Rate (mini-job)	0.054	–
Firm Size Distribution		
Mean of log firm size	3.745	4.136
Std. dev. of log firm size	0.928	2.187
Mean of log firm size (full-time)	3.550	4.147
Std. dev. of log firm size (full-time)	1.038	2.173
Mean of log firm size (part-time)	2.199	2.976
Std. dev. of log firm size (part-time)	0.753	2.039
Mean of log firm size (mini-job)	1.132	1.927
Std. dev. of log firm size (mini-job)	0.130	1.707

Notes: This table shows job vacancy rates and moments of the firm size distribution in the model and the data. Data: BHP, own calculations.

TABLE A.3: Model Fit – Employment Status

	$\Pr(e_{ft}\|e)$	$\Pr(e_{pt}\|e)$	$\Pr(e_{mj}\|e)$	$\Pr(u)$	$\Pr(lu\|u)$
Men, Single					
Data	0.832	0.100	0.068	0.109	0.526
Model	0.849	0.087	0.065	0.079	0.556
Men, Married					
Data	0.908	0.059	0.033	0.040	0.454
Model	0.900	0.078	0.022	0.053	0.482
Women, Single, No Kids					
Data	0.666	0.224	0.110	0.068	0.520
Model	0.697	0.193	0.110	0.081	0.526
Women, Single, Kids					
Data	0.330	0.534	0.136	0.140	0.552
Model	0.230	0.629	0.142	0.104	0.528
Women, Married					
Data	0.309	0.516	0.176	0.040	0.554
Model	0.195	0.633	0.172	0.085	0.488
Total					
Data	0.663	0.240	0.096	0.064	0.518
Model	0.636	0.273	0.091	0.074	0.512

Notes: This table shows the share of full-time, part-time and marginal jobs conditional on employment (columns 2-4), the unemployment rate (column 5) and the share of long-term unemployment conditional on unemployment (column 6) for each sociodemographic worker type and in the population (last panel). Data: SIAB.

Table A.4: Model Fit – Job finding Probabilities

	$\Pr(e'\|su)$	$\Pr(e'\|lu)$	$\Pr(e'\|e)$
Men, Single			
Data	0.286	0.062	–
Model	0.236	0.060	–
Men, Married			
Data	0.321	0.074	–
Model	0.306	0.080	–
Women, Single, No Kids			
Data	0.321	0.065	–
Model	0.255	0.067	–
Women, Single, Kids			
Data	0.303	0.082	–
Model	0.258	0.067	–
Women, Married			
Data	0.263	0.059	–
Model	0.322	0.078	–
Total			
Data	0.296	0.067	0.035
Model	0.285	0.071	0.043

Notes: This table shows the probability of finding a job out of short- and long-term unemployment as well as the job-to-job transition probability for each sociodemographic worker type and in the population (last panel). Data: SIAB.

TABLE A.5: Model Fit – Job Types and Sociodemographics by Wage Groups

	[0, 5.5)		**[5.5, 8.5)**		**[8.5, 12.5)**		**[12.5, 20)**		**[20, ∞)**		**Total**	
	Model	Data	Model	Data	Model	Data	Model	Data	Model	Data	Model	Data
Job Types												
Full-Time	0.028	0.269	0.358	0.331	0.519	0.558	0.678	0.708	0.766	0.802	0.636	0.663
Part-Time	0.000	0.239	0.290	0.257	0.332	0.289	0.277	0.262	0.229	0.191	0.273	0.240
Marginal	0.972	0.492	0.352	0.412	0.149	0.152	0.045	0.030	0.005	0.007	0.091	0.097
Gender												
Men	0.253	0.312	0.382	0.302	0.438	0.416	0.519	0.519	0.621	0.642	0.520	0.517
Women	0.747	0.688	0.618	0.698	0.562	0.584	0.481	0.481	0.379	0.358	0.480	0.483
Sociodemographics												
Men, Single	0.154	–	0.163	–	0.173	–	0.211	–	0.256	–	0.213	–
Men, Married	0.099	–	0.219	–	0.265	–	0.308	–	0.364	–	0.307	–
Women, Single, No Kids	0.229	–	0.196	–	0.182	–	0.172	–	0.142	–	0.167	–
Women, Single, Kids	0.063	–	0.060	–	0.054	–	0.043	–	0.035	–	0.045	–
Women, Married	0.455	–	0.363	–	0.326	–	0.266	–	0.202	–	0.269	–

Notes: This table shows the distribution of job-types, gender and sociodemographics within different parts of the wage distribution. Data: SIAB.

Table A.6: Model Fit – Wage Groups by Job Types and Sociodemographics

	$[0, 5.5)$		$[5.5, 8.5)$		$[8.5, 12.5)$		$[12.5, 20)$		$[20, \infty)$	
	Model	Data	Model	Data	Model	Data	Model	Data	Model	Data
Job Types										
Full-Time	0.000	0.007	0.057	0.048	0.180	0.156	0.358	0.367	0.404	0.422
Part-Time	0.000	0.018	0.108	0.103	0.269	0.224	0.342	0.376	0.282	0.279
Marginal	0.066	0.097	0.393	0.442	0.358	0.317	0.165	0.117	0.019	0.027
Gender										
Men	0.003	0.011	0.075	0.056	0.186	0.150	0.336	0.347	0.401	0.436
Women	0.010	0.025	0.131	0.140	0.258	0.227	0.336	0.346	0.265	0.262
Sociodemographics										
Men, Single	0.004	–	0.078	–	0.180	–	0.334	–	0.404	–
Men, Married	0.002	–	0.073	–	0.190	–	0.337	–	0.398	–
Women, Single, No Kids	0.008	–	0.119	–	0.240	–	0.346	–	0.286	–
Women, Single, Kids	0.009	–	0.136	–	0.265	–	0.328	–	0.262	–
Women, Married	0.010	–	0.138	–	0.267	–	0.332	–	0.252	–
Total	0.006	0.018	0.102	0.097	0.220	0.187	0.336	0.347	0.336	0.352

Notes: This table shows the share of workeres in different wage groups conditional on job types, gender and sociodemographics. Data: SIAB.

Table A.7: Model Fit – Worker Clustered AKM Fixed Effects

	Total		**Men**		**Women**	
	Model	Data	Model	Data	Model	Data
P05 / P50	0.507	0.622	0.507	0.640	0.551	0.601
P10 / P50	0.570	0.687	0.570	0.699	0.619	0.677
P20 / P50	0.690	0.778	0.690	0.783	0.717	0.769
P30 / P50	0.783	0.854	0.783	0.854	0.816	0.850
P40 / P50	0.884	0.925	0.884	0.924	0.922	0.926
P60 / P50	1.090	1.088	1.141	1.092	1.133	1.080
P70 / P50	1.259	1.203	1.329	1.215	1.300	1.178
P80 / P50	1.506	1.370	1.506	1.393	1.444	1.314
P90 / P50	1.777	1.651	1.777	1.653	1.763	1.546
P95 / P50	1.996	1.864	1.996	1.884	2.168	1.760

Notes: This table shows the median and selected percentile ratios of AKM worker fixed effects for full-time jobs. Data: SIAB.

Table A.8: Model Fit – Firm Clustered AKM Fixed Effects

	Full-Time		**Part-Time**		**Marginal**	
	Model	Data	Model	Data	Model	Data
P50 / P50$_{ft}$	1.000	1.000	0.931	0.993	0.823	0.851
P05 / P50	0.798	0.702	0.789	0.689	0.868	0.683
P10 / P50	0.851	0.762	0.811	0.784	0.868	0.789
P25 / P50	0.931	0.877	0.883	0.905	0.918	0.915
P75 / P50	1.145	1.084	1.074	1.125	1.132	1.231
P90 / P50	1.145	1.171	1.229	1.152	1.216	1.312
P95 / P50	1.145	1.171	1.229	1.245	1.392	1.344

Notes: This table shows the median and selected percentile ratios of (full-time) firm productivity for full-time, part-time and marginal jobs. The full-time firm productivity is the exponential of the AKM firm fixed effects estimated on wages of full-time workers only. Data: SIAB.

Table A.9: Model Fit – Wages

	Total		Full-Time		Part-Time		Marginal		Men		Women	
	Model	Data	Model	Data	Model	Data	Model	Data	Model	Data	Model	Data
Mean	18.11	18.54	19.87	20.38	16.85	9.23	9.65	9.23	19.68	20.78	16.42	16.13
Variance (logs)	0.49	0.51	0.46	0.47	0.46	0.48	0.35	0.37	0.49	0.50	0.48	0.49
P01	5.81	5.01	6.65	5.92	6.53	5.04	4.76	4.30	6.22	5.41	5.56	4.80
P05	7.27	6.79	8.44	8.37	7.61	6.81	5.34	4.97	7.80	7.90	6.92	6.24
P10	8.41	8.18	9.67	9.92	8.35	8.11	5.82	5.51	9.06	9.60	7.89	7.35
P15	9.43	9.32	10.90	11.21	9.15	9.13	6.21	5.92	10.21	11.00	8.70	8.25
P20	10.23	10.42	11.80	12.33	9.99	10.06	6.60	6.27	11.25	12.18	9.50	9.09
P30	12.00	12.52	13.57	14.47	11.32	11.79	7.37	6.92	13.44	14.41	11.00	10.82
P40	14.00	14.63	15.75	16.42	12.96	13.52	8.02	7.55	15.30	16.44	12.52	12.69
P50	15.96	16.67	17.70	18.36	14.77	15.30	8.85	8.20	17.39	18.56	14.26	14.70
P70	21.35	21.48	23.36	23.32	19.51	19.53	10.73	9.86	23.36	23.99	18.95	19.03
P90	31.18	31.66	33.20	33.72	28.46	27.59	14.34	13.86	33.23	34.77	28.29	26.84
P95	36.34	36.00	39.76	37.27	33.37	33.28	16.63	16.86	40.35	39.19	33.21	31.88

Notes: This table shows the mean wage, variance and selected percentile of hourly wages in the data and the estimated model. The variance is taken over the log wages. The moments for men and women were targeted in the estimation. Data: SIAB.

B Data and Target Moments

In this appendix, I detail how and with what data the target moments are computed.

B.1 Data

I mainly rely on high-quality administrative data from the IAB (*Institut für Arbeitsmarkt- und Berufsforschung*). In particular, I use the SIAB data which is a two percent random sample of the integrated employment biographies collected at the IAB.[37] I use the years 2011 to 2014. The data covers all individuals in Germany, which are employed, receive unemployment benefits, are officially registered as job-seeking at the German Federal Employment Agency or (plan to) participate in programs of active labor market policies. The only workers not included in the IAB data are civil servants as they are not subject to social security contributions. Marginally employed workers, however, are included in the data even though they are also not subject to social security contributions. Information on earnings is top-coded at the social security contribution limit. This affects about 10% of all workers each year. Following Card et al. (2013), I impute top-coded earnings using Tobit regressions by year, gender, east/west, age groups and education groups.

I complement the SIAB data with survey data from the German Socio-Economic Panel (SOEP) which contains annual information on more than 15 thousand workers. In the SOEP, I drop civil servants in order to be consistent with the IAB data.

B.2 Targeted Moments

Sociodemographics The distribution of sociodemographic types conditional on gender $\Pr(j|g)$ is taken from the SOEP. The distribution of gender $\Pr(g)$ is taken from the SIAB in order to use as much administrative information as possible.

Labor Market States As the SIAB data does not contain sociodemographic information for employed workers, I have to fill some gaps with information from the SOEP while ensuring that the joint distribution of gender and labor market status remains consistent with the administrative SIAB data.

I start by computing the unemployment rate conditional on j such that it is consistent with the gender-specific unemployment rate in the SIAB

$$\Pr(u|j) = \Pr(u|j,g) = \frac{\Pr(u,j|g)}{\Pr(j|g)} = \frac{\Pr(j|u,g)\Pr(u|g)}{\Pr(j|g)} \tag{B.1}$$

where only $\Pr(j|g)$ is taken from the SOEP. The probability of long-term unemployment conditional, $\Pr(lu|u,j)$, is taken directly from the SIAB.

[37] See Antoni et al. (2016) for a detailed description of the data.

Computing the share of type-j workers who have a type-x job requires slightly more information from the SOEP:

$$\Pr(e_x|j) = \Pr(e_x|j, g) = \frac{\Pr(j|e_x, g)\Pr(e_x|g)}{\Pr(j|g)} \tag{B.2}$$

Here, only $\Pr(j|e_x, g)$ and $\Pr(j|g)$ are taken from the SOEP.

Transition Probabilities The job finding rate out of short- and long-term unemployment, $\Pr(e'|su, j)$ and $\Pr(e'|lu, j)$, can be computed using SIAB data only. As I do not target job-to-job transition probabilities by sociodemographics, they are computed as the share of workers who change their employer or job type.

Hourly Wage Quantiles To compute hourly wages based on daily earnings reported in the SIAB data, I impute average hours worked per day using data from the SOEP and job-type dependent averages reported by Dustmann et al. (2020) who have confidential information on hours for the social security data in 2014.

The average adjusted hours for full- and part-time jobs in Dustmann et al. (2020) are almost identical to the averages in the SOEP and Structure of Earnings Survey (SES).[38] The only difference between the three data sets is that, for mini-jobs, average hours worked are substantially higher in the SOEP.

For full-time jobs, I set daily hours to 7.8 which corresponds to 39 hours per week. For par-time and mini-jobs, I use the joint distribution of hours and earnings from the SOEP to take into account that some of the variation in earnings is driven by heterogeneity in hours worked. To that end, I compute the mean and standard deviation of contractual hours worked within different earnings bins. I then draw hours worked from a Normal distribution with these parameters and impose that weekly hours for part-time and mini-jobs be in the interval $[5, 35]$ and $[2, 20]$ respectively.[39] Finally I rescale the hours worked such that, on average, part-time employees work 24 hours and mini-job employees 8.7 hours per week – as reported in Dustmann et al. (2020).

Hourly wages are then computed as earnings divided by imputed hours worked. I target the 0.01, 0.05, 0.1, 0.15, 0.2, 0.3, 0.4, 0.5, 0.7, 0.9, and 0.95 quantiles of the wage distributions conditional on job type and conditional on gender (separately). In addition, I target the share of part-time and mini-jobs and the share of men in the following five wage groups $(0, 6.5)$, $[6.5, 8.5)$, $[8.5, 12.5)$, $[12.5, 20)$, $[20, \infty)$.

[38] Dustmann et al. (2020) adjust the raw contracted working hours in the social security records to account for differences in whether sick leave and overtime are included in the contractual hours.

[39] For part-time jobs, I use 500, 750, 1000, 1500, ..., 4000, 5000, 10000 Euro as cutoffs to define the monthly earnings bins. For mini-jobs, I use the cutoffs 100, 150, ..., 500 Euro.

Worker and Firm Fixed Effects In the absence of a minimum wage, the wage equation in my model is very simple. As in Abowd et al. (1999) (henceforth AKM), the wage w of a full-time worker employed at firm with productivity p is log-additive in her skill h and the firm's productivity

$$\log(w) = \log(r) + \log(h) + \log(p) \tag{B.3}$$

where r is the exogenous piece-rate. I estimate the empirical distribution of worker and firm-class fixed effects using a clustered AKM approach (Bonhomme et al., 2019).

While the model is consistent with an AKM-style wage equation (Abowd et al., 1999; Card et al., 2013), I do not estimate the model by straight AKM because of two distinct reasons. First, while the SIAB data is large compared to survey data sets, it covers only 2% of all workers and the firms they are employed at. This implies that the connected set of firms and workers on which firm and worker fixed effects can be identified is too small. Second, estimation would suffer from severe incidental parameter bias as the number of movers between two firms tends to be low.

Instead, I estimate the empirical distributions of worker and firm heterogeneity using the approach recently proposed by Bonhomme et al. (2019) (henceforth BLM) which solves both of these issues using dimension reduction techniques. The proposed method is particularly useful as it can be applied to data sets that cover only few firm-to-firm moves. The key assumption is that unobserved firm heterogeneity operates on the level of discrete firm classes rather than on the level of individual firms. Given an estimated partition of all firms into classes, firm class and worker fixed effects are identified from job-to-job transitions between firms of different classes rather than between different firms. This allows estimation of worker and firm (class) effects on much smaller samples of linked employer-employee data such as the SIAB (2%).

Class membership is estimated using K-means clustering that minimizes the within-class variation of within-firm earnings distributions:

$$\min_{k(1),\ldots,k(J),H_1,\ldots,H_K} \sum_{j=1}^{J} \frac{1}{M} \sum_{m=1}^{M} \left(F_j^m - H_{k(j)}^m\right)^2$$

where $k(j)$ is the class of firm j, F_j^m is an observable characteristic of firm j and H_k^m is the average of that characteristic across all firms in class k. I classify firms based on information on the within-firm wage distribution. In particular, I use the mean, selected percentiles (25, 50, 75) and the share of workers with top-coded earnings for full-time employees.[40] Consistent with the model where firm productivity is deterministic, I average

[40]This information is made available for every firm such that the within-firm earnings distribution can be approximated without observing a representative sample of employees for each firm.

these characteristics at the firm level over the years 2011 to 2014. This yields a time-invariant classification of firms.

Given the firm classification, I estimate the worker and firm-class fixed effects conditional on the estimated, i.e. run a clustered AKM estimation without covariates (except time fixed effects).

$$\log(w_{it}) = \alpha_i + \psi_{k(j(it))} + \gamma_t + \varepsilon_{it} \tag{B.4}$$

I then target the distribution of α conditional on gender and the worker-weighted distribution of ψ to inform the distributions of human capital and firm productivity. In particular, I target the quantile ratios $q_x^k/q_x^{0.5}$ for $k = 0.01, 0.05, 0.1, 0.3, 0.7, 0.9, 0.95, 0.99$ and $x \in \{f, p, m\}$, where q_x^k is the k-quantile of the distribution of ψ weighted by the firm's number type-x workers. In addition, I target $q_x^{0.5}/q_f^{0.5}$ for $x \in \{p, m\}$. Finally, I target the shares of the variance of log wages explained by the worker and firm components as well as the correlation between worker and firm fixed effects.

Firm Size The mean and standard deviation of log firm size are computed using administrative data from the Establishment History Panel. For consistency with the worker moments, I only consider employees between 25 and 60 years of age and drop firms that do not have employees in this age range.

Job Vacancy Rate The job vacancy rate is the number of vacancies relative to the sum of vacancies and jobs. As many vacancies are not officially registered, I do not rely on the job vacancy rate reported by Eurostat but rather use the Job Vacancy Survey (JVS).[41] The JVS contains both registered and unregistered vacancies – each account for roughly half of the total number of vacancies. In 2014, around 900 thousand vacancies were open. With roughly 36 million jobs, this gives a job vacancy rate of 2.44%.[42]

C Computational Details

C.1 Steady State

In order to compute a stationary equilibrium in the economy, I discretize the state space by using a grid of values for human capital h (30 grid points) and firm productivity p (15 grid points). I solve for the equilibrium in each skill-segment separately using the following algorithm:

1. Guess an initial distribution of vacancies across firm productivities and employment levels ($n^0(x, p)$), and a level of labor market tightness (θ^0).

[41] See Brenzel et al. (2016) for details about the data.
[42] Source: Statistics of the Federal Employment Agency.

2. Set $i = 0$
3. Taking as given the vacancy shares $n^i(x,p)$ and labor market tightness θ^i
   (a) Use equations (9), (10), (11) and (12) to solve for workers' search policies $\ell^i(j,\sigma)$ and value functions $V^i(j,\sigma)$ where σ is a point in the state space of a worker with skills h and demographics j (policy function iteration).
   (b) Compute the implied distribution of workers across states, $F(j,\sigma)$, using the search policies, equation 13, and the exogenous transition probabilities.
   (c) Compute the implied total search mass S, the search mass willing to accept a (x,p) job offer, $S(x,p)$ and the vacancy filling probabilities $\eta(x,p)$ from equations 5 and 16 and the probability that a (x,p)-job filled by a type-j worker is destroyed from equation 14.
   (d) Solve for firms' optimal vacancy policies $v^i(x,p)$ using equation (20).
   (e) Compute the implied vacancy shares $n^{i+1}(x,p)$ and labor market tightness θ^{i+1}.
4. If $n^{i+1}(x,p) \approx n^i(x,p)$ for all $x \in \{ft, pt, mj\}$ and for all p on the firm productivity grid and if $\theta^{i+1} \approx \theta^i$, stop! If not, set increment i repeat step 3!

C.2 Transition Path

Starting from the terminal stationary equilibrium, I guess a path for all equilibrium objects and solve backwards. We focus on one generic skill segment h and drop h form the notation to improve readability.

Firm Problem Assuming that the new stationary equilibrium is reached after T periods, the firm's expected value of an employment relationship with a type-j worker starting in period t is:

$$W_t^j(x,p) = y + \sum_{s=t+1}^{T-1} \beta^{s-t} \underbrace{\prod_{k=t}^{s-1}(1-\delta_k^j(x,p))}_{\text{Pr(survival until } s)} y(x,p) + \beta^T \underbrace{\prod_{k=t}^{T-1}(1-\delta_k^j(x,p))}_{\text{Pr(survival until } T)} W_T^j(x,p)$$

$$\text{with} \quad W_T^j(x,p) \equiv W^{j*}(x,p) = \frac{y(x,p)}{1-\beta(1-\delta^{j*}(x,p))}$$

Given $W_{t+1}^j(x,p)$ and taking as given the vacancy filling rates $\eta_t^j(x,p)$, the firm optimally chooses the number of vacancies $v_t(x,p)$ to post in period t. Optimal vacancies

$$\begin{aligned}\kappa'(v_t(x,p),x,p) &= \beta^f \sum_j \eta_t^j(x,p) W_{t+1}^j(x,p) \\ &= \beta^f \Pi\big(\theta_t\big) \frac{S_t(x,p)}{S_t} \sum_j \frac{S_t^j(x,p)}{S_t(x,p)} W_{t+1}^j(x,p) \qquad \text{(C.1)}\end{aligned}$$

The firm's optimal policy thus depends on the workers' search policies and distribution over labor market states via $S_t(x,p)$, S_t and $S_t^j(x,p)$. It also depends on θ_t which is a function of the other firms' policies and S_t.

Worker Problem Workers take as given next period's value functions V_{t+1} – and hence the expected surplus of finding a job – as well as the job filling rate $\Lambda(\theta_t)$ and vacancy shares $N_t(x,p)/N_t$ and choose their optimal search effort according to the resulting first order condition.

$$\frac{\mathrm{d}d^j(\ell)}{\mathrm{d}\ell} = \beta\phi_s\Lambda(\theta_t)\left(\sum_{x,p}\frac{N_t(x,p)}{N_t}\max\left\{V^j_{e,t+1}(x,p), V^j_{s,t+1}(x,p)\right\} - V^j_{s,t+1}(x,p)\right) \tag{C.2}$$

The workers' optimal policies thus depends on the firms' vacancy policies and via $N_t(x,p)$ and N_t. It also depends on θ_t which is a function of the other workers' policies and N_t.

Algorithm Focus on one skill segment h and let F_t be the distribution of workers across labor market states in period $t = 0,\dots,T$. The economy is in the initial regime until period $t = -1$. We thus set F_0 equal to the stationary distribution in the initial regime. We assume that the economy has converged to the new regime by period T. All equilibrium objects in period T are thus the equilibrium objects in the stationary equilibrium. The main backward looking object is F_t. Search mass, vacancy mass and tightness can be adjusted instantly and are thus allowed to jump from $t = 0$ to $t = 1$. The distribution F_t only jumps due to non-employability.

Knowing the initial and terminal stationary equilibrium, we proceed as follows.

1. Guess a sequence $\{F_t^0\}_t$, e.g. a piece-wise linear interpolation between F_T and F_0 taking into account the employability constraint.
2. Set $i = 0$
3. Taking as given the sequence of distributions $\{F_t^i\}_t$ as well as the value functions W_T^j and $V^j_{s,T}$, solve backwards for the equilibrium sequence of policies $\{\ell_t^i, v_t^i\}_t$. Starting with $t = T-1$, solve for the equilibrium policies in t as follows:
   (a) guess vacancy shares and tightness: $N_t(x,p)$ and θ_t
   (b) solve for optimal search policies $\ell_t^i(j,x,p)$ using equation (C.2)
   (c) update $S_t(x,p)$, S_t, $S_t^j(x,p)$ and θ_t
   (d) solve for optimal vacancy policies $v_t^i(x,p)$ using equation (C.1)
   (e) compute implied vacancy shares and tightness
   (f) if equal to guess, stop, else update guess and go back to (b)
   (g) compute the workers' value: $V_t^j(\sigma) = u^j(\sigma,\ell^i) + \beta\mathbb{E}_{\sigma'|\sigma}\left[V^j_{t+1}(\sigma')|\sigma\right]$
   (h) compute the firm's values: $W_t^j(x,p) = y(x,p) + \beta(1-\delta_t^j(x,p))W^j_{t+1}(x,p)$
4. Set $t = t-1$ and and repeat until $t = 0$

5. Use the transition matrices P_t^i to iterate forward on the distribution starting from F_0 until F_T to get $\{F_t^{i+1}\}_t$
6. Check whether the implied sequence $\{F_t^{i+1}\}_t$ differs from the guess $\{F_t^i\}_t$. Stop if yes. Set $i = i + 1$ and go back to step (3)

Chapter 2

Falling Behind: Has Rising Inequality Fueled the American Debt Boom?

Joint with Fabian Greimel.

1 Introduction

Between 1980 and 2007, US household debt doubled relative to GDP. Mortgage debt was by far the most important driver of this household debt boom (see Figure 1a). In lockstep with mortgages, income inequality started to rise in 1980 and reached its peak in 2007 (see Figures 1b and 2). While real incomes stagnated for the bottom half of the population, incomes of the top 10% more than doubled over this time period (see Figure 2). The rise in household debt has drawn a lot of interest that has mostly focused on the role of falling interest rates following an increase in foreign and domestic *supply of credit* (e.g. Bernanke, 2005; Mian et al., 2021). In this paper, we investigate whether rising income inequality and *Keeping up with the richer Joneses* fueled the mortgage boom through an increase in the *demand for housing*.

We begin by documenting novel aspects of the US mortgage and housing boom which call for such a demand-side mechanism to complement supply-side drivers of household debt. First, we document that the aggregate link between top incomes and household debt extends to the level of US states. That is, household debt grew substantially more in US states where top incomes grew faster. While falling interest rates clearly fueled the rise in household debt, they are not able to account for this state-level link between top incomes and mortgage debt because financial markets are largely integrated and local demand for savings need not be compensated by local debt. Instead, arbitrage should lead to a uniform increase in debt of the non-rich across all states. Second, we show that higher state-level top incomes drive up mortgage debt but do not affect state-level non-mortgage debt which suggests that housing plays a key role in the transmission from rising top incomes to rising household debt. Third, we find that house prices grew faster in states with a stronger increase in top incomes pointing to an important role for housing demand. Most importantly, we show that an increase in the average incomes of the rich is associated with an increase in the mortgage-to-income ratio of the non-rich—a key prediction of of the *Keeping up with the richer Joneses* (KURJ) mechanism.

In the main part of the paper, we assess—analytically and quantitatively—the aggregate consequences of rising income inequality in the presence of social comparisons. The idea that people compare themselves to others was introduced into economics by Veblen (1899) and Duesenberry (1949) and is supported by plenty of recent empirical research.[1] We incorporate social comparisons into a heterogeneous agent model of the macro economy. In our model, households not only care about their own consumption and housing, but also about how their house compares to the benchmark set by the rich. When top incomes rise and the rich upgrade their houses, the non-rich lose some of their social status and substitute status-enhancing housing for status-neutral consumption to keep up with the richer Joneses. These houses are mortgage-financed, causing a boom in debt-to-income ratios across the entire income distribution as well as an increase in house prices.

In a stylized version without idiosyncratic income risk, we can show analytically how this status externality affects aggregate debt depending on who cares about whom in the network of social comparisons. Allowing for arbitrary connections among income types, we prove that a household's housing demand and her debt level is increasing in the incomes of her reference group (or the reference groups of her reference group). This is because these incomes determine the reference measure of housing. In the empirically relevant case where non-rich households care about the rich, the debt-to-income ratio of the non-rich is increasing in the incomes of the rich.[2] This implies that aggregate debt-to-income is increasing in top incomes.

We then calibrate the full model with idiosyncratic income risk in order to quantify the contribution of this mechanism to the observed increase in mortgages and house prices between 1980 and 2007. We discipline the social comparison motive using independent micro evidence on housing comparisons in the US by Bellet (2019) who estimates how strongly households' utility is affected by the housing of the local rich. The main experiment is to compare two steady states that differ only in the exogenous degree of income inequality. In

[1]See for example Luttmer (2005), Charles et al. (2009), Kuhn et al. (2011), Bursztyn et al. (2014), or De Giorgi et al. (2020). Importantly, there is recent micro evidence by Bellet (2019) showing that (i) comparisons are upward-looking such that only the houses of the rich positional externality.

[2]Bellet (2019) shows that households only care about the top end of the housing distribution. To our knowledge all papers that have tested for asymmetries in the comparison motive have found them to be upward-looking (e.g. Clark and Senik, 2010; Ferrer-i-Carbonell, 2005; Card et al., 2012).

particular, we scale the *permanent* component of income inequality to match the increase in cross-sectional income dispersion between 1980 and 2007.[3]

We find that in the presence of KURJ, the rise in income inequality can explain up to half of the observed 120%-increase in the mortgage-to-income ratio and up two thirds of the observed 60%-increase in house prices between 1980 and 2007. This effect can be decomposed into a direct effect and an indirect effect. On the one hand, social comparisons directly raise housing demand and thereby mortgage demand for non-rich households. On the other hand, rising inequality drives up house prices through growing demand for housing at the top of the income distribution. As housing and non-durable consumption complement each other, this increase in the equilibrium house price pushes up housing and mortgage demand of non-rich households. Even in the absence of KURJ, rising inequality can explain about a quarter of the observed mortgage boom through this general equilibrium effect. The model also accounts for up to half of the observed 65%-increase in the house-value-to-income ratio between 1980 and 2007.

Finally, we compare the effects of our demand-side mechanism to those of the *Global Saving Glut*, i.e. the surge in the foreign net debt position of the US from about 0% of GDP in 1980 to about 40% of GDP in 2007 (Bernanke, 2005; Justiniano et al., 2014).[4] In our model, this increase in the supply of credit can account for about 30% of the debt boom through lowering the real interest rate by approximately 40%. In contrast to rising inequality and KURJ, however, we find that the Global Saving Glut increases house prices by only 2% and the ratio of house values to income by only 4%. Both mechanisms together can explain up to three quarters of the observed 120%-increase in the mortgage-to-income ratio. Decomposing this total effect, we can attribute between one half and two thirds of the explained increase in debt to rising inequality and KURJ.

Related Literature. Our paper relates to several strands of the literature. First, we contribute to the literature that studies the drivers of the US household debt boom which was documented by Jordà et al. (2016) and Kuhn et al. (2017). A range of papers focuses on an increase in the foreign or domestic supply of credit that drives up household debt through a drop in the interest rate (Justiniano et al., 2014; Kumhof et al., 2015; Mian et al., 2020, 2021). Most notably, Mian et al. (2021) show that differences in saving rates out of permanent income can link rising income inequality to rising credit supply and falling interest rates. Other papers study the role of looser collateral constraints (e.g. Favilukis et al., 2017) and lending limits (Justiniano et al., 2019) as well as changes in house price expectations (Adam et al., 2012; Kaplan et al., 2020). This paper adds to this literature

[3]This focus on permanent income inequality is in line with evidence in Kopczuk et al. (2010) and Guvenen et al. (2018) who show that the rise in cross-sectional inequality is mostly driven by permanent income differences rather than increasing risk.

[4]Gourinchas et al. (2017) show that the net foreign debt position can be well approximated by the cumulative current account deficit.

Figure 1: The American Household Debt Boom and Rising Income Inequality

(a) Decomposition of the Debt Boom | (b) Mortgages (left axis) and Inequality (right)

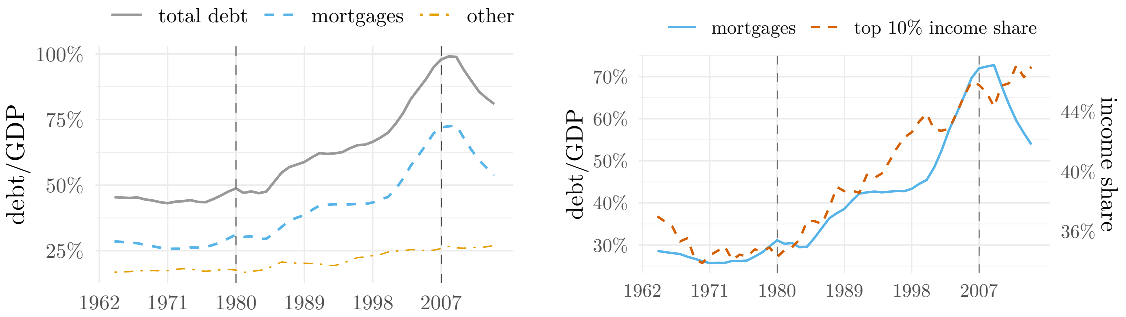

Notes: This figure shows the relationship of aggregate household debt (total, mortgage, non-mortgage) as share of GDP and the top 10% income share over time. Data sources: US Flow of Funds and Alvaredo et al. (2016). Details see Appendix B.

Figure 2: Distribution of Income Growth

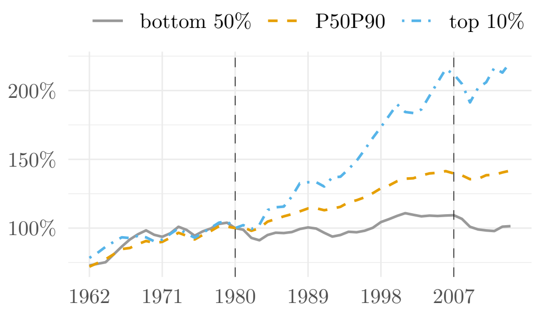

Notes: This figure shows real average pre-tax income growth from 1962 to 2014 in the US. Data are taken from Piketty et al. (2018). Growth rates are relative to the base year 1980.

Figure 3: House Prices in the US

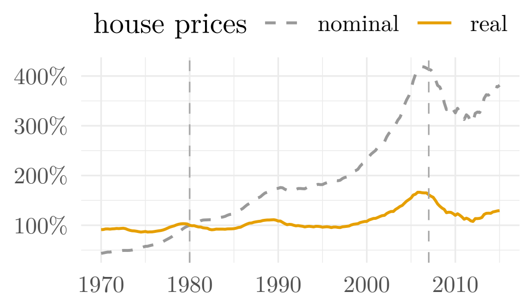

Notes: Nominal: Case-Shiller Home Price Index for the USA. Real: Deflated by the Consumer Price Index. Base year: 1980. Source: `http://www.econ.yale.edu/~shiller/data.htm`

in two ways: First, we document a tight link between (non-rich) mortgage debt and top incomes on the state level. Second, we explore a novel demand-side mechanism that can help rationalize the tight link between top incomes and non-rich debt and complements the existing supply-side mechanisms.[5]

Second, we contribute to the growing literature on the aggregate effects of rising income inequality. Heathcote et al. (2010) analyze how rising wage inequality affects human capital investment and labor supply. Auclert and Rognlie (2018) investigate how permanent and transitory income inequality differentially impact aggregate demand. Straub (2018) shows that rising permanent income inequality drives down interest rates in the presence of non-homothetic preferences. Fogli and Guerrieri (2019) analyze the interplay between residential segregation and income inequality in the presence of local spillovers that af-

[5] Note that the question whether rising top income inequality fueled the boom in household debt and amplified the Great Recession was also discussed in the public debate (e.g. Rajan, 2010; Stiglitz, 2009; Frank, 2013). See also the survey by van Treeck (2014).

fect the education returns. The approach of Straub (2018) and in particular Fogli and Guerrieri (2019) is similar to ours in the sense that we integrate insights from empirical microeconomic research into (non-homothetic preferences, local spillovers, social comparisons) into a macroeconomic heterogeneous agent model to analyze potential interactions with rising inequality. In our model, agents are linked not only through prices but also directly through social externalities of their consumption decisions.

Third, we contribute to the large literature on social comparisons (e.g. Luttmer, 2005; Card et al., 2012; Perez-Truglia, 2019) and economic choices (Charles et al., 2009; Kuhn et al., 2011; Bursztyn et al., 2014; Bertrand and Morse, 2016; Bursztyn et al., 2017; Bellet, 2019; De Giorgi et al., 2020). While the macroeconomic effects of keeping up with the Joneses have already been studied in the context of representative agent models (e.g. Abel, 1990; Campbell and Cochrane, 1999; Ljungqvist and Uhlig, 2000), we introduce social comparisons into a quantitative heterogeneous agents model. We build on the macroeconomic literature on keeping up with the Joneses and bring it closer to the empirical evidence. First, we distinguish between conspicuous and non-conspicuous goods. In our model households compare themselves only in their houses, arguable the most important conspicuous good (e.g. Solnick and Hemenway, 2005; Bertrand and Morse, 2016). And second, agents compare themselves to the rich (e.g. Card et al., 2012; Bellet, 2019). Households only lose satisfaction with their own house, when a big house is built.[6]

Fourth, our empirical results are most closely related to the the studies by Bertrand and Morse (2016) who use CEX data and state-year variation to document that consumption expenditures of non-rich households respond to the incomes and consumption expenditures of the rich. Coibion et al. (2020) investigate the relationship between zip-code level income inequality (P90-P10 ratio) and household debt between 2000 and 2012 and find heterogeneous effects by income rank. Mian et al. (2020) analyze whether increasing top incomes in a state lead to an increase in the amount of non-rich household debt held as an asset by the state's rich. We analyze whether the state's non-rich take on more debt and analyze the dynamic effects of increases in top incomes. In addition, we show that growing top incomes are also associated with higher state-level house prices.

Finally, our analytical results extend those by Ghiglino and Goyal (2010) and Ballester et al. (2006) who show that agents' choices depend on the strengths of social links in a one-period model. We extend their network models to infinite horizon and add a durable good (housing) to show that debt is increasing in the centrality of an agent. The centrality is reinterpreted as the weighted sum of incomes of the comparison group.

Structure of the paper The rest of the paper is structured as follows: In Section 2 we present empirical evidence on the relationship between household debt and top income

[6]These novel modeling choices distinguish our paper from Badarinza (2019), who shows that status externalities lead to inefficient debt levels in a lifecycle model.

inequality. In Section 3 we describe our model. In Section 4 we derive analytically how top incomes drive debt in a stylized version of the model. In Section 5 we describe the parameterization of the full model, followed by quantitative results in Section 6. Section 5 concludes.

2 Empirical Analysis: Top Incomes and Household Debt

In this Section, we use state-level distributional national accounts (DINA) data (Piketty et al., 2018; Mian et al., 2020) to study the relationship between US top incomes and household debt in more detail. We show that the aggregate link between top incomes and (non-rich) household debt goes beyond mere coincidence. Our empirical analysis exploits state-year variation in top incomes after controlling for aggregate shocks and time-invariant state heterogeneity. Let us emphasize at the outset that we do not use an explicit source of quasi-experimental variation in top incomes. Instead, we follow Mian et al. (2020) and argue that plenty of evidence in the literature supports the view that the rise in top inequality was triggered by shifts in technology and globalization that took place at the outset of the rise in inequality around 1980 (e.g. Katz and Murphy, 1992; Autor et al., 2008; Smith et al., 2019).

2.1 Data & Approach

We use state-level data on incomes and debt between 1980 and 2007 adapted from the data provided by Mian et al. (2020). These data are based on DINA data from Piketty et al. (2018). As state-level identifiers in the DINA data are suppressed for incomes above 200,000 US dollars, state identifiers are imputed using state-level data from the Internal Revenue Service (IRS) which include information on how many tax returns above 200,000 dollars come from each state.[7] Our main data set is a state-year panel for the period 1980–2007 covering income, outstanding mortgages, outstanding non-mortgage debt and outstanding total debt for different income groups such as the rich (top 10%) and the non-rich (bottom 90%). We complement these data with state-level data on house prices from the Federal Housing Finance Agency and consumer prices from Hazell et al. (2020).

[7]The imputation is based on the assumption that incomes above 200 thousand dollars follow a state-specific Pareto distribution with density $f_s(y) = \frac{\alpha_s 200{,}000^{\alpha_s}}{y^{\alpha_s+1}}$ where α_s can be computed from the state-level mean income of units with gross income above 200,000 dollars. The ratio of the state-specific and aggregate income density gives the relative likelihood that an observation comes from that state. This is then used to weight all observations when computing state averages.

Table 1: Top Incomes and Household Debt: Fixed Effects Regressions

	log($\text{debt}_{s,t}$)					
	(A) total		(B) mortgage		(C) non-mortgage	
	population	non-rich	population	non-rich	population	non-rich
	(1)	(2)	(3)	(4)	(5)	(6)
log(top $\text{incomes}_{s,t-2}$)	0.108*	0.215***	0.180**	0.306***	-0.156***	-0.054
	(0.059)	(0.070)	(0.079)	(0.099)	(0.050)	(0.054)
log(own $\text{income}_{s,t}$)	0.678***	0.796***	0.764***	0.934***	0.463***	0.519***
	(0.093)	(0.060)	(0.119)	(0.087)	(0.051)	(0.057)
State FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Estimator	OLS	OLS	OLS	OLS	OLS	OLS
N	1,224	1,224	1,224	1,224	1,224	1,224
R^2	0.983	0.979	0.974	0.965	0.985	0.986

Notes: This table shows the estimation results corresponding to equation 1. The dependent variable is either total, mortgage or non-mortgage debt in the population or among the non-rich. Robust standard errors in parentheses. The stars indicate the range of the p value: $^{***} \leq 0.01 \leq {}^{**} \leq 0.05 \leq {}^{*} \leq 0.1$.

2.2 Top Incomes and Household Debt: Fixed Effect Regressions

The main explanatory variable is the log of lagged top incomes measured as the average income in the top 10%.[8] Let $\text{debt}^g_{s,t}$ be either total, mortgage or non-mortgage debt in sub-population g in state s at time t. In our main estimation equation, we regress debt of group g on lagged top incomes, income of group g as well as state and year fixed effects.

$$\log(\text{debt}^g_{s,t}) = \beta \log(\text{top incomes}_{s,t-2}) + \gamma \log(\text{incomes}^g_{s,t}) + \delta_s + \delta_t + \epsilon_{s,t} \qquad (1)$$

If β is positive, higher top income levels are associated with higher levels of future (non-rich) debt when (non-rich) incomes are held constant and state and year effects (δ_s, δ_t) are controlled for. Table 1 reports the results. Columns (1) and (2) show that an increase in lagged top incomes has a statistically significant positive effect on total household debt—for both the population and non-rich households. Holding non-rich incomes constant, an increase in lagged top incomes by 1% is associated with an increase in non-rich debt by over 0.2%. Columns (3) and (4) show the results for mortgage debt. The effect on mortgage debt is even stronger. In contrast, columns (5) and (6) show that top incomes do not have a positive effect on non-mortgage debt. If anything, the relationship is negative.

Figure 4 visualizes the regression results using binned scatter plots of residualized log debt against residualized lagged log top incomes. The slope of the fit is equal to β in equation (1). The regression model is able to capture a substantial amount of the

[8] We use lagged top incomes for two reasons. First, building houses takes time. Second, if non-rich households *keep up with the richer Joneses*, they will only react once they see the houses of the rich. We use the second lag of top incomes, but results are robust to using lags greater than two.

FIGURE 4: Residualized Household Debt and Lagged Top Incomes

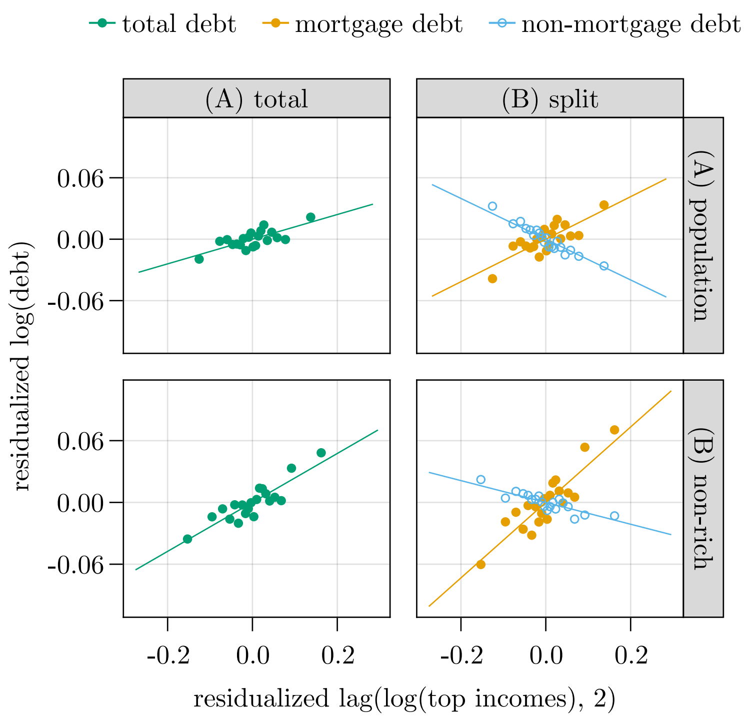

Notes: This figure shows the relationship between log debt (total, mortgage, non-mortgage) of all households and non-rich households and the second lag of log average top 10% incomes conditional on state and year fixed effects and non-rich income. All variables are residualized using state and year fixed effects. The slope of the regression line is the OLS estimate of β reported in Table 1. The figure shows averages in 20 equally sized bins of the x-variable.

relationship of household debt and top incomes after accounting for fixed effects and own income.

Figure 5 shows the raw relationship between the long-run changes (1982 to 2007) in top incomes and household debt-to-income ratios. While there seems to be no significant relationship between top incomes and the debt-to-income ratio in the total population, there is a significant positive relationship between the change in top incomes and the change in the debt-to-income ratio of *non-rich* households.[9] Note also that the slope in panel A is pulled towards zero by four states with exceptionally high income growth: New York, Massachussetts, Connecticut and the District of Columbia. Figure A.3 in Appendix A shows the slopes with and without these four states. Importantly, the picture is highly consistent with the above fixed effect regression results when looking at the relationship between top incomes and *non-rich* households' debt-to-income ratio irrespective of whether

[9]Table A.2 in the appendix reports the results of the bivariate regressions depicted in Figure 5.

FIGURE 5: Long-Run Changes in Household Debt and Top Incomes

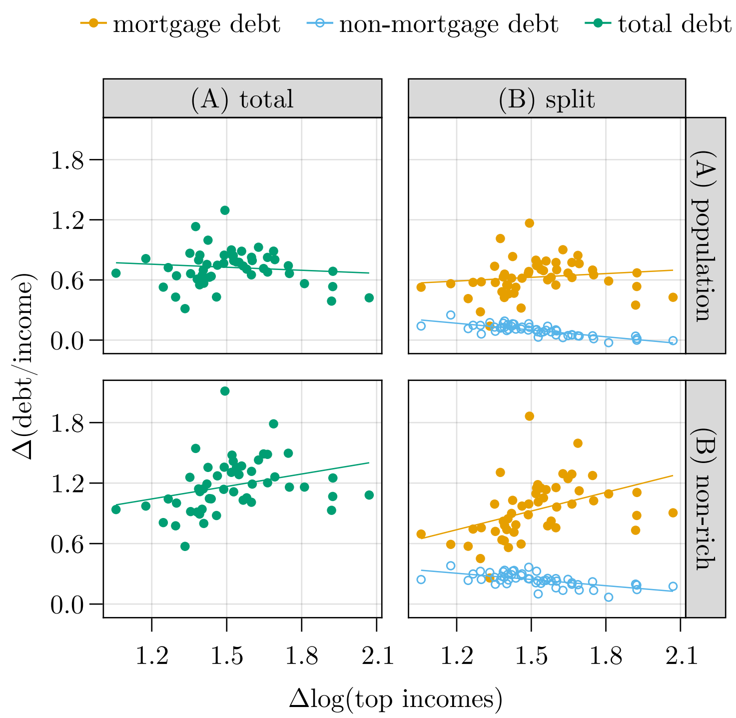

Notes: This figure shows the relationship between the change in debt-to-income (total, mortgage, non-mortgage) of all households (population) and non-rich households (bottom 90% of the income distribution) and the change in the log of average top 10% incomes between 1982 and 2007 across US states.

we include the four states with exceptionally high top income growth. Recall that this is a key prediction of the KURJ mechanism. In addition, Figure A.2 in the appendix shows that this positive relationship between long-run changes in top incomes and non-rich debt also holds for the middle 40% (P50 to P90) and bottom 50% of the income distribution.

2.3 Top Incomes and Household Debt: Dynamic Effects

To complement the two-way fixed effect regressions, we now analyze the dynamic response of household debt to changes in top incomes. In particular, we estimate how household debt changes from time $t-1$ to $t+h$ in response to a change in top incomes from $t-1$ to t using local projections of the form

$$\Delta^{h+1}\log(\text{debt}_{s,t+h}) = \alpha^h + \beta^h \Delta \log(\text{top incomes}_{st}) + \delta_t^h + \sum_{k=1}^{3}\Big(\gamma_k^h \log(\text{debt}^g_{s,t-k}) + \phi_k^h \log(\text{top incomes}_{s,t-k})\Big) + \epsilon_{st}^h \tag{2}$$

for each $h \in \{0, \ldots, 10\}$, where

$$\Delta^{h+1}\log(\text{debt}_{s,t+h}) \equiv \log(\text{debt}^g_{s,t+h}) - \log(\text{debt}^g_{s,t-1})$$
$$\Delta \log(\text{top incomes}_{st}) \equiv \log(\text{top incomes}_{s,t}) - \log(\text{top incomes}_{s,t-1})$$

The coefficients β^h give us the cumulative %-change in non-rich debt that is induced by a one-time change in top-incomes by 1%. By adding past debt and inequality measures as controls, specification (2) essentially compares states with the same pretrends in debt and top incomes, but where one state experiences a stronger increase in top incomes from $t-1$ to t.

Figure 6 plots the estimated impulse response function for total, mortgage and non-mortgage debt both in the population and among the non-rich. Consistent with the previous fixed effect regressions, top incomes substantially drive up mortgage debt over the following ten years while non-mortgage debt remains roughly constant.[10] For the non-rich, a 10% increase in top incomes from $t-1$ to t translates into a persistent increase in mortgage debt of roughly 10% after ten years. For non-mortgage debt, there are no (persistent) effects.

Overall, these results not only show that rising top incomes are associated with rising debt-to-income ratios in the population and in particular among the non-rich, they also show that housing plays a central role in the transmission. While the aggregate (i.e. nation-wide) non-mortgage-debt-to-income ratio has gone up along with top incomes from 1980 to 2007, this link does not hold up on the state-level. This asymmetry between mortgage and non-mortgage debt is consistent with social comparisons in housing given that housing comparisons have at least some spatial bias. Even in the presence of modern communication technology, the local rich are arguable more visible and thus impose a greater status externality on households in the same state compared to other households across the country.

2.4 Top Incomes and House Prices

Having documented that state-level top incomes are associated with a subsequent increase in (non-rich) household mortgage debt, we now ask how state-level house prices react to

[10] Controlling for lags of non-rich income or total income does not change the results.

FIGURE 6: Dynamic Effects of Top Incomes on Household Debt

Notes: This figure shows the cumulative effect of a 1% change in top 10% incomes on total, mortgage and non-mortgage debt in the population and among the non-rich estimated from equation 2. The confidence bands are constructed using a significance level of 5%.

rising top incomes. We estimate the following fixed effect regression

$$\log(\text{HPI}_{st}) = \alpha + \beta \log(\text{top incomes}_{s,t-1}) + \gamma \log(\text{incomes}_{st}) + \delta_s + \delta_t + \epsilon_{st} \qquad (3)$$

for two measures of state-level house prices. We deflate state-level nominal house prices from the Federal Housing Finance Agency by consumer prices. In Columns (1) and (2), we use novel state-level consumer price data from Hazell et al. (2020) to construct real state-level house prices. Unfortunately, these data are only available for a shorter sub-sample.[11] That is why we also report house prices deflated by nation-wide consumer prices in Columns (3) and (4).[12] The results are shown in Table 2. For both house price measures, we find that lagged top incomes have a statistically significant effect on house prices when total or non-rich income is held fixed.

[11]The state-level CPI data are only available starting in 1979 for only 21 states, starting in 1988 for another 13 and not available at all for the remaining states. Nominal state-level house price data are available for all states starting in 1979.

[12]The use of year fixed effects takes out the aggregate trend in consumer prices.

Table 2: Top Incomes and House Prices: Fixed Effects Regressions

	log(real house price$_{s,t}$)			
	State FE-level CPI		country-level CPI	
	(1)	(2)	(3)	(4)
log(top incomes$_{s,t-2}$)	0.588***	0.440***	0.686***	0.570***
	(0.112)	(0.108)	(0.116)	(0.117)
log(bottom incomes$_{s,t}$)	0.579***		0.516***	
	(0.132)		(0.108)	
log(total incomes$_{s,t}$)		0.677***		0.528***
		(0.097)		(0.079)
State FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Estimator	OLS	OLS	OLS	OLS
N	793	793	1,428	1,428
R^2	0.925	0.924	0.947	0.945

Notes: This table shows the results of regression model in equation (3). Robust standard errors in parentheses. The stars indicate the range of the p value: $^{***} \leq 0.01 \leq ^{**} \leq 0.05 \leq ^{*} \leq 0.1$. Data: DINA, IRS.

Figure 7 shows the dynamic effect of top incomes on house prices estimated using the following estimation equation:

$$\Delta^{h+1} \log(\text{HPI}_{s,t+h}) = \alpha^h + \beta^h \Delta \log(\text{top incomes}_{st}) + \delta_t^h + \sum_{k=1}^{3} \gamma_k^h \log(\text{HPI}^g_{s,t-k}) + \sum_{k=1}^{3} \phi_k^h \log(\text{top incomes}_{s,t-k}) + \sum_{k=1}^{3} \psi_k^h \log(\text{total incomes}_{s,t-k}) + \epsilon_{st}^h \quad (4)$$

We find a statistically significant hump-shaped response of house prices to an increase in top incomes. Without state-level CPI data but for the full sample (left sub-figure), we find that house prices are up by 1% after ten years following a 1% increase in top incomes. When using state-level CPI data for a subset of states and years, the effect is slightly lower, but the overall pattern is very similar.

The result that growing top incomes trigger not only an increase in mortgage debt but also an increase in house prices points to a role for housing-demand effects to complement credit-supply effects in order to understand the boom in household debt.

FIGURE 7: Dynamic Effects of Top Incomes on House Prices

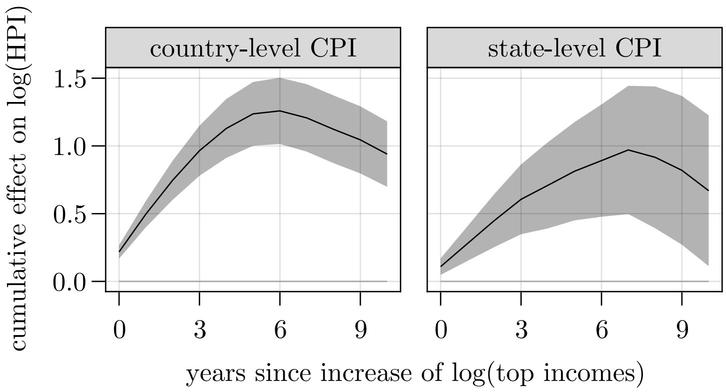

Notes: This figure shows the cumulative effect of a 1% change in top 10% incomes on house prices estimated from equation 4. The confidence bands are constructed using a significance level of 5%.

3 Model

Motivated by these empirical findings that call for a housing-demand channel, we now evaluate whether the combination of rising inequality and *Keeping up with the richer Joneses* can help rationalize the US mortgage and housing boom. To that end, we incorporate social comparisons into an otherwise standard macroeconomic model of housing. This section describes our model and sections 4 and 6 present our analytical and quantitative results.

Our model is a dynamic, incomplete markets general equilibrium model similar to the "canonical macroeconomic model with housing" in Piazzesi and Schneider (2016). We formulate our model in continuous time to take advantage of the fast solution methods of Achdou et al. (2021, in particular Section 4.3). We build our model with two aims in mind. First, we want to illustrate how rising top-incomes and social comparisons can lead to rising debt levels across the whole income distribution. And second, we want to quantify the effect of this channel on the increase in aggregate mortgage debt and house prices from 1980 to 2007.

3.1 Setup

Time is continuous and runs forever. There is a continuum of households that differ in their realizations of the earnings process. Households are indexed by their current portfolio holdings (a_t, h_t), where a_t denotes financial wealth and h_t denotes the housing stock, and their pre-tax earnings y_t. They supply labor inelastically to the non-durable consumption good and housing construction sectors. The financial intermediary collects households' savings and extends mortgages subject to a collateral constraint. The state of the economy is the joint distribution $\mu_t(\mathrm{d}a, \mathrm{d}h, \mathrm{d}y)$. There is no aggregate uncertainty.

3.2 Households

Households die at an exogenous mortality rate $m > 0$. The wealth of the deceased is redistributed to surviving individuals in proportion to their asset holdings (perfect annuity markets). Dead households are replaced by newborn households with zero initial wealth and earnings drawn from its ergodic distribution.[13] Households derive utility from a non-durable consumption good c and housing status s. They supply labor inelastically and receive earnings y. After-tax disposible earnings are given by

$$\tilde{y}_t = y_t - T(y_t),$$

where T is the tax function. Households choose streams of consumption $c_t > 0$, housing $h_t > 0$ and assets $a_t \in \mathbb{R}$ to maximize their expected discounted lifetime utility

$$\mathbb{E}_0 \int_0^\infty e^{-(\rho+m)t} \frac{\Big((1-\xi)c_t^\varepsilon + \xi s(h_t, \bar{h}_t)^\varepsilon\Big)^{\frac{1-\gamma}{\varepsilon}}}{1-\gamma} \mathrm{d}t, \tag{5}$$

where $\rho \geq 0$ is the discount rate and the expectation is taken over realizations of idiosyncratic earnings shocks. $1/\gamma > 0$ is the inter-temporal elasticity of substitution, $1/(1-\varepsilon) > 0$ is the intra-temporal elasticity of substitution between consumption and housing status and $\xi \in (0,1)$ is the relative utility-weight for housing status.

A household's utility from housing is a function of the *housing status* $s(h, \bar{h})$. Housing status increases in the household's housing stock h and decreases in reference housing $\bar{h}$ which is a function of the equilibrium distribution of housing as introduced in the next section.

Housing is both a consumption good and an asset. It is modeled as a homogenous, divisible good. As such, h represents a one-dimensional measure of housing quality (including size, location and amenities). An agent's housing stock depreciates at rate δ and can be adjusted frictionlessly.[14] Home improvements and maintenance expenditures x_t have the same price as housing (p) and go into the value of the housing stock one for one.

Households can save ($a > 0$) and borrow ($a < 0$) at the equilibrium interest rate r. Borrowers must post their house as collateral to satisfy an exogenous collateral constraint. The collateral constraint pins down the maximum possible loan-to-value ratio ω.

[13]This follows Kaplan et al. (2018).

[14]Frictionless adjustment is justified, because we will be comparing long-run changes (over a period of 27 years).

Households' assets evolve according to

$$\dot{a}_t = \tilde{y}_t + r_t a_t - c_t - p_t x_t, \tag{6}$$

$$\dot{h}_t = -\delta h_t + x_t, \tag{7}$$

subject to the constraints

$$a_t \geq -\omega p_t h_t, \tag{8}$$

$$h_t > 0.$$

3.3 Social Comparisons

We build on the macroeconomic literature (e.g. Abel, 1990; Gali, 1994; Campbell and Cochrane, 1999; Ljungqvist and Uhlig, 2000) on keeping up with the Joneses and bring it closer to the empirical evidence. These papers feature representative agent models with one good and one asset. Agents compare themselves in the single consumption good, and their reference measure is the average consumption in the economy.[15]

We depart from this literature in two ways. First, we assume that households compare themselves only in their houses. This captures that people compare themselves only in conspicuous goods and that housing is one of the most important conspicuous goods—both in terms of visibility and expenditure share (e.g. Solnick and Hemenway, 2005; Bertrand and Morse, 2016).

Second, we allow the reference measure to be a function of the distribution of houses (and not necessarily its mean): $\bar{h}_i = \bar{h}_i(\mu_h)$. This reflects that the comparison motive is asymmetric, being strongest (and best documented) with respect to the rich (e.g. Clark and Senik, 2010; Ferrer-i-Carbonell, 2005; Card et al., 2012, on self-reported well-being). People buy bigger cars when their neighbors win in the lottery (Kuhn et al., 2011); non-rich move their expenditures to visible goods (such as housing) when top incomes rise in their state (Bertrand and Morse, 2016); and construction of very big houses leads to substantially lower levels of self-reported housing satisfaction for other residents in the same area—while the construction of small houses does not (Bellet, 2019).

For our analytical results we assume that $\bar{h}$ is a weighted mean of the housing distribution and use $s(h, \bar{h}) = h - \phi\bar{h}$ for tractability. For the quantitative results, we set $\bar{h}$ to the 90th percentile of the housing distribution and use $s(h, \bar{h}) = \frac{h}{\bar{h}^{\phi}}$ based on empirical evidence (see Section 5).

[15] In equilibrium the reference measure has to be equal to the optimal choice of the representative agent.

3.4 Pre-Tax Earnings Process

In our main experiment, we want to adjust life-time (permanent) income inequality independently of income risk to capture the way income inequality has changed over time. We follow Guvenen et al. (2021), who estimate a pre-tax earnings process on administrative earnings data. The process consists of individual fixed effects ($\tilde{\alpha}_i \sim \mathcal{N}(\mu_\alpha, \sigma_\alpha)$), a persistent jump-drift process (z_{it}), a transitory jump-drift process (ϵ_{it}), and heterogeneous non-employment shocks ($\nu_{it} \in \{0, 1\}$).[16] We translate their estimated process to continuous time. Heterogeneity in $\tilde{\alpha}_i$ represents fixed ex-ante differences in earnings ability which is an important source of life-time inequality. If employed, individual pre-tax earnings are given by

$$y_{it}^{\text{pot}} = \exp(\tilde{\alpha}^i + z_{it} + \epsilon_{it}).$$

We will refer to y^{pot} as *potential earnings.* The actual pre-tax earnings (taking into account unemployment) are

$$y_{it} = (1 - \nu_{it}) y_{it}^{\text{pot}}.$$

See Appendix D for more details on the earnings process.

3.5 Production

There are two competitive production sectors producing the non-durable consumption good c and new housing investment I_h, respectively. Following Kaplan et al. (2020), there is no productive capital in this economy.

Non-Durable Consumption Sector The final consumption good is produced using a linear production function

$$Y_c = N_c \tag{9}$$

where N_c are units of labor working in the consumption good sector. As total labor supply is normalized to one, N_c is also the share of total labor working in this sector. The equilibrium wage per unit of labor is pinned down at $w = 1$.[17]

Construction Sector We model the housing sector following Kaplan et al. (2020) and Favilukis et al. (2017). Developers produce housing investment I_h from labor $N_h = 1 - N_c$

[16] We use version (7), where we take out the deterministic life-cycle profile. The only component that this version does not have are differences in deterministic income growth rates.

[17] Neither labor supply nor the wage appear in the earnings process, because there is no aggregate risk, households inelastically supply one unit of labor, and the wage is equal to 1.

and buildable land $\bar{L}$, $I_h = (\Theta N_h)^\alpha (\bar{L})^{1-\alpha}$ with $\alpha \in (0,1)$. Each period, the government issues new permits equivalent to $\bar{L}$ units of land, and these are sold at a competitive market price to developers. A developer solves

$$\max_{N_h} p_t I_h - w N_h \quad \text{s.t. } I_h = N_h^\alpha \bar{L}^{1-\alpha}$$

In equilibrium, this yields the following expression for optimal housing investment

$$I_h = (\alpha p)^{\frac{\alpha}{1-\alpha}} \bar{L}$$

which implies a price elasticity of aggregate housing supply of $\frac{\alpha}{1-\alpha}$.

3.6 Financial Markets

The financial intermediary collects savings from households and issues mortgages to households. Lending is limited by the households' exogenous collateral constraint (8).

In addition, the intermediary has an exogenous net asset position with the rest of the world a_t^S. The equilibrium interest ensures that bank profits are zero and the asset market clears,

$$\int a_t(a,h,y) \mathrm{d}\mu_t = a_t^S. \tag{10}$$

3.7 Stationary Equilibrium

A stationary equilibrium is a joint distribution $\mu(a,h,y)$, policy functions $c(a,h,y,\bar{h})$, $x(a,h,y,\bar{h})$, $h(a,h,y,\bar{h})$, $a(a,h,y,\bar{h})$, prices (p,r) and a reference measure $\bar{h}$ satisfying the following conditions

- Policy functions are consistent with agents' optimal choices $(c_t, h_t, a_t)_{t>0}$ given incomes $(y_t)_{t>0}$, prices p, r and the reference measure $\bar{h}$.
- Housing investment is such that the construction sector maximizes profits.
- $\mu(a,h,y)$ is stationary. That is, if the economy starts at μ, it will stay there.
- Asset market clears (10) and housing investment equals housing production $\int x(a,h,y)d\mu = I_h$.
- The reference measure is consistent with choices: $\bar{h} = \bar{h}(\mu)$.

4 Analytical Results

In this section we use a stylized version of the model described in section 3 to illustrate analytically how rising top incomes can lead to rising mortgage levels across the whole income distribution via social comparisons. In Proposition 1 we provide formulas for optimal housing and consumption, as functions of their permanent incomes, and the permanent incomes of the direct and indirect reference groups. In Proposition 2 we show that optimal debt is increasing in the incomes of the direct and indirect reference groups. In Proposition 3 we show that the impact of rising incomes $\tilde{y}_i$ on aggregate debt is increasing in type i's *popularity*. In Corollary 1 we show that total debt-to-income is increasing in top incomes if at least one person compares themselves to the rich. In Corollary 2 we show that under Cobb-Douglas aggregation ($\varepsilon = 0$), these results hold even under housing market clearing because they are independent of house prices p. In Corollary 3 we show that these results crucially depend on the fact the status good h is durable.

The assumptions needed to obtain tractability are that there is no idiosyncratic income risk; that the social status function is linear; and that the interest rate equals the discount rate (all of these assumptions are relaxed in the following sections).

Assumption 1. $r = \rho$.

Further, we assume that there is a finite number of types of households $i \in \{1, \dots, N\}$. Agents vary by their initial endowments a_0 and flow disposable income $\tilde{y}$.

Assumption 2. Flow income $\tilde{y}_i$ is deterministic and constant over time, but varies across types i.

Without loss of generality, we assume that types are ordered by their permanent income $\mathcal{Y}_i = r a_0^i + \tilde{y}_i$,

$$\mathcal{Y}_1 \leq \mathcal{Y}_2 \leq \ldots \leq \mathcal{Y}_N.$$

We use bold variables to denote the vector variables for each type using the above ordering, e.g. $\boldsymbol{h} = (h_1, \dots, h_N)^T$.

Assumption 3 (Tractable social comparisons)**.** The status function $s(h, \bar{h}) = h - \phi\bar{h}$ is linear and the reference measure $\bar{h}_i = \sum_{j \neq i} g_{ij} h_j$ is a weighted sum of other agent's housing stock (we assume $g_{ij} \geq 0$).

Note, that we can write the vector of reference measures as $\bar{\boldsymbol{h}} = (\bar{h}_1, \dots, \bar{h}_N)^T = G \cdot \boldsymbol{h} := (g_{ij})(h_i)$. The matrix G can be interpreted as the adjacency matrix of the network of types capturing the comparison links between agents of each type. g_{ij} measures how strongly agent i cares about agent j.

We further require the comparisons to satisfy the following regularity condition.

Assumption 4. The Leontief inverse $(I - \phi G)^{-1}$ exists and is equal to $\sum_{i=0}^{\infty} \phi^i G^i$ for ϕ from Assumption 3.

This assumption is not very strong. This assumption is satisfied whenever the power of the matrix converges, $G^i \to G^\infty$. For example, if G represents a Markov chain with a stationary distribution or if G is nilpotent.

4.1 Characterization of the Partial Equilibrium

We solve for a simplified version of the equilibrium in Section 3.7. Agents solve their optimization problem given prices and the reference measure; the reference measure is consistent; but for now, we don't require market clearing. We use a lifetime budget constraint instead of the implicit transversality condition.

Households optimal decisions are given in the following proposition.

Proposition 1. *Under assumptions 1, 2, 3 and 4 the optimal choices* $\boldsymbol{h} = (h_1, \ldots, h_N)^T$ *and* $\boldsymbol{a} = (a_1, \ldots, a_N)^T$ *are given by*

$$\boldsymbol{h} = \Big(\sum_{i=0}^{\infty}(\kappa_1 \phi G)^i\Big)\kappa_2 \boldsymbol{\mathcal{Y}}.$$

$$-r\boldsymbol{a} = \tilde{\boldsymbol{y}} - \kappa_3 \boldsymbol{\mathcal{Y}} + (1-\kappa_3)\Big(\sum_{i=1}^{\infty}(\kappa_1 \phi G)^i\Big)\boldsymbol{\mathcal{Y}} \tag{11}$$

where $\kappa_1 = \frac{1}{\frac{p(r+\delta)}{\kappa_0}+1} \in (0,1)$, $\kappa_2 = \frac{\kappa_1}{\kappa_0}$, $\kappa_3 = \frac{1}{1+\frac{pr}{\delta p + \kappa_0}} \in (0,1)$ *and* $\kappa_0 = \Big((r+\delta)\frac{1-\xi}{\xi}p\Big)^{\frac{1}{1-\varepsilon}}$.

Proof. See appendix C.2. □

Households' choices depend on a weighted average of the permanent incomes of their (direct and indirect) reference groups. The weights are positive, whenever there is a direct or indirect social link between those agents. This is captured by the *income-weighted Bonacich centrality*, $B = \sum_{i=0}^{\infty}(C_1\phi G)^i \boldsymbol{\mathcal{Y}}$. If the weight B_{ij} is positive, household j's lifetime income affects household i's choices. This is the case whenever j is in i's reference group (there is a direct link $g_{ij} > 0$), or if j is in the reference group of some agent k who is in the reference group of agent i (there is an indirect link of length two, $g_{ik}g_{kj} > 0$) or if there is any other indirect link ($\prod_{n=1}^{N-1} g_{\ell_n,\ell_{n+1}}$ where $\ell_1 = i$ and $\ell_{N-1} = j$).

These results are reminiscent of those in Ballester et al. (2006). They showed that the unique Nash equilibrium in a large class of network games is proportional to the (standard) Bonacich centrality.

4.2 Comparative Statics

First, we show that optimal debt and optimal housing are increasing in incomes of the direct and indirect comparison groups.

Proposition 2. *For each type j in i's reference group (that is, $g_{ij} > 0$) and for each k that is in the reference group of the reference group (etc.) of i (that is, there is $j_1, j_2, \ldots, j_n$ such that $g_{ij_1}g_{j_1j_2}\cdots g_{j_{n-1}j_n}g_{j_nk} > 0$), then h_i is increasing and a_i is decreasing in $\mathcal{Y}_j$ (or $\mathcal{Y}_k$).*

Proof. G is non-negative, so $\sum_i c^i G^i$ is non-negative for all $c \geq 0$. From the definition of the Leontief inverse, being the discounted sum of direct and indirect links it follows,

$$\frac{\partial h_i}{\partial \tilde{y}_j} > \kappa_2 \kappa_1 \phi g_{ij} > 0 \quad \text{and} \quad \frac{\partial h_i}{\partial \tilde{y}_k} > \kappa_2 (\kappa_1 \phi)^{n-1} g_{ij_1} g_{j_1 j_2} \cdots g_{j_{n-1} j_n} g_{j_n k} > 0.$$

Similarly

$$-\frac{\partial a_i}{\partial \tilde{y}_j} > (1-\kappa_3)\kappa_1 \phi g_{ij} > 0 \quad \text{and} \quad -\frac{\partial a_i}{\partial \tilde{y}_k} > (1-\kappa_3)(\kappa_1 \phi)^{n-1} \phi g_{ij_1} g_{j_1 j_2} \cdots g_{j_{n-1} j_n} g_{j_n k} > 0.$$

□

Agent A's debt increases if agent B's lifetime income increases—as long as there is a direct or indirect link from A to B. That link exists, if agent A cares about agent B, or if agent A cares about some agent C who cares about agent B.

Second, we show how aggregage housing and debt react to changes in type j's income $\mathcal{Y}_j$. We first define the popularity of a type.

Definition 1 (Popularity)**.** We define the vector of popularities as

$$\boldsymbol{b}^T = \mathbf{1}^T \sum_{i=1}^{\infty} (\kappa_1 \phi G)^i,$$

and type i's popularity b_i as the ith component of $\boldsymbol{b}$.

The popularity is the sum of all paths that end at individual i. It measures how many agents compare themselves with i (directly and indirectly) and how strongly they do. The popularity of a type is crucial in determining how strongly their income will affect economic aggregates.

Proposition 3. *The impact of a change in type j's on aggregate housing and aggregate debt is proportional to its popularity.*

$$\frac{\partial}{\partial \tilde{y}_j} \sum_i h_i = \kappa_2 (1 + b_j)$$
$$\frac{\partial}{\partial \tilde{y}_j} \sum_i r a_i = (1-\kappa_3)(1 + b_j).$$

Proof. Take the expressions from proposition 1 and plug in the definitions for $\mathcal{Y}$ and b (Definition 1), aggregate housing can be written as $\sum_{i=1}^N h_i = \kappa_2 \sum_{i=1}^N (1+b_i)(\tilde{y}_i + r a_0^i)$ and aggregate debt can be written as $-\sum_{i=1}^N r a_i = (1-\kappa_3)\sum \tilde{y}_i - \kappa_3 \sum a_0^i + (1-\kappa_3)\sum_{i=1}^N b_i(\tilde{y}_i + r a_0^i)$. The derivatives follow immediately. □

Corollary 1. *If all types $i \neq j$ are connected to agent j and $\tilde{y}_j$ increases, then debt-to-income increases for all types $i \neq j$.*

Proof. By Proposition 2 debt of types $i \neq j$ increases, while their income is unchanged. It follows that debt-to-income rises. □

Corollary 2. *Under Cobb-Douglas aggregation, the results for a in Propositions 1, 2 and 3 are independent of house prices.*

Proof. Under Cobb-Douglas κ_0 is divisible by p. This means that p cancels in κ_1 and κ_3. Thus, all p cancel in the expression for a in Proposition 1 and consequently doesn't show up in the respective expressions in Propositions 2 and 3. □

The results on optimal debt in Propositions 2 and 3 and Corollary 1 break down if houses are non-durable. For any small time interval Δ, the depreciation rate has to be $\delta = \frac{1}{\Delta}$, so that the housing stock depreciates immediately,

$$(1 - \Delta\delta)h_t = 0.$$

Note the familiar special case when $\Delta = 1$ ("discrete time"), then the depreciation rate must be $\delta = 1$ for goods to be non-durable.

Corollary 3. *When houses are non-durable, optimal debt does not depend on others' incomes.*

Proof. In continuous time $\Delta \to 0$, so $\delta \to \infty$. It can be easily seen that $\kappa_3 \to 1$ as $\delta \to \infty$, thus $(1 - \kappa_3) \to 0$. Since all other terms in expression (11) are bounded, the part containing the Leontief inverse vanishes and becomes $-r\boldsymbol{a} = \tilde{\boldsymbol{y}} - \boldsymbol{\mathcal{Y}} = -r\boldsymbol{a}_0$. □

Note that this result does not depend on continuous time. The same result works in a discrete time version of the model, where $\Delta = \delta = 1$ and no limit argument is involved.

4.3 How Rising Top Incomes Fuel the Mortgage Boom: Intuition

It is at the heart of the mechanism that there is a complementarity between a household's housing stock and their reference measure. When top incomes $\mathcal{Y}_N$ rise, households of type N will improve (or upsize) their housing stock h_N, increasing the reference measure $\bar{h}_i$ for all types i that care about type N directly or indirectly. Each of these agents will optimally substitute durable, status-enhancing housing for non-durable status neutral consumption.

For debt to be affected it is key that the status good is durable and the status-neutral good is non-durable (see Corollary 3). Households want their stock of the durable good to be constant over time. Therefore, they need to pay for the entire house ph upfront and only replace the depreciation δph in the future. In other words, households need to shift some of their lifetime income forward to finance their house and take on mortgage debt to achieve that. The greater the value of the house, the bigger is the necessary mortgage.

4.4 Implications for Renters and Non-Mortage debt

We can use the limit case of non-durable housing to analyze the model implications for *non-mortgage debt.* If housing is non-durable, then housing services are essentially rented from real estate owners outside the model. As there is no house to finance, debt should now be interpreted as unsecured debt, smoothing out variations in earnings (a_0 vs y).

When top incomes rise and the rich scale up their (rental) housing, the other groups still substitute status enhancing housing services for the status-neutral consumption good, but there is *no effect on debt* (see Corollary 3).

4.5 Example: Upward Comparisons with Three Types of Agents

We now illustrate the results for the simple case of three types of agents, poor P, middle class M, and rich R. The poor type compares himself with both other types, the middle type compares himself only with the rich type, and the rich type not at all. Figure 8 shows the corresponding graph and its adjacency matrix.

$$\begin{tikzcd} P \arrow[r, "g_{PM}"] \arrow[rr, bend left, "g_{PR}"] & M \arrow[r, "g_{MR}"'] & R \end{tikzcd}$$

(A) The Graph

$$G = \begin{array}{c c} & \begin{array}{ccc} P & M & R \end{array} \\ \begin{array}{c} P \\ M \\ R \end{array} & \begin{pmatrix} 0 & g_{PM} & g_{PR} \\ 0 & 0 & g_{MR} \\ 0 & 0 & 0 \end{pmatrix} \end{array}$$

(B) The Adjacency Matrix

FIGURE 8: The social network structure with three types, assuming upward comparisons. The network can be represented as a graph and as its adjacency matrix.

Since G is a triangular matrix with only zeros on the diagonal, it is nilpotent ($G^3 = \mathbf{0}$), and thus the Leontief inverse exists.

$$G^2 = \begin{array}{c c} & \begin{array}{ccc} P & M & R \end{array} \\ \begin{array}{c} P \\ M \\ R \end{array} & \begin{pmatrix} 0 & 0 & g_{PM}g_{MR} \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \end{array}, \quad G^3 = \mathbf{0}$$

The matrix G^2 counts the paths of length 2. In our example there is only one such path—from type P to type R. Defining $\tilde{\phi} = \kappa_1 \phi$, the vector of Bonacich centralities is given by

$$\sum_{i=0}^{\infty} \alpha^i G^i = I + \sum_{i=1}^{2} \alpha^i G^i = I + \begin{pmatrix} 0 & \alpha \cdot g_{PM} & \alpha \cdot g_{PR} + \alpha^2 \cdot g_{PM} \cdot g_{MR} \\ 0 & 0 & \alpha \cdot g_{MR} \\ 0 & 0 & 0 \end{pmatrix}$$

The partial equilibrium choices for housing and debt are now given by

$$\begin{pmatrix} h_P \\ h_M \\ h_R \end{pmatrix} = \kappa_2 \begin{pmatrix} 1 & \tilde{\phi} \cdot g_{PM} & \tilde{\phi} \cdot g_{PR} + \tilde{\phi}^2 \cdot g_{PM} \cdot g_{MR} \\ 0 & 1 & \tilde{\phi} \cdot g_{MR} \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \mathcal{Y}_P \\ \mathcal{Y}_M \\ \mathcal{Y}_R \end{pmatrix}$$

$$-r \begin{pmatrix} a_P \\ a_M \\ a_R \end{pmatrix} = \tilde{\boldsymbol{y}} - \kappa_3 \boldsymbol{\mathcal{Y}} + (1-\kappa_3) \begin{pmatrix} 0 & \tilde{\phi} \cdot g_{PM} & \tilde{\phi} \cdot g_{PR} + \tilde{\phi}^2 \cdot g_{PM} \cdot g_{MR} \\ 0 & 0 & \tilde{\phi} \cdot g_{MR} \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} \mathcal{Y}_P \\ \mathcal{Y}_M \\ \mathcal{Y}_R \end{pmatrix}$$

An agent's housing choice increases linearly in own permanent income, $\mathcal{Y} = \tilde{y} + ra_0$, and on the permanent income of agents *in the reference group.* The poor agent's consumption increases through the direct links, but also indirect links (which are discounted more strongly). Agents' decisions to save or borrow depend on the ratio of initial wealth a_0 and income $\tilde{y}$. The higher the income relative to initial wealth, the greater the need to borrow.

5 Parameterization

Now we return to the full model. We parameterize the model to be consistent with the aggregate relationships of mortgage debt, house value and income in the US at the beginning of the 1980s. We use the estimated income process from Guvenen et al. (2021) and assign eight other parameters externally. The remaining two parameters (the discount rate ρ and the utility weight of housing status ξ) are calibrated internally so that in general equilibrium the aggregate net-worth-to-income ratio and aggregate loan-to-value ratio match these aggregate moments in the 1983 Survey of Consumer Finances.

Income Process We translate the estimated income process from Guvenen et al. (2021) to continuous time. It has a permanent, a persistent and a transitory component and state-dependent unemployment risk. Guvenen et al. (2021) estimate it to data from the time period 1994–2013. In order to construct the income process for the baseline economy $\mathcal{E}$ (corresponding to the year 1980) we rescale the permanent component following evidence on the changes in the income distribution from Kopczuk et al. (2010), Guvenen et al. (2014) and Guvenen et al. (2018).

The cross-sectional dispersion of incomes has increased substantially between 1980 and 2007. Figure 9 (taken from Guvenen et al., 2018, Figure 12) shows the variation of three common measures over time: the P90/P50 ratio, the P90/P10 ratio and the standard deviation of log-earnings. These changes in the variation of incomes can come from either component of the income process, or even a combination of them.

FIGURE 9: US Earnings Distribution

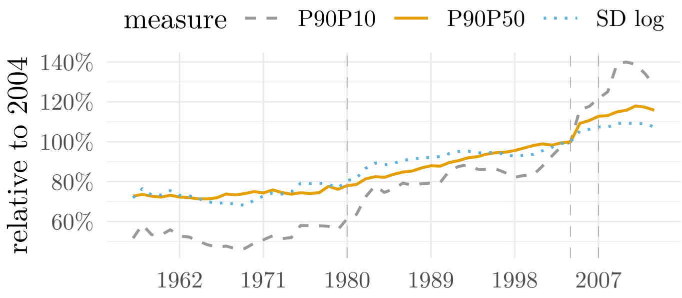

Notes: This figure shows the change in the cross-sectional distribution of male earnings in the US. Vertical bars in 1980, 2004 and 2007. Source: Guvenen et al. (2018).

While there is no consensus yet,[18] as to which of those factors contributed how much, there is evidence that rising permanent inequality explains a substantial share in increased cross-sectional variation. Kopczuk et al. (2010, Figure V) find that almost all of the change in earnings variation came from increases in permanent inequality. This finding is supported by Guvenen et al. (2014, Figure 5) who show that the variances of earnings shocks have had a slight downward trend since 1980.

Given this evidence, we attribute all change in inequality to changes in permanent inequality (σ_α). In our income process, permanent income inequality is represented by the permanent component $\tilde{\alpha}$. So, given the discretized version of the process, we stretch the upper half of the $\tilde{\alpha}$-grid to match the changes in the cross-sectional P90/P50 ratio.

When translating the process to continuous time, we assume that shocks arrive on average once a year (instead of every year). Moreover, we replace the discrete time *iid* process by a jump-drift process (ϵ_{it}) that is re-centered around zero whenever a shock hits so that shocks do not accumulate. The mean reversion rate of the persistent process (z_{it}) is the negative log of the discrete time persistence parameter which preserves the same annual autocorrelation. The exit rate out of nonemployment is chosen to match the average duration of nonemployment stays in the discrete time process. As households in our infinite horizon model die at a constant rate, we remove all age-dependence by setting the age profile constant (to the value at the mean age $\bar{t}$).[19] Table D.1 in the appendix shows all parameters of our continuous time earnings process.

We put the process on a discrete state space, using the approach of Kaplan et al. (2018). We discretize each component separately, obtaining continuous-time Markov chains[20] for the persistent and transitory components and combining them afterwards. Finally, we add the state-dependent non-employment risk.

[18] Carr and Wiemers (2016, 2018) show that depending on data source, sample selection, and statistical model one can find substantial differences in the decomposition into risk and permanent inequality.

[19] This affects the mean of log earnings as well as the arrival rate of nonemployment shocks.

[20] Mostly called Poisson processes in the literature.

Table 3: Baseline Parameters

Parameter description		Source	Value
Preferences			
ϕ	strength of keeping up motive	Bellet (2019)	0.7
ρ	discount rate	internal	0.02
ξ	utility weight of housing	internal	0.277
$\frac{1}{1-\varepsilon}$	intra-temporal elasticity of substitution	Flavin and Nakagawa (2008, AER)	0.15
γ	inverse intertemporal elasticity of substitution	standard	1.5
$\frac{1}{m}$	constant mortality rate	45 years worklife	45.0
Housing and financial technogy			
$\frac{\alpha}{1-\alpha}$	price elasticity of housing supply	Saiz (2010, QJE)	1.5
δ	depreciation rate of housing	Bureau of Economic Analysis	0.021
ω	maximum loan-to-value ratio	P95 of LTV	0.85
$a^S/\bar{y}$	exogenous net asst supply	cum. current account	-0.01
Taxation and Unemployment Insurance			
τ_0	level of taxes	internal	0.932
τ_1	progressivity	Heathcote et al. (2017)	0.15
b	replacement rate	Dept of Labor	0.32

Income Taxation We use the progressive income tax function from Heathcote et al. (2017),

$$T(y) = y - \tau_0 y^{1-\tau_1}.$$

If non-employed, households receive a fraction b of their potential earnings from unemployment insurance. Thus, the post-tax disposable income is given by

$$\tilde{y}_t = \begin{cases} y_{it}^{\text{pot}} - T(y_{it}^{\text{pot}}) & \text{if employed} \\ b y_{it}^{\text{pot}} & \text{otherwise.} \end{cases}$$

We follow Kaplan et al. (2020) in our choice of the parameters τ_0, τ_1. The progressivity parameter τ_1 is an estimate from Heathcote et al. (2017) and the scale parameter τ_0 is set to match the tax revenue from personal income tax and social security contribution as a share of GDP in 1980 (14.4%).[21] We set the replacement rate to 32%, matching average unemployment insurance benefits, as a fraction of average wage, as reported by the US Department of Labor.[22]

Preferences and Demographics The discount rate ρ and the utility weight of housing status ξ are internally calibrated to match the economy-wide mortgage-debt-to-income and loan-to-value ratios from the 1983 SCF. The interpretation of the utility weight ξ differs from other models, because ξ is the utility weight of housing status (not housing stock).

[21] Retrieved from `https://taxfoundation.org/federal-tax-revenue-source-1934-2018/`.

[22] Retrieved from `https://oui.doleta.gov/unemploy/DataDashboard.asp`.

The literature has not yet converged to a common value for the intratemporal elasticity of substitution $\frac{1}{1-\varepsilon}$. Estimates range from 0.13–0.24 (from structural models; e.g. Flavin and Nakagawa, 2008; Bajari et al., 2013) up to 1.25 (Ogaki and Reinhart, 1998; Piazzesi et al., 2007, using estimates from aggregate data). Many papers have picked parameters out of this range.[23] We follow the evidence from structurally estimated models and set the elasticity to 0.15.

The inverse intertemporal elasticity of substitution γ is set to the standard value 1.5. The constant annual mortality rate $m = 1/45$ is set to get an expected (working) lifetime of 45 years.

Social Comparisons For the status function we use a ratio-specification $s(h, \bar{h}) = \frac{h}{\bar{h}^{\phi}}$ as in Abel (1990). Bellet (2019) shows that this functional form captures the empirical finding that the utility loss from big houses decreases with own house size. Households with a medium sized house are more affected by top housing than households living in a small house.[24]

We define the reference measure as the 90^{th} percentile of the (endogenous) housing distribution, $\bar{h} = h_{\text{P90}}$. This follows Bellet (2019) who shows that households are only sensitive to changes in the top quintile of the house (size) distribution and strongest when the reference measure is defined as the 90^{th} percentile.[25]

The parameter ϕ pins down the strength of the comparison motive. It is the ratio of two utility elasticities

$$\phi = -\frac{\text{elasticity of utility w.r.t. } \bar{h}}{\text{elasticity of utility w.r.t. } h}. \tag{12}$$

If reference housing improves by 1%, then agents would have to improve their own house ϕ% to keep utility constant. Bellet (2019) estimates ϕ to be between 0.6 and 0.8 when setting $\bar{h}$ equal to the 90^{th} percentile of the housing distribution. We thus choose $\phi = 0.7$.[26] Note that Bellet (2019) estimates exactly this sensitivity using data on housing satisfaction which allows us to take his estimates without an intermediate indirect inference procedure. However, we note that this value for ϕ is likely an upper bound as our model does not have a spatial dimension. The implicit assumption is that the rise in top incomes and hence reference housing is equally spread across space.[27]

[23] Garriga and Hedlund (2020) use 0.13, Garriga et al. (2019) use 0.5, many papers use Cobb-Douglas (that is, an elasticity of 1.0, e.g. Berger et al., 2018; Landvoigt, 2017) and Kaplan et al. (2020) use 1.25.

[24] Note that the more tractable linear specification $(h - \phi\bar{h})$ as used in Campbell and Cochrane (1999), Ljungqvist and Uhlig (2000) and Section 4 would imply the opposite relationship between own house size and comparison strength.

[25] See Figure 6 in Bellet (2019).

[26] See Table 2 in Bellet (2019)

[27] Bellet (2019) shows that the estimate of ϕ depends on the distance between the reference house and one's own house. The point estimate is greater or equal to 0.7 when a big house is built within a radius of up to 25 miles and close to zero for changes in the reference house that happen further away.

TABLE 4: Targeted Moments

Moment	Model	Data (80/83)
aggregate loan-to-value	0.24	0.24
aggregate networth-to-income	4.63	4.60
tax-revenue-to-income	0.14	0.14

Technology and Financial Markets The construction technology parameter α is set to 0.6 so that the price elasticity of housing supply ($\frac{\alpha}{1-\alpha}$) equals 1.5, which is the median value across MSAs estimated by Saiz (2010). The maximum admissible loan-to-value ratio (ω) is set to 0.85, to match the 95th percentile of the LTV distribution in the SCF (Kaplan et al., 2020, use a similar approach for setting the debt-service-to-income constraint). Finally, we specify the exogenous net supply of assets a^S to match the net foreign debt position of the US. The net foreign debt position can be well approximated by the cumulative current account deficit of the US (Gourinchas et al., 2017), which was 1% of GDP in 1980 (see also Figure 13).

5.1 Internal Calibration and Model Fit

For the internal calibration we target the aggregate networth-to-income ratio (4.6) and the aggregate loan-to-value ratio (0.24) from the first wave of the Survey of Consumer Finances in 1983. We pick the utility weight of housing ξ and the the discount rate ρ so that simulated moments match their counterparts in the data. Table 4 shows that the model fits the data very well.

6 Quantitative Results

In this section we study how the model economy reacts to changes in the environment in the long-run. We compare the initial stationary equilibrium (corresponding to 1980) with the stationary equilibria where we increase income inequality to the level of 2007. Afterwards we set these results into perspective by comparing the effects of rising inequality and KURJ to the effect of increasing credit supply (Global Saving Glut). Lastly, we study the effects of both mechanisms combined.

6.1 Effects of Rising Inequality & KURJ

We now move to the main experiment of the paper. We start from the steady-state calibrated to the U.S. economy in 1980. We then raise income inequality to match the level in 2007 and solve for the new general equilibrium. Before getting to the results, we describe how we model the increase in income inequality.

Figure 10: Steady State Effects – 1980 vs. 2007

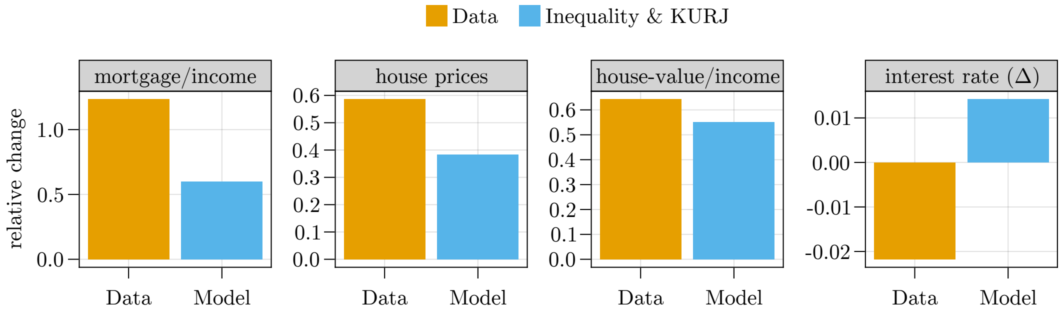

Notes: This figure shows relative changes in aggregate variables between the steady states in 1980 and 2007 and the corresponding changes in the data. Data: DINA.

As we discuss in Section 5, the cross-sectional dispersion of income has increased substantially between 1980 and 2007. Given the evidence in Kopczuk et al. (2010) and Guvenen et al. (2014), we attribute this change in cross-sectional inequality to changes in permanent inequality. In our model permanent inequality is reflected by the standard deviation of the distribution of the permanent component σ_α of the income process. Hence, we increase σ_α to match the increase in the cross-sectional P90/P50 ratio.

Figure 10 summarizes the steady state comparison by comparing the changes in the mortgage-to-income ratio, the house-value-to-income ratio, house prices and interest rates from 1980 to 2007 in the model and the data. In the model, rising inequality and the presence of KURJ create both a mortgage boom and a house price boom. The left panel of Figure 10 shows that this mechanism generates an increase in the mortgage-to-income ratio of about 60%—roughly half of the increase that is observed in the data where the mortgage-to-income ratio went up by 123%. The increase in housing demand puts upward pressure on house prices which increase by 38% in the model. This corresponds to about 60% of the 58% house price increase observed in the data. The shift towards housing is also reflected in the house-value-to-income ratio which goes up by 55% versus 64% in the data.

The housing and mortgage boom in the model is the result of two channels. The first channel is the direct comparison effect. Households increase their housing demand to keep up with the new reference measure set by the rich. This channel raises the demand for mortgage debt through an increase in housing demand. Second, rising top incomes raise the demand for housing and thus house prices because the richer households want to live in bigger houses. Since housing and non-durable consumption are complements, the expenditure share on housing goes up for all households as the house price increases. This indirect effect works even in the absence of KURJ.

FIGURE 11: Decomposition of Steady State Effect – The Importance of KURJ

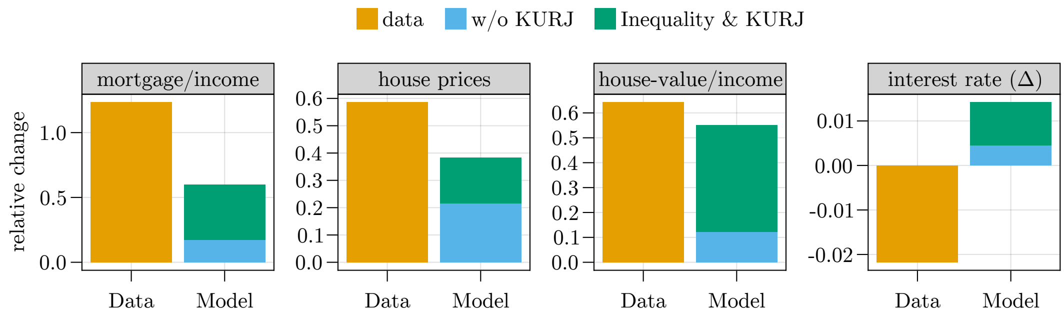

Notes: Comparison simulated changes in aggregate variables between the steady states in 1980 and 2007. "w/o KURJ" shows the changes when the reference measure $\bar{h}$ is kept fixed at level $\bar{h}_{1980}$ from the initial stationary equilibrium. Data: DINA.

Figure 11 shows that the social comparison motive plays a quantitatively important role. In particular, we show how much of the overall effect can be obtained with rising inequality in the absence of KURJ.[28] Without KURJ, rising inequality generates a 20% increase in the mortgage-to-income ratio and a 22% increase in house prices. That implies that KURJ is required to generate most of the mortgage boom and almost half of the house price boom.

Note that the case without KURJ corresponds to setting $\phi = 0$. To the extent that our baseline parameter $\phi = 0.7$ (Bellet, 2019) is an upper bound, the graph also shows the range of effects for other choices of $\phi \in (0, 0.7]$. Importantly, even for intermediate values of ϕ, the contribution of KURJ is quantitatively important.

While rising top inequality and KURJ generate a sizable mortgage and housing boom, no one mechanism can explain all aspects of the data. Here, the increasing demand for mortgages results in a counterfactual prediction about the interest rate which increases in the model by 1.4 percentage points, but decreases in the data by 2 percentage points.

Figure 12 shows the relative change in the mortgage-to-income ratio, the house-value-to-income ratio, and the loan-to-value ratio (leverage) across the income distribution. Both in the data and the model, the mortgage and housing boom spans the entire income distribution. For mortgages relative to income, the increase is especially strong in the bottom half of the income distribution. The model does a better job at the top of the income distribution relative to the bottom. The hump-shaped pattern is even more pronounced for the increase in house-values relative to incomes. Here, the model is broadly consistent

[28]Instead of re-calibrating the model with $s(h, \bar{h}) = h$ one can use that for a given reference measure $\bar{h}$ that is constant across the population, the initial equilibrium $\mathcal{E}$ is equivalent to a parameterization with $s(h, \bar{h}) = h$ and housing weight $\tilde{\xi}$ such that $\frac{\tilde{\xi}}{1-\tilde{\xi}} = \frac{\xi}{1-\xi}\frac{1}{\bar{h}^{\phi}}$. This holds because our specification of social comparisons, just re-weights the utility of housing and consumption.

FIGURE 12: Heterogeneity Across the Income Distribution

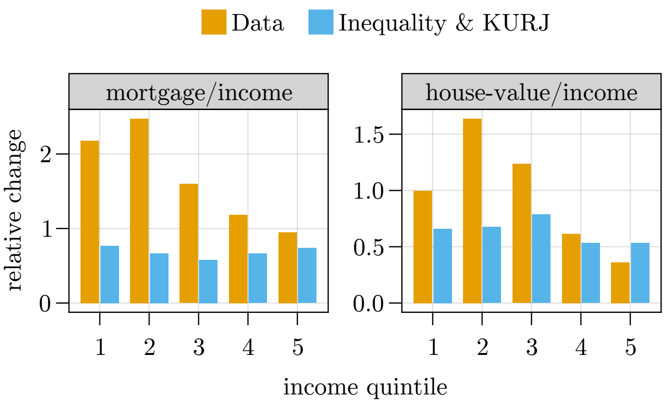

Notes: This figure shows the relative changes in the mortgage-to-income ratio, the house-value-to-income ratio and leverage (mortgage-to-house-value ratio) across the income distribution. Data: DINA.

with the heterogeneity across the income distribution as the increase is strongest in the middle and weakest at the top.

6.2 Comparison with the Global Saving Glut

Rising inequality and KURJ was certainly not the only driver of mortgage and housing boom. In order to put the quantitative results into perspective, we use our model to simulate the effects of the major supply side mechanism—the Global Saving Glut.

The Global Saving Glut refers to the accumulation of external debt, i.e. the cumulative current account deficit which is depicted in Figure 13. The cumulative current account was roughly zero in 1980, then started to rise and reached -40% of GDP in 2006. That is, the US was a net debtor with net debt amounting to 40% of GDP.[29] Bernanke (2005) proposes the steep increase in the global demand for savings—especially from China and India—as a potential explanation for this rise in foreign debt. He argues that these savings flowed into the US economy, building up the US debt position.

Through the lens of our model, the Global Saving Glut changes the market clearing condition (10) of the asset and mortgage market. Exogenous asset supply is given by a_t^S, where $a_t^S/\bar{y}_t$ is the cumulative current account from Figure 13 ($\bar{y}_t$ is average pre-tax earnings, our measure of GDP).

Figure 14 shows that the Global Saving Glut indeed causes a substantial increase in the mortgage-to-income ratio that is of the same order of magnitude. In contrast to the combination of inequality and KURJ, however, the Global Saving Glut can only account for a weak increase in house prices if inequality is held fixed at the 1980 level. If we also

[29] Gourinchas et al. (2017) estimate that the precise net foreign asset position was less negative due to valuation effects.

Figure 13: US Current Account Deficit

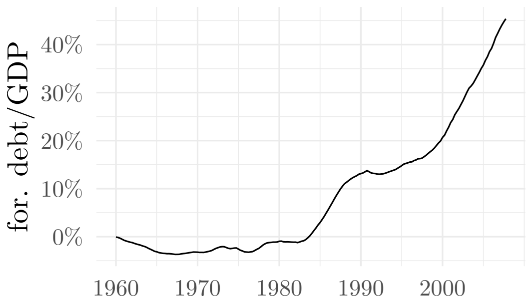

Notes: This figure shows the cumulative current account deficit (which is approximately US external debt) as a fraction of GDP. Source: BEA and FRED. For details, see Appendix B.

change inequality to the level in 2007, the combination of rising income inequality and the Global Saving Glut generates a moderate house price increase by 22%. The increase in the mortgage-to-income ratio does not change significantly when switching on the change in income inequality (37% instead of 32%).

Figure 15 shows the change in mortgages and house-values relative to income across the income distribution for the Global Saving Glut experiment. We find that the mortgage boom induced by the Global Saving Glut is mainly concentrated at the top of the income distribution. Independent of the average size of the effect, the same holds for house values.

Combining Inequality and KURJ with the Global Saving Glut

The results so far suggest that these demand- and supply-side mechanisms complement each other quite well. Rising inequality and KURJ generate a debt boom across the income distribution and drive up house prices and house values relative to income. The Global Saving Glut also contributes to the debt boom and puts downward pressure on the real interest rate. In the final part of the quantitative analysis, we therefore combine rising inequality and KURJ with the Global Saving Glut.

Figure 16 shows the joint effect of both mechanisms combined as well as the contribution of each mechanism. In particular, we first solve the model with only the Global Saving Glut (blue bar) and then add rising inequality and KURJ and re-solve the model to get the joint effect.

Both mechanisms together generate an increase in the mortgage-to-income ratio of 77%, an increase in house prices of 38% and an increase in the house-value-to-income ratio of 55%. We further find that the Global Saving Glut contributes slightly less to the increase in mortgages than the combination of rising inequality and KURJ. We further find that the contributions of the Global Saving Glut and the combination of rising inequality and KURJ are of the same order of magnitude. However, virtually all of the increase in house prices and house values relative to income can be attributed to rising inequality and

Figure 14: Steady State Effects for the Global Saving Glut

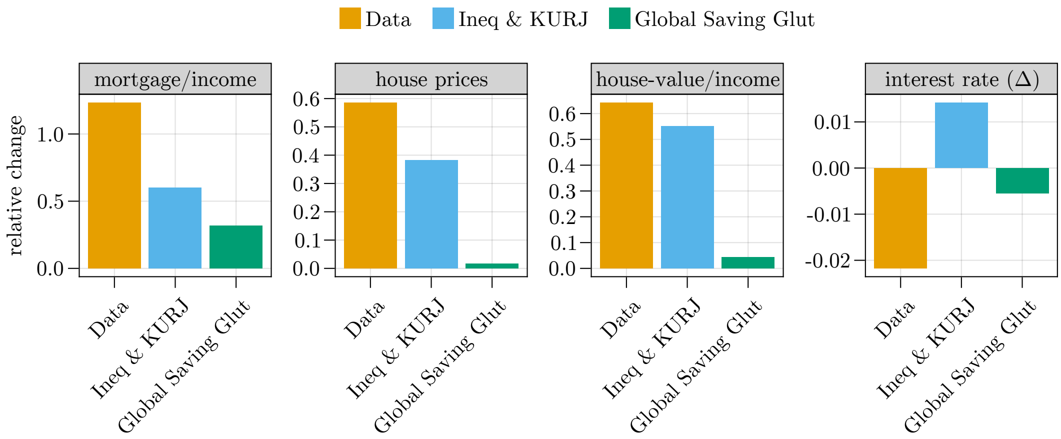

Notes: This figure shows relative changes in aggregate variables between the steady states in 1980 and 2007 for different scenarios and the corresponding changes in the data. Data: DINA. *Saving Glut*: Constant inequality and reference measure $\bar{h}$, varying a^S to match net foreign debt position (see Figure 13). Data: DINA.

KURJ. The total effect on the interest rate is positive even though the Global Saving Glut pulls it down.

Figure 15: Heterogeneity Across the Income Distribution

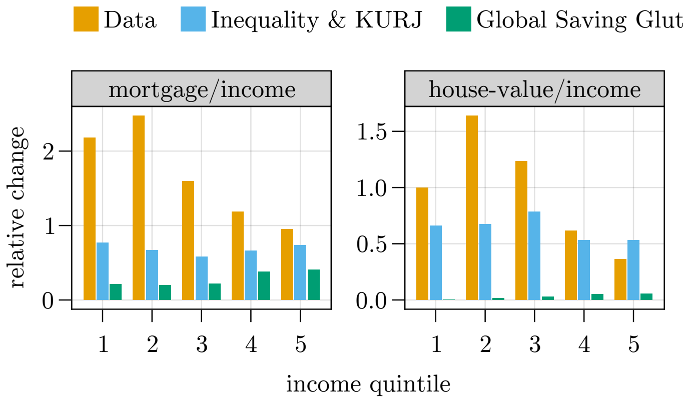

Notes: This figure shows the relative changes in the mortgage-to-income ratio, the house-value-to-income ratio and leverage (mortgage-to-house-value ratio) across the income distribution. Data: DINA.

Figure 16: Decomposition of the three mechanisms

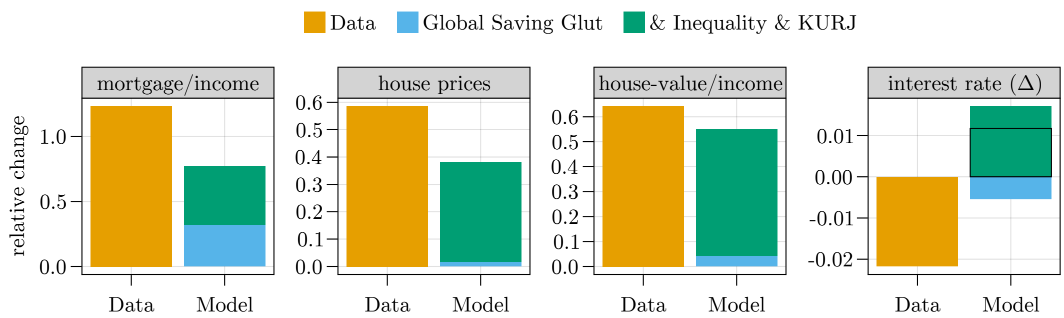

Notes: This figure shows the joint effect of the Global Saving Glut and the combination of rising inequality and KURJ. The blue bars correspond to the effects of the Global Saving Glut and the green bars show the marginal effect adding rising inequality and KURJ. The sum of the blue and green bars gives the total effect. Data: DINA.

7 Conclusion

This analysis was motivated by the parallel increase in top income inequality and mortgage debt in the US between 1980 and 2007. We first document novel aspects of the mortgage and housing boom. Using state-year variation, we find a strong positive relationship between lagged top incomes and (non-rich) household debt. Importantly, this state-level relationship is entirely driven by mortgage debt suggesting that housing played an important role in the transmission of rising top income inequality to rising household debt. Our finding that rising top incomes also drive up house prices underscores this and implies that rising housing demand contributed to the increase in mortgage debt.

Attempting to rationalize the sharp increase in household debt, previous studies have focused on supply-side mechanisms and the role of falling interest rates. Motivated by our empirical findings, we investigate—analytically and quantitatively—a demand-side mechanism to complement existing supply-side theories and rationalize the state-level findings. Our model, where households care not only about own consumption and housing but also about the housing benchmark set by the rich, is consistent with the findings of our empirical analysis in Section 2. The model predicts that *mortgage debt of the non-rich* and *house prices* rise as top incomes rise, while *non-mortage debt of the non-rich* is not affected by changes in top incomes (see Sections 4.3 and 4.4).

While the mechanism is consistent with the evidence along many dimensions, it generates a counter-factual prediction regarding interest rates. While the Global Saving Glut can rationalize falling interest rates, we find that it cannot rationalize the shift towards housing, particularly in the bottom half of the income distribution. We emphasize that we see rising inequality and KURJ as an important complement to supply-side drivers of the household debt boom. Our analysis suggests that the combination of supply- and demand-side factors is important in order to paint a complete picture of the US debt boom.

Bibliography

Abel, A. B. (1990): "Asset Prices under Habit Formation and Catching Up with the Joneses," *American Economic Review*, 80, 38–42.

Achdou, Y., J. Han, J.-M. Lasry, P.-L. Lions, and B. Moll (2021): "Income and Wealth Distribution in Macroeconomics: A Continuous-Time Approach," *Review of Economic Studies*, rdab002.

Adam, K., P. Kuang, and A. Marcet (2012): "House Price Booms and the Current Account," *NBER Macroeconomics Annual*, 26, 77–122.

Alvaredo, F., A. B. Atkinson, T. Piketty, E. Saez, and G. Zucman (2016): "The World Wealth and Income Database," .

Auclert, A. and M. Rognlie (2018): "Inequality and aggregate demand," Tech. rep., National Bureau of Economic Research.

Autor, D. H., L. F. Katz, and M. S. Kearney (2008): "Trends in Us Wage Inequality: Revising the Revisionists," *The Review of Economics and Statistics*, 90, 300–323.

Badarinza, C. (2019): "Mortgage debt and social externalities," *Review of Economic Dynamics*, 34, 43–60.

Bajari, P., P. Chan, D. Krueger, and D. Miller (2013): "A dynamic model of housing demand: Estimation and policy implications," *International Economic Review*, 54, 409–442.

Ballester, C., A. Calvó-Armengol, and Y. Zenou (2006): "Who's Who in Networks. Wanted: The Key Player," *Econometrica*, 74, 1403–1417.

Bellet, C. (2019): "The McMansion Effect: Top Size Inequality, House Satisfaction and Home Improvement in US Suburbs," Working paper, Erasmus University Rotterdam.

Berger, D., V. Guerrieri, G. Lorenzoni, and J. Vavra (2018): "House Prices and Consumer Spending," *Review of Economic Studies*, 85, 1502–1542.

Bernanke, B. S. (2005): "The Global Saving Glut and the U.S. Current Account Deficit," Speech 77, Board of Governors of the Federal Reserve System (U.S.).

Bertrand, M. and A. Morse (2016): "Trickle-Down Consumption," *Review of Economics and Statistics*, 98, 863–879.

Bursztyn, L., F. Ederer, B. Ferman, and N. Yuchtman (2014): "Understanding mechanisms underlying peer effects: Evidence from a field experiment on financial decisions," *Econometrica*, 82, 1273–1301.

Bursztyn, L., B. Ferman, S. Fiorin, M. Kanz, and G. Rao (2017): "Status Goods: Experimental Evidence from Platinum Credit Cards," *The Quarterly Journal of Economics*, qjx048.

Campbell, J. Y. and J. H. Cochrane (1999): "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, 107, 205–251.

Card, D., A. Mas, E. Moretti, and E. Saez (2012): “Inequality at Work: The Effect of Peer Salaries on Job Satisfaction,” *American Economic Review*, 102, 2981–3003.

Carr, M. D. and E. E. Wiemers (2016): “Recent Trends in the Variability of Men’s Earnings: Evidence from Administrative and Survey Data,” .

——— (2018): “New evidence on earnings volatility in survey and administrative data,” in *AEA Papers and Proceedings*, vol. 108, 287–91.

Charles, K. K., E. Hurst, and N. Roussanov (2009): “Conspicuous Consumption and Race*,” *The Quarterly Journal of Economics*, 124, 425–467.

Clark, A. E. and C. Senik (2010): “Who Compares to Whom? The Anatomy of Income Comparisons in Europe,” *Economic Journal*, 120, 573–594.

Coibion, O., Y. Gorodnichenko, M. Kudlyak, and J. Mondragon (2020): “Greater Inequality and Household Borrowing: New Evidence from Household Data,” *Journal of the European Economic Association*, 18, 2922–2971.

De Giorgi, G., A. Frederiksen, and L. Pistaferri (2020): “Consumption network effects,” *Review of Economic Studies*, 87, 130–163.

Duesenberry, J. (1949): *Income, Saving and the Theory of Consumer Behavior*, Cambridge, MA, USA: Harvard University Press.

Favilukis, J., S. C. Ludvigson, and S. van Nieuwerburgh (2017): “The macroeconomic effects of housing wealth, housing finance, and limited risk sharing in general equilibrium,” *Journal of Political Economy*, 125, 140–223.

Ferrer-i-Carbonell, A. (2005): “Income and Well-being: An Empirical Analysis of the Comparison Income Effect,” *Journal of Public Economics*, 89, 997–1019.

Flavin, M. and S. Nakagawa (2008): “A model of housing in the presence of adjustment costs: A structural interpretation of habit persistence,” *American Economic Review*, 98, 474–95.

Fogli, A. and V. Guerrieri (2019): “The End of the American Dream? Inequality and Segregation in US Cities,” Working Paper 26143, National Bureau of Economic Research.

Frank, R. (2013): *Falling behind: How Rising Inequality Harms the Middle Class*, vol. 4, Univ of California Press.

Gali, J. (1994): “Keeping up with the Joneses: Consumption externalities, portfolio choice, and asset prices,” *Journal of Money, Credit and Banking*, 26, 1–8.

Garriga, C. and A. Hedlund (2020): “Mortgage Debt, Consumption, and Illiquid Housing Markets in the Great Recession,” *American Economic Review*, 110, 1603–34.

Garriga, C., R. Manuelli, and A. Peralta-Alva (2019): “A macroeconomic model of price swings in the housing market,” *American Economic Review*, 109, 2036–72.

Ghiglino, C. and S. Goyal (2010): "Keeping up with the Neighbors: Social Interaction in a Market Economy," *Journal of the European Economic Association*, 8, 90–119.

Gourinchas, P.-O., H. Rey, and N. Govillot (2017): "Exorbitant privilege and exorbitant duty," Working paper, University of California, Berkeley.

Greimel, F. and F. Zadow (2019): "Understanding Housing Wealth Effects: Debt, Homeownership and the Lifecycle," Working paper, University of Mannheim.

Guvenen, F., G. Kaplan, J. Song, and J. Weidner (2018): "Lifetime incomes in the United States over six decades," .

Guvenen, F., F. Karahan, S. Ozkan, and J. Song (2021): "What Do Data on Millions of U.S. Workers Reveal About Life-Cycle Earnings Dynamics?" *Econometrica*.

Guvenen, F., S. Ozkan, and J. Song (2014): "The Nature of Countercyclical Income Risk," *Journal of Political Economy*, 122, 621–660.

Hazell, J., J. Herreño, E. Nakamura, and J. Steinsson (2020): "The Slope of the Phillips Curve: Evidence from U.S. States," Working Paper 28005, National Bureau of Economic Research.

Heathcote, J., K. Storesletten, and G. L. Violante (2010): "The Macroeconomic Implications of Rising Wage Inequality in the United States," *Journal of Political Economy*, 118, 681–722.

——— (2017): "Optimal tax progressivity: An analytical framework," *Quarterly Journal of Economics*, 132, 1693–1754.

Jordà, Ò., M. Schularick, and A. M. Taylor (2016): "The Great Mortgaging: Housing Finance, Crises and Business Cycles," *Economic Policy*, 31, 107–152.

Justiniano, A., G. E. Primiceri, and A. Tambalotti (2014): "The effects of the saving and banking glut on the US economy," *Journal of International Economics*, 92, S52–S67.

——— (2019): "Credit Supply and the Housing Boom," *Journal of Political Economy*, 127, 1317–1350.

Kaplan, G., K. Mitman, and G. L. Violante (2020): "The housing boom and bust: Model meets evidence," *Journal of Poltical Economy*.

Kaplan, G., B. Moll, and G. L. Violante (2018): "Monetary policy according to HANK," *American Economic Review*, 108, 697–743.

Katz, L. F. and K. M. Murphy (1992): "Changes in Relative Wages, 1963–1987: Supply and Demand Factors," *The Quarterly Journal of Economics*, 107, 35–78.

Kopczuk, W., E. Saez, and J. Song (2010): "Earnings Inequality and Mobility in the United States: Evidence from Social Security Data since 1937," *The Quarterly Journal of Economics*, 125, 91–128.

KUHN, M., M. SCHULARICK, AND U. I. STEINS (2017): "The Great American Debt Boom, 1949–2013," Working paper, University of Bonn.

KUHN, P., P. KOOREMAN, A. SOETEVENT, AND A. KAPTEYN (2011): "The effects of lottery prizes on winners and their neighbors: Evidence from the Dutch postcode lottery," *The American Economic Review*, 101, 2226–2247.

KUMHOF, M., R. RANCIÈRE, AND P. WINANT (2015): "Inequality, Leverage, and Crises," *American Economic Review*, 105, 1217–45.

LANDVOIGT, T. (2017): "Housing demand during the boom: The role of expectations and credit constraints," *Review of Financial Studies*, 30, 1865–1902.

LJUNGQVIST, L. AND H. UHLIG (2000): "Tax policy and aggregate demand management under catching up with the Joneses," *American Economic Review*, 356–366.

LUTTMER, E. F. P. (2005): "Neighbors as Negatives: Relative Earnings and Well-Being," *Quarterly Journal of Economics*, 120, 963–1002.

MIAN, A., L. STRAUB, AND A. SUFI (2021): "Indebted Demand," *The Quarterly Journal of Economics*, qjab007.

MIAN, A. R., L. STRAUB, AND A. SUFI (2020): "The Saving Glut of the Rich and the Rise in Household Debt," Working Paper 26941, National Bureau of Economic Research.

OGAKI, M. AND C. M. REINHART (1998): "Measuring intertemporal substitution: The role of durable goods," *Journal of Political Economy*, 106, 1078–1098.

PEREZ-TRUGLIA, R. (2019): "The Effects of Income Transparency on Well-Being: Evidence from a Natural Experiment," Working Paper 25622, National Bureau of Economic Research.

PIAZZESI, M. AND M. SCHNEIDER (2016): "Chapter 19 - Housing and Macroeconomics ," Elsevier, vol. 2 of *Handbook of Macroeconomics*, 1547 – 1640.

PIAZZESI, M., M. SCHNEIDER, AND S. TUZEL (2007): "Housing, Consumption and Asset Pricing," *Journal of Financial Economics*, 83, 531 – 569.

PIKETTY, T., E. SAEZ, AND G. ZUCMAN (2018): "Distributional national accounts: methods and estimates for the United States," *Quarterly Journal of Economics*, 133, 553–609.

RAJAN, R. (2010): *Fault Lines: How Hidden Fractures Still Threaten the World Economy*, Princeton, NJ, USA: Princeton University Press.

SAIZ, A. (2010): "The geographic determinants of housing supply," *Quarterly Journal of Economics*, 125, 1253–1296.

SMITH, M., D. YAGAN, O. ZIDAR, AND E. ZWICK (2019): "Capitalists in the Twenty-first Century," *The Quarterly Journal of Economics*, 134, 1675–1745.

SOLNICK, S. AND D. HEMENWAY (2005): "Are Positional Concerns Stronger in Some Domains Than in Others?" *American Economic Review*, 95, 147–151.

Stiglitz, J. (2009): "The global crisis, social protection and jobs," *International Labour Review*, 148, 1–13.

Straub, L. (2018): "Consumption, Savings, and the Distribution of Permanent Income," *Unpublished manuscript, Harvard University.*

van Treeck, T. (2014): "Did Inequality Cause the U.S. Financial Crisis?" *Journal of Economic Surveys*, 28, 421–448.

Veblen, T. (1899): *The Theory of the Leisure Class*, Macmillan: George Allen and Unwin.

A Additional Figures and Tables for Empirical Analysis

Figure A.1: Residualized Household Debt and Lagged Top Incomes by Income Groups

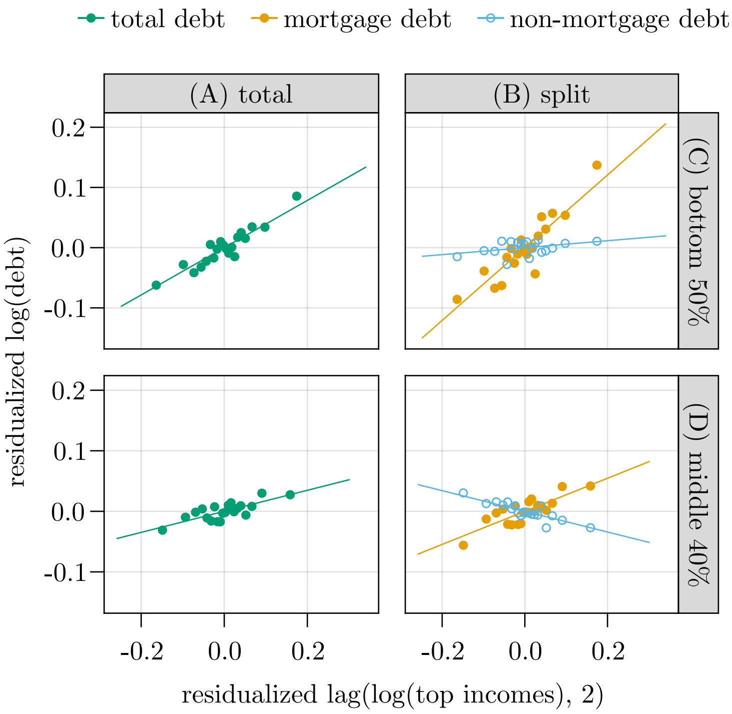

Notes: This figure shows the relationship between log debt (total, mortgage, non-mortgage) of households in the middle 40 and bottom 50 percent of the income distribution and the second lag of log average top 10% incomes conditional on state and year fixed effects and non-rich income. All variables are residualized using state and year fixed effects. The slope of the regression line is the OLS estimate of β reported in Table A.1. The figure shows averages in 20 equally sized bins of the x-variable.

TABLE A.1: Top Incomes and Household Debt: Fixed Effects Regressions by Income Groups

	log(debt$_{s,t}$)					
	(A) total		(B) mortgage		(C) non-mortgage	
	bottom 50%	middle 40%	bottom 50%	middle 40%	bottom 50%	middle 40%
	(1)	(2)	(3)	(4)	(5)	(6)
log(top incomes$_{s,t-2}$)	0.391***	0.148**	0.567***	0.211**	0.099	-0.115
	(0.112)	(0.068)	(0.163)	(0.098)	(0.069)	(0.084)
log(own income$_{s,t}$)	0.429***	0.949***	0.484***	1.122***	0.365***	0.560***
	(0.044)	(0.082)	(0.065)	(0.107)	(0.027)	(0.092)
State FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Estimator	OLS	OLS	OLS	OLS	OLS	OLS
N	1,224	1,224	1,224	1,224	1,224	1,224
R^2	0.950	0.976	0.895	0.963	0.970	0.972

Notes: This table shows the estimation results corresponding to equation 1. The dependent variable is either total, mortgage or non-mortgage debt in the middle 40 and bottom 50 percent of the income distribution. Robust standard errors in parentheses. The stars indicate the range of the p value: $^{***} \leq 0.01 \leq {}^{**} \leq 0.05 \leq {}^{*} \leq 0.1$.

FIGURE A.2: Long-Run Changes in Household Debt and Top Incomes by Income Groups

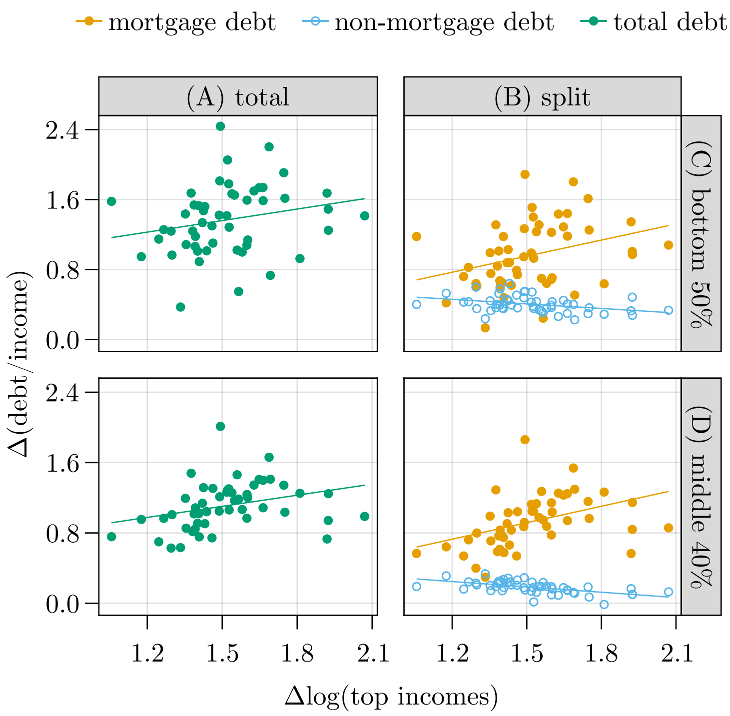

Notes: This figure shows the relationship between the change in debt-to-income (total, mortgage, non-mortgage) of households in the middle 40 and bottom 50 percent of the income distribution and the change in the log of average top 10% incomes between 1982 and 2007 across US states.

Table A.2: Top Incomes and Household Debt: Long Run Changes

	Δlog(debt/income)											
	(A) total				(B) mortgage				(C) non-mortgage			
	all	bottom 90	bottom 50	middle 40	all	bottom 90	bottom 50	middle 40	all	bottom 90	bottom 50	middle 40
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Δlog(top incomes)	-0.101	0.412**	0.440	0.421**	0.124	0.618***	0.613**	0.624***	-0.225***	-0.206***	-0.173***	-0.203***
	(0.128)	(0.184)	(0.279)	(0.185)	(0.128)	(0.187)	(0.258)	(0.193)	(0.025)	(0.038)	(0.063)	(0.041)
Estimator	OLS	OLS	OLS	OLS	OLS	OLS	OLS	OLS	OLS	OLS	OLS	OLS
N	51	51	51	51	51	51	51	51	51	51	51	51
R^2	0.013	0.093	0.048	0.096	0.019	0.182	0.103	0.176	0.633	0.376	0.135	0.332

Notes: This table shows the results of OLS regressions of the long-run change in the debt-to-income ratio (total, mortgage, non-mortgage) of different groups (all households, bottom 90%, middle 40%, bottom 50%) on the long-run change in the log of average top 10% incomes.

TABLE A.3: Top Incomes and Household Debt: Long Run Changes *Without Outliers*

	Δlog(debt/income)					
	(A) total		(B) mortgage		(C) non-mortgage	
	population	non-rich	population	non-rich	population	non-rich
	(1)	(2)	(3)	(4)	(5)	(6)
Δlog(top incomes)	0.207	0.891***	0.460***	1.147***	-0.253***	-0.256***
	(0.161)	(0.230)	(0.158)	(0.228)	(0.033)	(0.051)
Estimator	OLS	OLS	OLS	OLS	OLS	OLS
N	47	47	47	47	47	47
R^2	0.036	0.250	0.159	0.359	0.561	0.357

Notes: This table shows the results of OLS regressions of the long-run change in the debt-to-income ratio (total, mortgage, non-mortgage) of all households and non-rich households on the long-run change in the log of average top 10% incomes using all states except for the following 4 states with especially high growth in top incomes: New York, Massachusetts, Connecticut and the District of Columbia.

FIGURE A.3: Long-Run Changes in Household Debt and Top Incomes by Income Groups *Without Outliers*

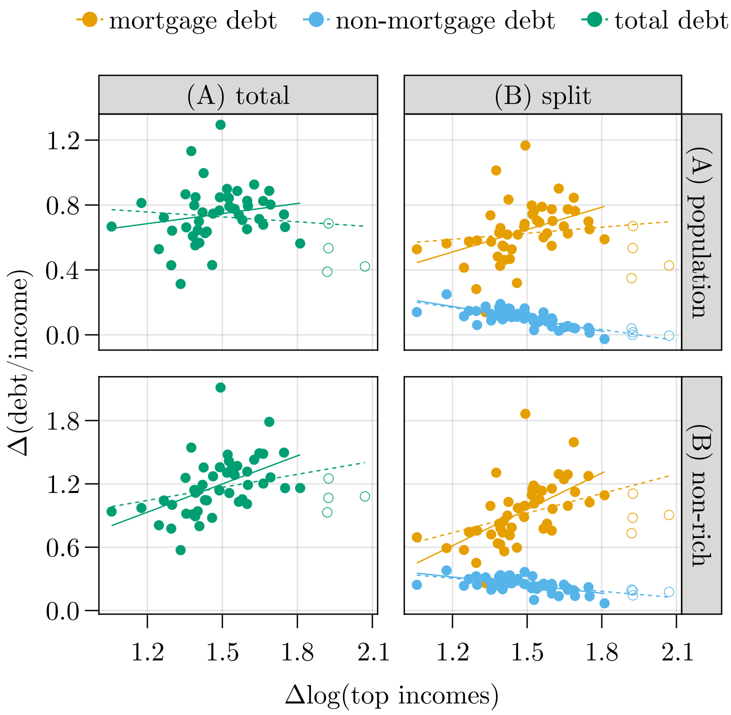

Notes: This figure shows the relationship between the change in debt-to-income (total, mortgage, non-mortgage) of all households and non-rich households and the change in the log of average top 10% incomes between 1982 and 2007 across all US states except for the following 4 states with especially high growth in top incomes: New York, Massachusetts, Connecticut and the District of Columbia. These states are depicted using hollow markers.

Figure A.4: Dynamic Effects of Top Incomes on Household Debt by Income Groups

Notes: This figure shows the cumulative effect of a 1% change in top 10% incomes on total, mortgage and non-mortgage debt of different income groups estimated from equation 2. The confidence bands are constructed using a significance level of 5%.

B Data Sources

Figure 1: Aggregate debt and inequality We use data on outstanding household debt from the US Flow of Funds, retrieved from FRED: total debt (`TLBSHNO`) and mortgages (`HMLBSHNO`). *Other debt* is constructed as the difference between total debt and mortgages. Debt is displayed as a share of nominal GDP (Bureau of Economic Analysis, BEA, via FRED: `GDP`).

The top 10% income share is from the World Wealth and Income Database (Alvaredo et al., 2016).

Figure 13: Net foreign debt position of the US We use the current account and GDP series from the BEA, retrieved via FRED (`BOPBCA`, `GDP`). Following Gourinchas et al. (2017) we compute the cumulative sum of the current account

$$\text{cum CA}_t = \sum_{i=1960}^{t} \text{CA}_t$$

and show it as a fraction of GDP in that given year $\frac{\text{cum CA}_t}{\text{GDP}_t}$.

C Proofs

C.1 Lemmas

Lemma 1. *The necessary conditions for an optimum of the households' problem are*

$$u_c\big(c_t, s(h_t, \bar{h}_t)\big) = \lambda_t \tag{C.1}$$

$$u_s\big(c_t, s(h_t, \bar{h}_t)\big) s_h(h_t, \bar{h}_t) = \lambda_t(r+\delta)p \tag{C.2}$$

$$\dot{\lambda}_t - \rho\lambda_t = -r\lambda_t \tag{C.3}$$

where λ is the co-state in the continuous time optimization problem.

Proof. Without adjustment costs, the two endogenous state variables a_t and h_t collapse into one state variable net worth w_t.

$$\dot{w}_t = rw_t + y_t - (r+\delta)ph_t - c_t$$

The present-value Hamiltonian is

$$H(w, h, c, \lambda) = u\big(c, s(h, \bar{h})\big) + \lambda\big(rw_t + y_t - (r+\delta)ph_t - c_t\big),$$

where w is the state, c and h are the controls and λ is the co-state. The necessary conditions are

$$\begin{aligned}
\frac{\partial H(w_t,h_t,c_t,\lambda_t)}{\partial c} &= u_c\big(c_t, s(h_t,\bar{h}_t)\big) - \lambda_t = 0 \\
\frac{\partial H(w_t,h_t,c_t,\lambda_t)}{\partial h} &= u_s\big(c_t, s(h_t,\bar{h}_t)\big)s_h(h_t,\bar{h}_t) - \lambda_t(r+\delta)p = 0 \\
\dot{\lambda}_t - \rho\lambda_t = \frac{\partial H(w_t,h_t,c_t,\lambda_t)}{\partial w} &= -r\lambda_t. \qquad \square
\end{aligned}$$

Lemma 2. *Under our assumption of CRRA-CES preferences, the optimal relation of c_t and h_t is given by*

$$\frac{\xi}{1-\xi}\Big(\frac{s(h_t,\bar{h}_t)}{c_t}\Big)^{\varepsilon-1} s_h(h_t,\bar{h}_t) = (r+\delta)p. \tag{C.4}$$

Further assuming Assumption 3 yields

$$c_t = \kappa_0 h_t - \kappa_0\phi\bar{h}_t, \quad \textit{where } \kappa_0 = \Big((r+\delta)p\frac{1-\xi}{\xi}\Big)^{\frac{1}{1-\varepsilon}}. \tag{C.5}$$

Proof. Combining conditions (C.1) and (C.2) yields

$$\frac{u_s(c_t,s_t)}{u_c(c_t,s_t)} s_h(h_t,\bar{h}_t) \stackrel{!}{=} (r+\delta)p. \tag{C.6}$$

For the given CRRA-CES preferences the marginal utilites are given by

$$u_c(c_t,s_t) = ((1-\xi)c_t^\varepsilon + \xi s_t^\varepsilon)^{\frac{1-\gamma}{\varepsilon}-1}(1-\xi)c_t^{\varepsilon-1} \tag{C.7}$$

$$u_s(c_t,s_t) = ((1-\xi)c_t^\varepsilon + \xi s_t^\varepsilon)^{\frac{1-\gamma}{\varepsilon}-1}\xi s_t^{\varepsilon-1}. \tag{C.8}$$

Thus,

$$\frac{u_s(c_t,s_t)}{u_c(c_t,s_t)} = \frac{\xi}{1-\xi}\Big(\frac{s_t}{c_t}\Big)^{\varepsilon-1}. \tag{C.9}$$

Plugging in above yields the first statement. Using Assumption 3 we get

$$\frac{\xi}{1-\xi}\Big(\frac{h_t-\phi\bar{h}}{c_t}\Big)^{\varepsilon-1} = (r+\delta)p. \tag{C.10}$$

$$\Big(\frac{c_t}{h_t-\phi\bar{h}}\Big) = \Big((r+\delta)p\frac{1-\xi}{\xi}\Big)^{\frac{1}{1-\varepsilon}} = \kappa_0 \tag{C.11}$$

$$c_t = \kappa_0 h_t - \kappa_0\phi\bar{h}_t \qquad \square$$

Lemma 3. *Under the assumption of time-constant house prices p, and all previous assumptions of this section, individual choices c_t, h_t are constant over time.*

Proof. The costate λ is constant over time. This follows from using Assumption 1 in condition (C.3), which gives $\dot{\lambda}_t = 0$.

Plugging in (C.5) in condition (C.2) one gets that an decreasing function of h is constant over time, thus h_t is constant over time. Knowing that h_t constant over time, and a similar argument for condition (C.1) it follows that c_t is constant over time. $\square$

C.2 Proof of Proposition 1

From the lemmas above we get that

$$c = \kappa_0 s(h, \bar{h}) = \kappa_0 h - \kappa_0 \phi \bar{h}. \tag{C.12}$$

Using the lifetime budget constraint we get

$$\mathcal{Y} := ra_0 + y = ph(r+\delta) + c \tag{C.13}$$

$$= h\Big(p(r+\delta) + \kappa_0\Big) - \kappa_0 \phi \bar{h} \tag{C.14}$$

$$\implies h = \frac{\mathcal{Y} + \kappa_0 \phi \bar{h}}{p(r+\delta)+\kappa_0} = \underbrace{\frac{1}{p(r+\delta)+\kappa_0}}_{\kappa_2} \mathcal{Y} + \underbrace{\frac{\kappa_0}{p(r+\delta)+\kappa_0}}_{\kappa_1} \phi \bar{h} = \kappa_2 \mathcal{Y} + \kappa_1 \phi \bar{h} \tag{C.15}$$

where

$$\kappa_1 := \frac{\kappa_0}{p(r+\delta)+\kappa_0} = \frac{1}{\frac{p(r+\delta)}{\kappa_0} + 1} \in (0,1) \tag{C.16}$$

since

$$\frac{p(r+\delta)}{\kappa_0} = \Big(\frac{1}{(r+\delta)p}\Big)^{\frac{1}{1-\varepsilon}-1} \Big(\frac{\xi}{1-\xi}\Big)^{\frac{1}{1-\varepsilon}} > 0. \tag{C.17}$$

Stacking equations (C.15) for and using $\bar{\boldsymbol{h}} = G\boldsymbol{h}$

$$\boldsymbol{h} = \kappa_2 \boldsymbol{\mathcal{Y}} + \kappa_1 \phi G \boldsymbol{h} \tag{C.18}$$

$$\boldsymbol{h} = (I - \kappa_1 \phi G)^{-1} \kappa_2 \boldsymbol{\mathcal{Y}} = \Big(\sum_{i=0}^{\infty} (\kappa_1 \phi G)^i\Big) \kappa_2 \boldsymbol{\mathcal{Y}}. \tag{C.19}$$

Moreover,

$$\bar{\boldsymbol{h}} = G\boldsymbol{h} = \frac{\kappa_1\phi}{\kappa_1\phi} G\Big(\sum_{i=0}^{\infty}(\kappa_1\phi G)^i\Big)\kappa_2\boldsymbol{\mathcal{Y}} \tag{C.20}$$

$$= \frac{1}{\kappa_1\phi}\Big(\sum_{i=1}^{\infty}(\kappa_1\phi G)^i\Big)\kappa_2\boldsymbol{\mathcal{Y}} \tag{C.21}$$

$$= \frac{1}{\kappa_0\phi}\Big(\sum_{i=1}^{\infty}(\kappa_1\phi G)^i\Big)\boldsymbol{\mathcal{Y}} \tag{C.22}$$

$(I - \kappa_1\phi G)^{-1}$ is a Leontief inverse. It exists if the matrix power series $\sum_{i=0}^{\infty}(\kappa_1\phi G)^i$ converges[30]. In that case

$$(I - \kappa_1\phi G)^{-1} = \sum_{i=0}^{\infty}(\kappa_1\phi G)^i.$$

Now, we calculate debt.

$$-ra = y - \delta ph - c$$

using C.12,

$$\begin{aligned}
&= y - \delta ph - \kappa_0 h + \kappa_0\phi\bar{h} \\
&= y - (\delta p + \kappa_0)h + \kappa_0\phi\bar{h} \\
-r\boldsymbol{a} &= \boldsymbol{y} - (\delta p + \kappa_0)\quad \underbrace{\Big(\sum_{i=0}^{\infty}(\kappa_1\phi G)^i\Big)}_{=I+\left(\sum_{i=1}^{\infty}(\kappa_1\phi G)^i\right)}\quad \kappa_2\boldsymbol{\mathcal{Y}} + \Big(\sum_{i=1}^{\infty}(\kappa_1\phi G)^i\Big)\boldsymbol{\mathcal{Y}} \\
&= \boldsymbol{y} - \kappa_3\boldsymbol{\mathcal{Y}} + (1-\kappa_3)\Big(\sum_{i=1}^{\infty}(\kappa_1\phi G)^i\Big)\boldsymbol{\mathcal{Y}}
\end{aligned}$$

where

$$\kappa_3 = (\delta p + \kappa_0)\kappa_2 = \frac{\delta p + \kappa_0}{p(r+\delta)+\kappa_0} = \frac{1}{1 + \frac{pr}{\delta p + \kappa_0}} \in (0,1).$$

D Details on the Earnings Process

The innovations of both the transitory and persistent process are drawn from mixture distributions to match higher order moments of income risk and impulse response functions.

[30] This is the case for all nilpotent matrices (there exists a power p such that $G^p = 0I$) (there are no infinitely-long paths in the network) or if all eigenvalues of $\kappa_1\phi G$ are between 0 and 1. This holds whenever G can be interpreted as a Markov Chain.

Finally, Guvenen et al. (2021) show that a non-employment shock with z-dependent shock probabilities greatly improves the model fit.[31]

If employed, individual pre-tax earnings are given by

$$y_{it}^{\text{pot}} = \exp(\tilde{\alpha}^i + z_{it} + \epsilon_{it}).$$

We will refer to y^{pot} as *potential earnings.* The actual pre-tax earnings (taking into account unemployment) are

$$y_{it} = (1 - \nu_{it}) y_{it}^{\text{pot}},$$

where

$$\begin{aligned}
\tilde{\alpha}_i &\sim \mathcal{N}(\mu_\alpha, \sigma_\alpha), \\
\mathrm{d}z_{it} &= -\theta^z z_{it}\mathrm{d}t + \mathrm{d}J_{it}^z, \\
\mathrm{d}\epsilon_{it} &= -\theta^\epsilon z_{it}\mathrm{d}t + \mathrm{d}J_{it}^\epsilon.
\end{aligned}$$

J_{it}^z is a jump-process that arrives at rate λ^z. The size of the jump, η_{it}^z is drawn from a mixture of two normal distributions,

$$\eta_{it}^z = \begin{cases} \mathcal{N}\big(\mu^z(1-p^z), \sigma_1^z\big) & \text{with prob. } p^z \\ \mathcal{N}\big(-p^z\mu^z, \sigma_2^z\big) & \text{with prob. } 1-p^z. \end{cases}$$

Similarly, the jump process for the transitory process arrives at rate λ^ϵ and the jump size, η_{it}^z is drawn from a mixture of two normal distributions,

$$\eta_{it}^\epsilon = \begin{cases} \mathcal{N}\big(-\epsilon_{it} + \mu^\epsilon(1-p^\epsilon), \sigma_1^\epsilon\big) & \text{with prob. } p^\epsilon \\ \mathcal{N}\big(-\epsilon_{it} - p^\epsilon\mu^\epsilon, \sigma_2^\epsilon\big) & \text{with prob. } 1-p^\epsilon. \end{cases}$$

The key difference between the persistent and the transitory process is that the jumps in the former are added to the current state whereas the jumps in the latter process reset the process such that the post-jump state is centered around zero.

The nonemployment shock arrives at rate $\lambda_0^\nu(z_{it})$ and has average duration $1/\lambda_1^\nu$. Specifically, the arrival probability as a function of the current state of the persistent process is modeled as

$$\lambda_0^\nu(z_{it})\mathrm{d}t = \frac{\exp\big(a + bz_{it}\big)}{1 + \exp\big(a + bz_{it}\big)}.$$

[31] The only component that is missing compared to the Benchmark process is fixed heterogeneous income profiles, i.e. ex-ante permanent heterogeneity in lifecycle income growth rates.

TABLE D.1: Earnings Process Parameters

Parameter		Value
Fixed Effects		
μ_α	mean	$2.7408 + 0.4989\bar{t} - 0.1137\bar{t}^2$
σ_α	standard deviation	0.467
Persistent Process		
λ^z	arrival rate	1.0
θ^z	mean reversion rate	$-\log(0.983)$
p^z	mixture probability	0.267
μ^z	location parameter	-0.194
σ_1^z	std. dev. of first Normal	0.444
σ_2^z	std. dev. of second Normal	0.076
σ_0^z	std. dev. of z_{i0}	0.495
Transitory Shocks		
λ^ϵ	arrival rate	1.0
θ^ϵ	mean reversion rate	0.0
p^ϵ	mixture probability	0.092
μ^ϵ	location parameter	0.352
σ_1^ϵ	std. dev. of first Normal	0.294
σ_2^ϵ	std. dev. of second Normal	0.065
Nonemployment Shocks		
a	constant	$-3.2740 - 0.8935\bar{t}$
b	slope	$-4.5692 - 2.9203\bar{t}$
λ_1^ν	exit rate	1/0.9784

Table D.1 shows all parameters of our continuous time earnings process.

E Numerical Solution for a Stationary Equilibrium

We first describe how we discretize the complex income process, then we show how to solve the partial equilibrium using a finite difference method from Achdou et al. (2021). Finally we present the algorithm used to compute equilibrium prices and reference measure.

The model was solved using version 1.2 of the `Julia` language. For a given parameterization, 200 endogenous grid points and 2000 exogenous grid points solving for a general equilibrium takes about 30 minutes on standard laptop using just one core.

For the calibration we ran the code in parallel (using 30 nodes with 16 cores) for 12 hours on a high performance cluster.

E.1 Discretizing the Income Process

Pre-tax earnings depend on four exogenous states $\theta = (\tilde{\alpha}, z, \epsilon, \nu)$,

$$y(\theta) = (1-\nu)\exp(\tilde{\alpha} + z + \epsilon). \tag{E.1}$$

We first discretize the two jump-drift processes z and ϵ following the procedure of Kaplan et al. (2018). We discretize them separately, creating two continuous time Markov chains and combining them. The state space of the combined continuous time Markov Chain is given by

$$\{z_1, \dots, z_{N_z}\} \times \{\epsilon_1 \dots \epsilon_{N_\epsilon}\}. \tag{E.2}$$

Then we add non-employment states for each state, where the transition probabilities into the non-employment state are state-dependent. The state space of the CTMC with non-employment becomes

$$\{z_1, \dots, z_{N_z}\} \times \{\epsilon_1 \dots \epsilon_{N_\epsilon}\} \times \{0, 1\}. \tag{E.3}$$

Finally we add the permanent component $\tilde{\alpha}$. We choose $N_\alpha = 10$ grid points, where each of those grid points represents a decile of $\tilde{\alpha}$'s distribution. Conditional on drawing $\tilde{\alpha}_i$, the other three components follow the same CTMC with $N_z \cdot N_\epsilon \cdot 2$ states. Denote the changing states by $\tilde{\theta} = (z, \epsilon, \nu)$

The transition between states $\tilde{\theta}$ is given by the intensities q_{ij}. For an agent at state $\tilde{\theta}_i$ the probability of jumping to a new state $\tilde{\theta}_j$ within the time short time period Δ is approximately given by $p_{ij}(\Delta) \approx q_{ij}\Delta$. More precisely, given the intensity matrix $Q = (q_{ij})$ where $q_{ij} \geq 0$ for $i \neq j$ and $q_{ii} = -\sum_{k \neq i} q_{ik}$, the matrix of transition probabilities is given by

$$P(\Delta) = \exp(-\Delta Q),$$

where exp is the matrix exponential. $P(\Delta)$ is a stochastic matrix.

E.2 Partial Equilibrium given p, r, $\bar{h}$

Given prices (p, r) and reference measure $\bar{h}$ the households' problem can be characterized by a coupled system of partial differential equations: the Hamilton-Jacobi-Bellman (HJB) equation and the Kolmogorov forward (KF) equation. The HJB equation describes the optimization problem of the households and the KF equation describes the evolution of the cross-sectional distribution $\mu(\mathrm{d}a, \mathrm{d}h, \mathrm{d}y)$.

We solve these two equations using the finite difference method from Achdou et al. (2021). The discretized system can be written as

$$\rho \boldsymbol{v} = \boldsymbol{u}(\boldsymbol{v}) + A(\boldsymbol{v}; r, p, \bar{h})\boldsymbol{v} \tag{HJB}$$

$$\mathbf{0} = \big(A(\boldsymbol{v}; r, p, \bar{h}) + M\big)^T \boldsymbol{g}, \tag{KF}$$

where $\boldsymbol{v}$ is the discretized value function, $\boldsymbol{g}$ is the discretized cross-sectional distribution, $\boldsymbol{u}(\boldsymbol{v})$ is the discretized flow utility, $A(\boldsymbol{v}; r, p, \bar{h})$ is the discretized infinitesimal generator of the HJB equation (a very sparse matrix) and M is a matrix that corrects the intensities for births and deaths. The discretized system reveals how tightly coupled the HJB and KF equations are. The matrix $A(\boldsymbol{v}; r, p, \bar{h})$ shows up in both equation. Once it is known from the solution of the HJB equation, it can be directly used to get the distribution $\boldsymbol{g}$ from the KF equation.

Solving the Hamilton-Jacobi-Bellman equation

We assume that housing h can be adjusted frictionlessly. So the two states h and a collapse into one, "net worth"

$$w_t = a_t + ph_t, \tag{E.4}$$

with its law of motion

$$\dot{w}_t = rw_t + y_t - (r + \delta)ph_t - c_t. \tag{E.5}$$

The collateral constraint can be rewritten in terms of w

$$w_t = ph_t + a_t \geq ph_t - \omega ph_t \tag{E.6}$$

$$\implies ph_t \leq \frac{w_t}{1-\omega}. \tag{E.7}$$

The households' HJB equation is

$$(\rho + m)v(w, \theta_i) = \max_{c, h \leq \frac{w}{1-\omega}} u(c, s(h, \bar{h})) \tag{E.8}$$

$$+ v_w(w, \theta_i)(rw + \theta_i - (r + \delta)ph - c) \tag{E.9}$$

$$+ \sum_{k \neq i} q_{ik}(v(w, \theta_k) - v(w, \theta_i)). \tag{E.10}$$

The intensities q_{ij} are the intensities of the continuous time Markov chain from Section E.1. In order to solve this equation, we need to replace the maximum operator with the maximized Hamiltonian. That is, we need to plug in the optimal policy functions $c^*(w, y)$, $h^*(w, y)$ which are given in Corollary 4 below. The result depends on the following lemma.

Lemma 4. *When the collateral constraint is slack, we get the optimality conditions*

$$h(w,y) = \left(\frac{1}{\tau_2}\left(\bar{h}^{\phi}(\rho+\delta)pv_w(w,y)\right)\right)^{-\frac{1}{\gamma}}\bar{h}^{\phi} \tag{E.11}$$

$$c(w,y) = s(h(w,y),\bar{h})\tau_1, \tag{E.12}$$

where $\tau_1 = \left((r+\delta)p\frac{1-\xi}{\xi}\bar{h}^{\phi}\right)^{\frac{1}{1-\varepsilon}}$ *and* $\tau_2 = ((1-\xi)\tau_1^{\varepsilon}+\xi)^{\frac{1-\gamma-\varepsilon}{\varepsilon}}\xi$.

Proof. Using the optimality conditions (C.4) and (C.2) with (C.8) we get

$$(r+\delta)p = \frac{u_s(c,s)}{u_c(c,s)}s_h(h,\bar{h}) = \frac{\xi}{1-\xi}\left(\frac{s(h,\bar{h})}{c}\right)^{\varepsilon-1}s_h(h,\bar{h}) \tag{E.13}$$

$$(\rho+\delta)pv_w(w,y) = u_s(c,s)s_h = ((1-\xi)c^{\varepsilon}+\xi s^{\varepsilon})^{\frac{1-\gamma}{\varepsilon}-1}\xi s^{\varepsilon-1}s_h. \tag{E.14}$$

Using (E.13) we express optimal c as a function of optimal s

$$c(h,\bar{h}) = s(h,\bar{h})\left((r+\delta)p\frac{1-\xi}{\xi}\frac{1}{s_h(h,\bar{h})}\right)^{\frac{1}{1-\varepsilon}} \tag{E.15}$$

using the ratio specification for s

$$= s(h,\bar{h})\left((r+\delta)p\frac{1-\xi}{\xi}\bar{h}^{\phi}\right)^{\frac{1}{1-\varepsilon}} =: s(h,\bar{h})\tau_1. \tag{E.16}$$

Then we can plug this expression into (E.14) and get

$$(\rho+\delta)pv_w(w,y) = ((1-\xi)(\tau_1 s)^{\varepsilon}+\xi s^{\varepsilon})^{\frac{1-\gamma-\varepsilon}{\varepsilon}}\xi s^{\varepsilon-1}s_h \tag{E.17}$$

$$= \underbrace{((1-\xi)\tau_1^{\varepsilon}+\xi)^{\frac{1-\gamma-\varepsilon}{\varepsilon}}\xi}_{=:\tau_2} s^{1-\gamma-\varepsilon}s^{\varepsilon-1}s_h \tag{E.18}$$

$$= \tau_2 s^{-\gamma}s_h \tag{E.19}$$

Thus we get

$$s(h,\bar{h}) = \left(\frac{(\rho+\delta)pv_w(w,y)}{\tau_2 s_h}\right)^{-\frac{1}{\gamma}}, \tag{E.20}$$

and using ratio-specification for s,

$$h = \left(\frac{1}{\tau_2}\left((\rho+\delta)pv_w(w,y)\bar{h}^{\phi}\right)\right)^{-\frac{1}{\gamma}}\bar{h}^{\phi}. \qquad \square$$

Corollary 4. *The optimal policies are given by*

$$h^*(w,y) = \begin{cases} h(w,y) & \textit{if } h(w,y) < \frac{w}{p(1-\omega)} \\ \frac{w}{p(1-\omega)} & \textit{otherwise} \end{cases}, \quad c^*(w,y) = \begin{cases} c(w,y) & \textit{if } h(w,y) < \frac{w}{p(1-\omega)} \\ \tilde{c}(w,y) & \textit{otherwise} \end{cases} \tag{E.21}$$

where $h(w,y)$ and $c(w,y)$ are from Lemma 4 and $\tilde{c}(w,y)$ is the solution to the optimality condtion for c, given $h = \frac{w}{p(1-\omega)}$,

$$v_w(w,y) = ((1-\xi)c^\varepsilon + \xi s^\varepsilon)^{\frac{1-\gamma-\varepsilon}{\varepsilon}}(1-\xi)c^{\varepsilon-1}, \tag{E.22}$$

which is solved numerically.

Given the optimal policies, it is straight-forward to solve the HJB using the implicit upwind scheme in Achdou et al. (2021).

Solving the Kolmogorov forward equation

We construct the birth and death matrix M as in Kaplan et al. (2018) and solve for the distribution using the implicit scheme from Achdou et al. (2021).

E.3 General equilibrium: Solving for r, p and $\bar{h}$

We use the following algorithm to compute general equilibria.

0. Guess r_0, p_0 and $\bar{h}_0$
(i) Clear housing markets given r_{n-1} and $\bar{h}_{n-1}$
    (i) Use Newton steps until the sign of the excess demand for housing changes
    (ii) Use Bisection to find the market clearing price p_n
(ii) Compute the excess demand on the asset market
(iii) Use a Newton step to update the interest rate r_n
(iv) Compute the implied reference measure $\bar{h}_x$ and update $\bar{h}_n = \bar{h}_{n-1} + a(\bar{h}_x - \bar{h}_{n-1})$
(v) If $r_n \approx r_{n-1}$ and $\bar{h}_n \approx \bar{h}_{n-1}$, an equilibrium has been found. If not, go back to step 1.

Chapter 3

Gender Gaps and the Role of Bosses

Joint with Felix Holub.

1 Introduction

While the gender wage gap has declined considerably, convergence has slowed down and substantial gender disparities persist. In the US, for example, the unadjusted gender gap has stagnated at around 19% since the turn of the century (BLS, 2019). The adjusted wage gap is even more persistent as, for the past three decades, women have been earning about 9% less than men after adjusting for differences in education, experience, industry and occupation (Blau and Kahn, 2017). Gender pay gaps thus continue to receive significant attention as policy makers discuss gender quotas and many firms and large organizations train managers to be more aware of gender-related biases (Chang et al., 2019). It has been suggested that the adjusted wage gap is driven by differences in productivity, negotiation prowess, temporal flexibility, or (unconscious) discrimination (e.g. Azmat and Ferrer, 2017; Babcock et al., 2003; Goldin, 2014; Sarsons, 2018, respectively). All of these explanations could be closely related to the behavior of bosses. The direct superiors of workers affect productivity, negotiate salaries, shape work environments, and evaluate performance and thereby determine bonus payments (e.g. Lazear et al., 2015; Hoffman and Tadelis, forthcoming; Frederiksen et al., 2019).

In this paper, we use novel personnel data of a large multinational firm in order to shed light on the role of managers for gender gaps. In particular, we ask two novel questions connecting the literature on managers to that on gender gaps. First, do women work for "worse" bosses than men? Similarly to the distribution across occupations, sorting of workers and managers may explain a part of the gender wage gap. In other words, if male workers in the same occupation work for better-paying bosses than their female colleagues, this will drive a wedge between the earnings of equally qualified men and women. Second, are male bosses bad for women? Three main mechanisms come to mind. First, men may undervalue the performance of women due to conscious or unconscious biases. Second, male managers may create work environments that make it harder for women to shine. Third, men might be more productive under male managers due to gender-specific complementarities in productivity.[1] The persistent under-representation of women in powerful positions may therefore be cause and consequence at the same time if—for some reason—female employees are systematically disadvantaged when having male superiors.

In order to separate out the different factors explaining gender disparities and investigate the importance of (male) managers for gender gaps, we bring in unique personnel data provided by one of the largest European manufacturing firms. The panel dataset covers the multinational's entire workforce in the period 2014–2019 and has several key advantages allowing us to address these questions. First, the data contain detailed information on job characteristics, sociodemographics, compensation, and performance evaluations. This allows us to identify the performance-related component of earnings. Second, we are able to trace out the organizational hierarchy and identify every employee's coworkers, superiors, and subordinates. Third, we can condition on time-invariant unobservable characteristics of both employees and their managers.[2]

The paper has two sets of results. The first quantifies to which extent the sorting of male and female workers to different managers can explain gender gaps while also taking into account the contributions of other observables such as sociodemographics and job characteristics. To that end, we implement a Kitagawa-Oaxaca-Blinder-decomposition for base salaries, bonus payouts, contracted bonus targets, and performance ratings (Kitagawa, 1955; Oaxaca, 1973; Blinder, 1973). For male and female workers, we run separate regressions of the outcome of interest on job characteristics, manager indicators, age, tenure, and location controls. The decomposition reveals the following three findings.

First, the raw gender gaps in base salary and bonus payouts are 12.3 log points and 22.2 log points, respectively. Men's contracted bonus targets are on average 2.8% greater than those of women. The raw gender gap in performance ratings is negative as men are two percentage points less likely than women to receive a high performance rating.

Second, 25% of the raw gender gap in base salary and 19% of the gender bonus gap are attributed to the sorting of male and female workers to different managers. The unexplained component of the gender gap is larger for bonuses (16.9%) than for base salaries (8.8%). For performance ratings, we find that a large part of the gender gap cannot be explained. However, while the contribution of standard observables such as age

[1]The opposite could also be true if women hinder other women, as for example has been documented by Bagues and Esteve-Volart (2010)

[2]We will refer to a worker's direct superior as manager. A manager is also a worker from the perspective of his or her manager.

or job characteristics drops substantially, the impact of managers remains sizable. Worker-manager sorting increases the performance gap (in favor of men) by 2.1 log points. This means that if women were assigned to the same jobs and managers, the performance gap would be even more negative.

Third, comparing similar employees doing the same job under the same manager, we find significant residual gender gaps in base salaries (1.1%), bonus targets (2.0%), and bonus payouts (3.8%). While all of these gaps favor men, the opposite holds for performance evaluations. Women are 3.5 percentage points (14.2% relative to the mean) *more* likely to receive high a performance rating, which implies above-target bonus payouts. The performance-corrected gap in bonus payouts is thus even larger than the raw gap. We find no evidence that women simply receive better ratings because they would cost the firm less in terms of implied bonus payments. We do not find evidence in favor of the interpretation that residual pay gaps are majorly related to child care obligations.

The second set of results answers the question whether the gender gap is different under male and female managers. Intuitively, we do so by defining within-manager gender gaps and comparing their average sizes between male and female managers in a difference-in-differences framework. As we take into account unobserved worker heterogeneity, identification comes from workers working for managers of different genders over time. As the worker fixed effects fully absorb the absolute level of the gender gaps, we identify by how much the expected gender gap *changes* when the manager is male rather than female.

We find that the over-representation of male managers implies a structural disadvantage for women. In particular, male managers cause the gender gap in bonus payouts to increase by 5.1%. This is driven by a relative increase of the gender gap in performance evaluations of 2.7 percentage points comparing male to female bosses. Hence, while in general men receive lower ratings, this gap closes considerably when the manager is male.

We evaluate which mechanism is likely to drive these findings. While more productive men could work more often for male bosses, this mechanism cannot rationalize our findings as we control for unobserved worker characteristics. Another potential explanation are within-gender complementarities. However, gender gaps do not increase with the share of male coworkers in a team. Assuming that potential within-gender complementarities would also exist among coworkers, we can rule out that the productivity channel drives the results. In contrast, we find suggestive evidence consistent with (unconscious) discrimination, as manager gender tends to matter more for less knowledgeable decision makers. In particular, we split managers into groups who should be more or less informed about the true quality or the needs of their subordinates. While not statistically significant, we find that for all proxies of manager knowledge the effect of manager gender is smaller. The observation that a manager's experience, team size, spatial proximity and the time that a manager has worked with a subordinate all are correlated with a lower effect of manager gender is consistent with discrimination due to biased beliefs of less knowledgeable managers. This mechanism relates to Bohren et al. (2019), who show that decision makers resort to their biased beliefs if little information about workers is available to them.

Related Literature. This paper primarily contributes to the vast literature on gender inequality in the labor market summarized among others by Altonji and Blank (1999), or more recently by Bertrand (2011) and Blau and Kahn (2017). One set of papers tries to understand (raw) gender pay gaps. Early papers focused on the role of education and human capital (Altonji and Blank, 1999). As the gender gap in human capital has vanished over time, recent studies have highlighted the role of children as well as differences in occupation and industry. Kleven et al. (2019b) use Danish data to show that the arrival of children creates a substantial long-run gender gap in earnings driven by hours worked, participation, and wage rates (see also Kleven et al., 2019a, for evidence on other countries). Blau and Kahn (2017) document that differences in occupation continue to account for parts of the gender wage gap, and Goldin (2014) finds that work environments rewarding working long hours prevent female wages from fully catching up. Based on an AKM-model (Abowd et al., 1999) in which workers are sorted to firms, Card et al. (2016) find that part of the gender wage gap can be attributed to women working for firms that pay lower premiums. We add to the literature on raw gender wage gaps by showing that the sorting of men and women to different managers in part explains the gender wage gap. To our knowledge, our work is the first to focus on the impact of worker-manager sorting on pay gaps. We bring in a data source—a large manufacturing firm's personnel data—that allows for better control over job characteristics. Research on gender gaps using personnel data dates back to Malkiel and Malkiel (1973) and is vast.[3] However, our data are unique as they identify hierarchical relations between workers and managers while also containing highly detailed information about pay, performance ratings, ranks, and occupations

A second strand of the literature investigates the reasons behind the persistence of the adjusted gender wage gap, in particular with a focus on performance and evaluations. Previous studies come from very specific settings and may therefore lead to different conclusions. In the context of academia—where performance is relatively easy to measure—Sarsons et al. (forthcoming) show that female researchers get less credit for joint work than male co-authors. Card et al. (2020) conclude that journal editors and referees hand out too few revise-and-resubmit decisions to female-authored papers relative to a citation-maximizing benchmark. Similarly, Hospido and Sanz (2019) find that all-female authored papers are less likely to be accepted for major economics conferences. Outside of academia, female surgeons and financial advisors have been found to be more heavily penalized for bad performances or misconduct (Sarsons, 2018; Egan et al., 2017, respec-

[3]For example, Sorensen (1986), Kahn (1992), Ransom and Oaxaca (2005), Barnet-Verzat and Wolff (2008), Dohmen et al. (2008), Ichino and Moretti (2009), Pema and Mehay (2010) or Pekkarinen and Vartiainen (2016) study within-organization gender gaps, mainly based on public sector data.

tively). Mengel et al. (2019) show that female university tutors receive systematically lower teaching evaluations. Azmat and Ferrer (2017) find that gender performance differences exist as male lawyers actually outperform their female colleagues and that accounting for this substantially alters the interpretation of the gender wage gap. What distinguishes our work is that we observe wages, bonus payouts and performance evaluations in the context of a large multinational enterprise in the manufacturing sector, i.e. a setting that is highly relevant to many workers in developed economies. In this context, we find that the over-representation of men in top positions does harm women as gender gaps in bonus payments and performance ratings increase substantially under male managers. In addition, our findings challenge the view that productivity differences can account for adjusted gender wage gaps in a wide range of occupations. We further provide suggestive evidence in favor of the biased-beliefs mechanism proposed by Bohren et al. (2019).

We also add to the literature on the effect of male leadership on gender gaps. While the gender composition at the very top of firms does not affect gender gaps (Bertrand et al., 2019; Maida and Weber, 2019), a number of papers show that gender compositions matter when the distance between superior and subordinate is smaller. Kunze and Miller (2017) and Kurtulus and Tomaskovic-Devey (2011) find that a larger share of women at higher ranks increases women's chances of being promoted. However, it is unclear whether individual interactions between workers and managers or firm-wide policies drive these observations. Our paper differs in that we *directly* link workers to their managers at all levels of the firm hierarchy rather than only at the very top. Using cross-sectional survey data, several authors have documented that gender gaps tend to be greater under male superiors (Ragan and Tremblay, 1988; Rothstein, 1997; Abendroth et al., 2017).[4] In the context of schools, Biasi and Sarsons (2020) find that gender pay gaps among teachers increase when principals or superintendents are male. A recent study by Cullen and Perez-Truglia (2019) shows that the gender promotion gap in a Southeast Asian bank widens when the direct superior is male. Relative to their study, we focus on how manager gender impacts within-team gaps in wages and performance evaluations in a wide range of occupations and countries. Our approach also takes into account unobserved manager characteristics, which prove critical to the finding that managers affect gender pay gaps.

Outline The remainder of the paper is structured as follows. Section 2 provides more information about the data and the firm. Section 3 describes our empirical strategy, Section 4 documents the findings, and Section 5 concludes.

[4]Another cross-sectional study by Halldén et al. (2018), based on Swedish survey data, finds that women earn less when their superior is female. However, the data do not allow to make claims about the gender gap.

2 Data and Setting

The Firm We use the personnel data of a large firm in the manufacturing sector.[5] The firm is among the 250 largest European firms in terms of sales and employment and an industry leader in an R&D-intensive sector. A quarter of the multinational's workforce is located in the firm's home country but it has establishments in over 50 countries. For example, around 20% of the workforce is located in the United States.

While a single firm can hardly be representative of the economy as a whole, its size, international representation, range of occupations, and diversity of skills required ensure that its internal labor market is typical for what a worker would encounter at any large firm. Key variables of our study are earnings and the share of male workers, which we can use to compare the firm to other firms in the same sector. In the US, workers in the same three-digit NAICS industry earned during the sample period around 7.2% less than workers in our data.[6] At the firm, the share of female workers in the US is about 6 percentage points higher than the sectoral average.[7]

By now, administrative data matching employees to firms are widely available. While such kind of data would be preferable to use for a more holistic view, our dataset has several unique features. As opposed to administrative data where workers would be linked by working in the same firm and same occupation, we observe precisely who works in the same team and who is each worker's responsible superior. Furthermore, the description of jobs and hierarchy provided to us by the firm goes beyond typical definitions of occupations. This allows us to control much more precisely for the nature of the job. We also do not only observe earnings but the detailed variables determining compensation, including performance ratings.

Personnel Data We were provided with an anonymized monthly panel of all personnel records between January 2014 and March 2019. The data includes information on employees' compensation (base salaries and bonus payments), performance ratings, occupation, hierarchical rank, location, tenure, and some sociodemographic characteristics such as gender, age, or nationality.[8] Importantly, the data also indicate the identity of each worker's superior, to which we refer as *manager* or *boss*. Therefore, we can trace out the organizational hierarchy and identify employees' coworkers, superiors, and subordinates. Employees can be superiors and subordinates at the same time as the firm has many hierarchical layers.

[5]We are not allowed to reveal the identity of the firm.

[6]Bureau of Economic Analysis, Wages and Salaries Per Full-Time Equivalent Employee by Industry, `https://apps.bea.gov/iTable/iTable.cfm?reqid=13&step=3&isuri=1&nipa_table_list=201&keyword_index=w`

[7]Bureau of Labor Statistics, Labor Force Statistics from the Current Population Survey, `https://www.bls.gov/cps/cpsaat18.htm`

[8]The number of employees reporting education and previous employers is very small.

Table 1: Gender Distribution across Hierarchical Ranks

Rank	Share [%]	Salary [€]	Bonus [€]	Any Bonus [%]	Share Male [%]	Salary Gap [€]	Bonus Gap [€]
1	0.5	18 002.58	13 036.89	33.67	63.85	−5324.59	5792.14
2	7.65	27 131.86	4007.53	53.26	66.3	5958.01	2095.94
3	14.76	32 181.96	2479.71	54.02	67.62	2395.46	454.9
4	25.54	32 771.42	3002.1	47.4	57.1	−860.15	−162.59
5	27.79	44 937.26	5931.84	44.28	57.01	654.7	138.78
6	17.25	76 957.14	14 907.22	71.98	61.8	3002.29	131.61
7	3.71	117 899.9	31 514.63	80.15	68.5	−8527.77	−135.85
8	2.41	154 211.52	57 869.84	81.39	75.94	−5702.43	−1890.47
9	0.34	227 647.17	137 310.0	74.71	87.77	−16 322.27	−24 453.82
10	0.05	335 444.2	276 739.82	69.79	72.92	59 631.59	17 396.89
All	100.0	50 090.26	10 781.41	54.27	61.16	3917.02	2374.62

Notes: The table displays mean values of base salaries and bonus payouts in € with base year 2010 for nine out of ten hierarchical ranks (1 being lowest and 10 highest). The table also displays the share of workers, the share of men, and the gender gaps (male outcome - female outcome) in base salaries and bonus payouts.

We observe ten thousands of full-time employees.[9] The unadjusted gender gap (the average difference between male and female outcomes) in base salaries (before taxes) is 3917€, or 7.6%. One would typically control for experience, location, and job characteristics to determine the adjusted gender pay gap. Such an adjustment yields a gender gap in base salaries of 1.7%. 26 404 workers are paid a bonus. In this set of workers, the gender gap in annual bonus payout (before taxes) is 2375€, or 20.3%. Controlling for experience, location, and job characteristics results in an adjusted bonus gap of 4.6%.

Job Definitions Jobs are classified into 109 different occupations, for example *Electrical Engineering*, *Scientific Technical Assistance*, *Fire Brigade*, or *Web Design*.[10] In addition, jobs are classified by hierarchical rank on a scale from one to ten, representing a range from low-skilled helpers to executives. Ranks one to three are to a large extent blue-collar production jobs. While workers of different hierarchical rank can work in the same occupation, occupations only comprise a limited number of ranks. Starting at rank four, white-collar occupations are more prevalent. We henceforth characterize a job by the combination of occupation and hierarchical rank.

From Table 1 we can see that the majority of workers works on jobs of intermediate hierarchical rank. In general, higher ranks pay higher salaries and bonuses, but not all workers receive bonuses. While participating in a bonus program seems to be more common at high ranks, a number of workers opts out. At very high ranks, other long term incentives also play a greater role. The share of men at low ranks is relatively large as many production-related occupations are ranked here. Women represent around 40% of the workforce in the middle range of the hierarchy. But the higher the rank, the lower becomes the share of women. While there is no clear pattern of raw gender pay gaps at

[9] We cannot report the exact number to protect the identity of the firm.
[10] We cannot show a full list of occupations as it could reveal the identity of the firm.

separate hierarchical levels, it seems that men face an advantage at the most common ranks in the firm.

Salaries, Bonus Payments and Performance Ratings According to the firm, conditional on individual productivity, experience, and location, employees working in the same job should be compensated equivalently. It might well be the case that a certain occupation pays more at a lower rank than another occupation at a higher rank. Each job has a salary band set by the human resource department which is only known by the worker's manager and not part of the data we obtain. Managers and workers negotiate a base salary within such a band. Usually, pay raises within the same job are only negotiated once a year, namely when an employee's performance is evaluated.

The contract also specifies a bonus target, which is the amount that will be paid out in addition to the base salary as an annual bonus. The target is expressed as a percentage of the base salary. For example, if the employee's target is 5% and the base salary is 50 000€, she can expect annual earnings of 50 000€ + 2500€ = 52 500€.

However, the bonus is supposed to incentivize effort. Workers showing satisfactory but not outstanding performance are paid as just described. Less is paid out as a bonus if the employee performs poorly, and more is paid out if the employee does especially well. This means that workers with a higher base salary, higher bonus target, or higher performance will receive a larger annual bonus. Workers' effort has a significant influence on the amount eventually paid out as a bonus. Workers know the function $f(\text{rating})$, which maps their grading into a factor multiplying base salary and bonus target. In principle, earnings can be expressed as the sum of base salary and the performance-dependent component, i.e. the product of base salary, target and a function of performance.

$$\text{Earnings} \approx \text{Base Salary} + \text{Base Salary} \times \text{Target} \times f(\text{rating})$$

Performance ratings are handed out by the employee's direct superior once per year, evaluating the previous twelve months. An evaluation scheme of six grades is applied across all jobs and countries. The firm considers a ranking to be high if an employee achieves at least the second-best grade. A *high* ranking will ceteris paribus result in a bonus payout greater than the contracted target. As an employee's output is hard to measure, the mapping from effort to performance ratings cannot be contracted.[11] Hence, performance ratings are potentially subject to conscious or unconscious gender biases of the manager (e.g. Bordalo et al., 2019). Through their performance evaluations, managers thus have a substantial impact on the total earnings of their subordinates.

[11] Employees for which individual output can easily be measured are sales workers. They receive a separate sales bonus in addition to the general bonus which reflects the generated revenue. We do not have access to these data.

FIGURE 1: Distribution of Ratings by Gender

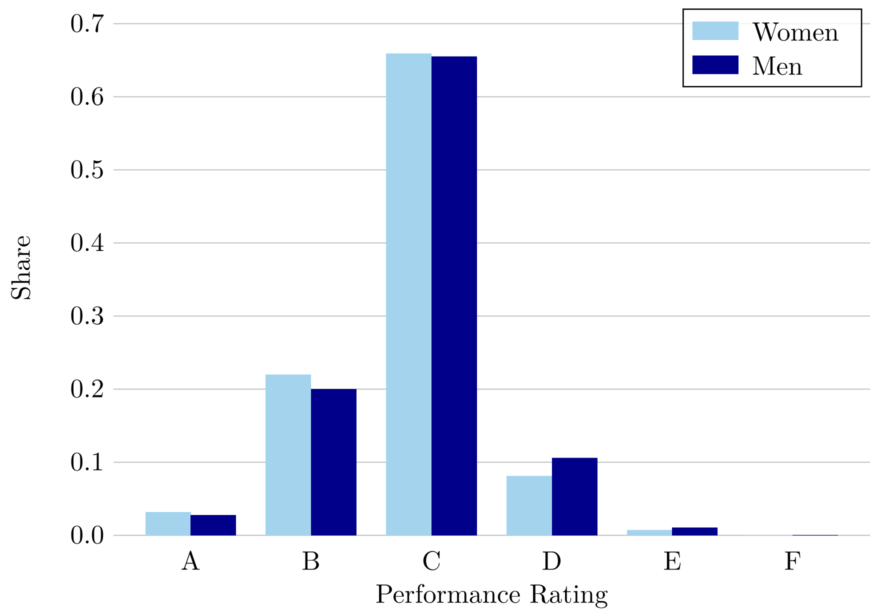

Note: The figure plots a histogram of annual performance ratings received by male and female workers. The best rating is A, the worst rating is F.

Figure 1 displays the distribution of performance ratings for men and women. The majority of workers receives a C, which implies no adjustment to the contracted bonus. A high performance rating is achieved by workers receiving ratings A or B. The graph also shows that in the raw data women are more likely than men to receive a high rating.

Observed bonuses are not equivalent to Base Salary × Target × f(rating). There are several reasons for this. First, department-wide achievements also affect bonus payout. Second, we are not provided with the mapping for workers with changes in workplace characteristics. Third, performance ratings are very detailed, while we only use a simple approximation, an indicator for receiving a high rating. Importantly, it holds that all else equal, workers with high performance should receive a larger bonus.

Sample Description We focus on full-time employees aged 25 to 60 and for which we observe gender, superior, job, age, and tenure. We aggregate the data to annual frequency for 2014-2019 as bonus payments and performance ratings are only determined once per year. Monetary variables are converted to Euros with 2010 as the base year. The resulting dataset is summarized in Table 2.[12]

Base salaries are on average 50 000€. We observe positive bonus payouts only for a subset of workers because not every worker receives performance pay and due to data limitations. The size of bonus payments is significant with a mean annual payout of

[12] Extreme values are omitted for confidentiality reasons.

TABLE 2: Descriptive Statistics

	Mean	SD	10th Perc.	Median	90th Perc.	Observations
Salary [€]	50 090.26	53 096.49	10 077.67	40 404.73	94 234.79	178 377
Bonus [€]	10 781.41	20 324.33	1289.45	4233.43	24 812.78	96 803
High Performance	0.22	0.42	0.0	0.0	1.0	154 790
Low Performance	0.13	0.33	0.0	0.0	1.0	154 790
Bonus Target [%]	11.82	8.81	4.0	10.0	22.51	108 839
Male	0.61	0.49	0.0	1.0	1.0	178 377
Age	41.53	9.24	29.0	41.0	55.0	178 377
Tenure	9.82	9.19	1.0	7.0	24.0	178 377
Span of Control	1.16	3.77	0.0	0.0	5.0	178 377
Coworkers	9.65	20.59	1.0	5.0	19.0	178 377
New Manager in Same Job	0.25	0.44	0.0	0.0	1.0	128 521
New Manager in New Job	0.07	0.25	0.0	0.0	0.0	128 521
Male Manager	0.72	0.45	0.0	1.0	1.0	178 377
Male & Male Manager	0.49	0.5	0.0	0.0	1.0	178 377
Male & Female Manager	0.12	0.33	0.0	0.0	1.0	178 377
Female & Male Manager	0.23	0.42	0.0	0.0	1.0	178 377
Female & Female Manager	0.16	0.36	0.0	0.0	1.0	178 377
Age of Manager	44.97	8.08	34.0	45.0	56.0	178 377

Notes: Extreme values omitted for confidentiality reasons. Unbalanced panel based on ten thousands of workers and the years 2014-2019. Monetary variables normalized to € in 2010.

almost 11 000€. Approximately 20% of workers receive a high performance rating, so are entitled to a bonus payout greater above their contracted bonus target. 13% of workers underperform, implying a below-target bonus payout. Some workers receive a rating even though their contracts do not include performance pay. 60% of employees are male and the average age is 42 years. The span of control measures the number of direct subordinates.

Identifying each worker's manager is key to studying managers' impact on gender gap. We do so by matching each worker to his or her direct superior. Each employee has on average 9.7 coworkers who work under the same manager. Employees are more likely to have a change in their manager due to managers rotating than due to a job change of the worker. One quarter of employees stay in their current job but work for a new manager. Managers are more likely to be male and on average three years older than workers.

The data reveal that male and female workers are sorted to different managers based on gender. The share of male workers sorted to a male manager is 80%. Women are more likely to work for female managers as the share of female workers sorted to a male manager is only 59%.

3 Empirical Strategy

In the first part of the empirical analysis, we examine how the sorting of workers to managers can explain gender gaps. Evidence for the impact of managers on worker productivity has been provided by Frederiksen et al. (2019) or Lazear et al. (2015). For example, women may work for managers who are less productive. This could happen if male workers have stronger social networks within the firm which provide them with better information about managers. In Table 2 we saw that workers are sorted to managers based on gender. If female managers are less productive on average (in line with the gender gaps in performance found by Azmat and Ferrer, 2017) and manager productivity affects worker productivity, female workers are disadvantaged.

There are other explanations why managers could matter. Male employees are often less reluctant to negotiate (e.g. Babcock et al., 2003). This could drive them towards more generous managers who are open to negotiation. Similarly, the fact that women tend to shy away from competition (Niederle and Vesterlund, 2007) could drive female workers to work more often for managers who create less competitive work environments, with the effect that workers are on average less productive. Women might be also driven to managers who—at the cost of lost productivity—offer a more family friendly environment, for example by allowing for more flexible work hours or permitting working from home on a regular basis (Goldin and Katz, 2015). This explanation would imply that women might actually prefer to work for "worse" bosses, i.e. bosses creating a environment which makes workers less productive.

In the second part of the analysis, we hold these manager effects fixed and study whether there is evidence that managers affect their *within-team pay* gaps. Observing any residual gender gap within teams does not necessarily imply that managers are to blame for gender pay gaps. But if within-team gaps vary across managers with different characteristics, we can conclude that managers do affect workplace equity.

3.1 Explaining Gender Gaps

We aim to answer the question what portion of gender gaps can be explained by observable characteristics. In particular, we want to quantify the contribution of the matching of managers and workers to gender gaps. To do so, we implement traditional Kitagawa-Oaxaca-Blinder-decompositions of differences between male and female workers in log salaries, log bonus payouts, high performance indicators, and contracted bonus targets (Kitagawa, 1955; Oaxaca, 1973; Blinder, 1973).

The decomposition classifies differences between two groups into a composition component that accounts for different characteristics, e.g. tenure or occupation, and an unexplained, or wage structure, component. Such Kitagawa-Oaxaca-Blinder-decompositions are commonly used in the estimation of gender gaps (e.g. Bertrand et al., 2010; Blau and Kahn, 2017; Card et al., 2016; Juhn and McCue, 2017).

The unexplained component is often interpreted as a measure of discrimination, as it implies differences in pay or other outcomes for observationally identical workers. However, discrimination might also stem from different characteristics, i.e. the explained part of the gender gap. Women might be discriminated against by being allocated to worse-paying occupations. By controlling for the type of occupation, we take this allocation as given while it could already be the result of discriminatory treatment. It is also not clear that all unexplained differences are the result of discrimination. If men perform better than women (as found by Azmat and Ferrer, 2017), productivity differences contribute to the residual term.

The Kitagawa-Oaxaca-Blinder-decomposition is implemented as follows. Each worker i is either male m or female f and observed in year t. We estimate ordinary least squares (OLS) regressions for both genders separately.

$$Y_{it}^{m} = X_{it}^{m}\beta^{m} + u_{it}^{m} \tag{1}$$

$$Y_{it}^{f} = X_{it}^{f}\beta^{f} + u_{it}^{f} \tag{2}$$

Y_{it} is the outcome of interest of worker i in year t. X_{it} is a vector of variables observed at the worker level, including a constant. β^{m} and β^{f} are the gender-specific returns to these characteristics. u_{it} is the error term.

We obtain the estimates of returns $\widehat{\beta}^{m}$ and $\widehat{\beta}^{f}$ from OLS. We then calculate means of the outcomes and characteristics for both genders, denoted by a bar over the respective term. The difference of the means of equations (1) and (2) is the observed gender gap.

$$\overline{Y}^{m} - \overline{Y}^{f} = \overline{X}^{m}\widehat{\beta}^{m} - \overline{X}^{f}\widehat{\beta}^{f} = (\overline{X}^{m} - \overline{X}^{f})\widehat{\beta}^{m} + \overline{X}^{f}(\widehat{\beta}^{m} - \widehat{\beta}^{f}) \tag{3}$$

The residuals drop out when taking the mean. The first term of the decomposition in equation (3) is the difference in male and female outcomes due to different characteristics, based on male coefficients. The second term is the unexplained difference in outcomes due to different returns for men and women. One could perform the decomposition using the coefficients on female returns as well. However, here we are interested in how the outcomes of women would change if the firm is required to treat women like men.

We decompose log base salaries, log bonus payouts, high performance indicators, and log bonus targets. Due to the richness of the firm's personnel data, we can include a wide range of variables in X_{it}. For age and tenure, we create bins for every six years of age. We include indicators for each year and each country. Job characteristics are controlled for by including the combination of occupation and hierarchical rank. As we are particularly interested in the role of bosses, X_{it} also includes an indicator for worker i's boss at time t.

The high resolution of worker-level controls comes at a cost. The matrices X^{m} and X^{f}, which respectively collect the vectors $X_{i,t}^{m}$ and $X_{i,t}^{f}$, need to have full rank. This is not the case if, for example, there is a certain job which is only done by men. In such a case the column of the matrix X^{f} indicating working for this manager would always be zero and hence perfectly correlated with the column indicating the constant. We therefore only include observations with a characteristic observed among men and women.

Another requirement for full rank matrices is that the different categorical variables are connected. This is identical to the condition explained by Abowd et al. (2002), that fixed effects in an AKM-model (Abowd, Kramarz, and Margolis, 1999) are only identified within a connected set. A set of observations is unconnected if a categorical variable is nested in another categorical variable. Consider an example where we include a constant, occupation, and location. Assume that engineers and accountants always work in Spain, cleaners always work in France, and there are no other occupations or countries. When comparing a worker in France to a worker in Spain it is unclear whether their pay difference is due to location or occupation. But if accountants work in both countries, they identify differences due to location. Once these differences are determined, residual differences in pay can be attributed to occupations. Note that even in this case we still require a linear restriction due to the inclusion of a constant. For this reason, we only keep the largest connected set of categorical variables for men and women. In practice, we iterate over the algorithm proposed by Abowd et al. (2002) for different combinations of categorical variables until the matrices X^m and X^f are full rank and only include observations with a characteristic observed both among men and women.

Assuring that matrices are full rank reduces the sample size. To maintain a constant size of data, we also impose that workers participate in bonus schemes and that bonus targets and performance ratings are observed. The resulting dataset consists of 59 813 observations based on 20 048 workers and 4327 managers from 58 countries and 421 jobs. Appendix Table A.1 shows that in this dataset workers earn a bit more which is due to the fact that some low-rank jobs without performance pay are excluded. Besides that, the subset of data used for estimation is very similar.

3.2 Understanding Residual Gender Gaps

Residual gender gaps could be interpreted as productivity differentials between men or women. In the decomposition, managers can contribute to gender gaps through the channel that male and female workers work for different managers. But it does not tell us anything about how individual managers treat men and women who actually work for them. For example, these residual gaps could persist if the majority of bosses for some reason favor male workers. At the firm, 72% of managers are male.

Illustration of Within-Manager-Gender Gender Gaps Figure 2a is an illustration of how gender gaps could look like, separating gaps at male and female managers. The light dots indicate the earnings of women, the dark dots earnings of men. The filled dots on the left plot earnings at female managers, the filled dots on the right at male managers. The labels MF, FF, etc. also refer to the four possible combinations of worker and manager gender. For example, MF stands for male workers working for female managers. Focusing on female managers only, the difference between MF and FF is the gender pay gap. Workers at male bosses earn more ($FM - FF$), but the size of the gender gap

$MM - FM$ is equal to the gap under female managers. The gender pay gap $MF - FF$ is a quantity of interest, but it is unclear why this difference exists.

In the example from Figure 2a, all managers could be discriminating against women, or women could be less productive. In such a setting we cannot draw any conclusions on the impact of managers on gender gaps.

But if we find that there is variation in gender gaps across teams managed by managers of different gender, we know that manager gender and hence managers affect gender gaps. This is depicted in Figure 2b. In this example, the gender gap at male bosses is greater than the gender gap at female bosses. If the difference in gender gaps, ω, is significant, we can conclude that managers affect gender gaps.

Estimation To estimate the change in the gender gap depicted in Figure 2b, we run a difference-in-difference estimation for outcome Y of the following form.

$$Y_{it} = \gamma_0 + \gamma_1 \times male_i + \gamma_2 \times male_{M(i,t)} + \omega \times male_i \times male_{M(i,t)} + X_{it}\beta + \epsilon_{it} \tag{4}$$

$male_i$ is a dummy taking value 1 if worker i is male. Its coefficient represents the gender gap under female managers, i.e. $MF - FF$ in Figure 2b. $male_{M(i,t)}$ is a dummy taking value 1 if the manager M of worker i at time t is male. Its coefficient can be interpreted as the difference in earnings among women when working for a male instead of a female manager, i.e. $FM - FF$ in Figure 2b. $X_{i,t}$ controls for age-bin, tenure-bin, year, country, and job as before. We are interested in the difference-in-differences coefficient ω on the interaction of the dummy $male_i$ with the dummy $male_{M(i,t)}$. This product is 1 if worker and manager are male and 0 otherwise. As shown in Figure 2b, ω is the difference in gender gaps between male and female managers. A positive ω indicates that the gender gap moves in favor of men when the manager is male.

Replacing all gender dummies by $female_i = 1 - male_i$ etc., results in exactly the same estimate of ω, but now on the product of dummies $female_i \times female_{M(i,t)}$.[13] The natural interpretation would be as follows. The inverse gender gap (female outcome - male outcome) increases if the manager is female. This shows that we only can quantify by how much the gender gap changes, but not whether male or female managers are to blame.

A major problem when estimating Equation (4) can be unobserved heterogeneity of workers and managers. If workers are sorted to managers based on these unobserved characteristics, the estimate of ω is biased. For example, if good male workers tend to work

[13] This can be shown by a simple replacement of variables:

$$\begin{aligned}
&\gamma_0 + \gamma_1\ male_i + \gamma_2\ male_{M(i,t)} + \omega\ male_i \times male_{M(i,t)} + \dots \\
&\quad = \gamma_0 + \gamma_1\ (1 - female_i) + \gamma_2\ (1 - female_{M(i,t)}) + \omega\ (1 - female_i) \times (1 - female_{M(i,t)}) + \dots \\
&= (\gamma_0 + \gamma_1 + \gamma_2 + \omega) - (\gamma_1 + \omega)\ female_i - (\gamma_2 + \omega)\ female_{M(i,t)} + \omega\ female_i \times female_{M(i,t)} + \dots
\end{aligned}$$

for good male managers, or good female workers often work for good female managers, identification is compromised. Therefore, we introduce in Equation (5) a worker fixed effect, α_i, and a manager fixed effect, $\Psi_{M(i,t)}$, for i's manager M at time t.

$$Y_{it} = \omega \times male_i \times male_{M(i,t)} + \alpha_i + \Psi_{M(i,t)} + X_{it}\beta + \epsilon_{it} \tag{5}$$

Estimating ω in Equation (5) yields the change in within-manager pay gaps, adjusted for worker quality, when the manager is male instead of female. Intuitively, this does the following. We residualize each worker's outcome based on the controls X_{it} and remove the worker mean. Then, we calculate the within-manager pay gap based on all workers who worked for the manager. Again, the level of the gap has no interpretation, because we subtract the mean from each worker's outcome. ω is the mean difference of the adjusted within-manager pay gap between male and female managers.

Identification Identification of ω in Equation (5) comes from "movers", i.e. workers who work for different managers. There are two reasons why a worker experiences a manager change. First, workers who switch their job will face a new manager.[14] Second, workers who do not change positions receive a new manager if the previous manager rotates to another job or leaves the firm. In the present setting, 32% of workers observe a change in manager in a given year, and almost 80% of these switches are due to managers rotating jobs.

We require an exogeneity assumption regarding the changes of managers. Sorting of workers to managers based on time-varying performance would bias estimations of equation (5). The fact that the majority of changes is due to managers rotating limits this concern because managers would need to be assigned based on the potential future performance of workers. From our interviews with the human resources department of the firm this is extremely unlikely. Usually, managers only get to know their subordinates after having started a new position. However, one might still worry that moves of workers correlate with time-varying performance or that workers only move if they get a better deal. While one cannot prove exogeneity, we examine whether there is evidence that sorting could bias the estimation results. As suggested by Card et al. (2013) in a setting where workers move across firms, we implement event studies on pay for workers changing managers.

To do so, we classify managers by the average of salaries or bonuses. To account for job characteristics, we first residualize these outcomes, taking into account job characteristics, location, tenure, age, and year. Then, we calculate the leave-one-out mean of the residual to avoid selection based on worker i's own productivity This is the average residual observed under i's current manager $M(i,t)$, excluding the contribution of i. The

[14]Excluding the possibility that boss and worker jointly switch teams.

leave-one-out means are then used to classify each worker-year observation into one of four quartiles. Next, we calculate the average residuals of workers two years before and two years after changing managers. We do so for salaries and bonuses of workers in all 16 possible transitions, e.g. workers changing from a category 4 to a category 1 manager, workers changing from a category 4 to a category 3 manager, etc. For clarity, we focus on workers who previously worked for category 4 or category 1 managers. Note that we include changes in manager due to job changes of workers and due to manager rotation.

Figure 3 shows that the groups have different pay levels before (years -2 and -1) and after changing managers (years 0 and 1). For example, salaries of workers with coworkers in the fourth quartile who move to a quartile 1 manager have lower salaries prior to a change compared to workers who change from quartile 4 to another quartile 4 manager. Moving to a manager with higher-paid coworkers, e.g. from quartile 1 to quartile 4, increases pay. Workers who stay in the same quartile have relatively constant pay, although bonus pay seems to increase quite a bit for workers switching from a quartile 1 to another quartile 1 manager. Workers who change from a quartile 4 manager to a lower quartile manager lose pay, with larger losses for more extreme changes.

Pay changes in Figure 3 look symmetric for workers moving between quartile 1 and quartile 4 managers. This suggests that a simple additive model is a reasonable approximation of base salaries and bonus payments. It implies that workers do not only change managers if higher residual pay is expected.

The pay profiles in Figure 3 also look relatively flat before and after changing manager. While there is some variation in pre-change pay, for example among workers moving from bonus quartile 1 to 4, these changes are small compared to the jumps we observe. This suggests that a static model as in equation (5) should be a sufficient approximation.

Figure 2: Illustration of Gender Gaps

(a) Equal Gender Gaps at Male and Female Managers

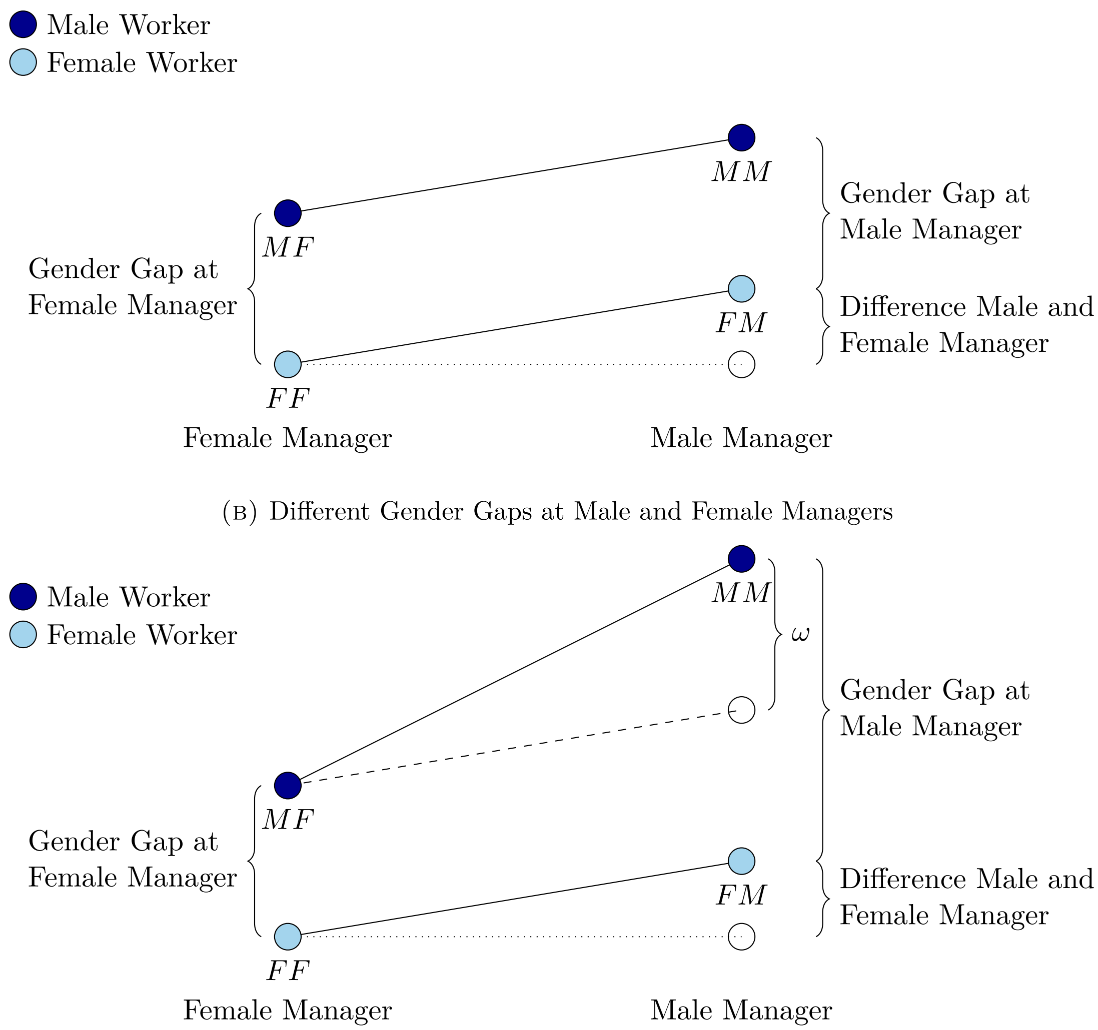

(b) Different Gender Gaps at Male and Female Managers

Notes: These figure provides a graphical illustration of gender gaps under male and female managers. FF stands for a female worker working for a female manager, MF for a male worker working for a female manager, FM for a female worker working for a male manager, and MM for a male worker working for a male manager. ω is the difference in gender gaps between male and female managers, i.e. $\omega = (MM - FM) - (MF - FF)$.

FIGURE 3: Mean Log Pay of Manager Changers by Quartiles of Mean Coworker Pay

(A) Salary (B) Bonus

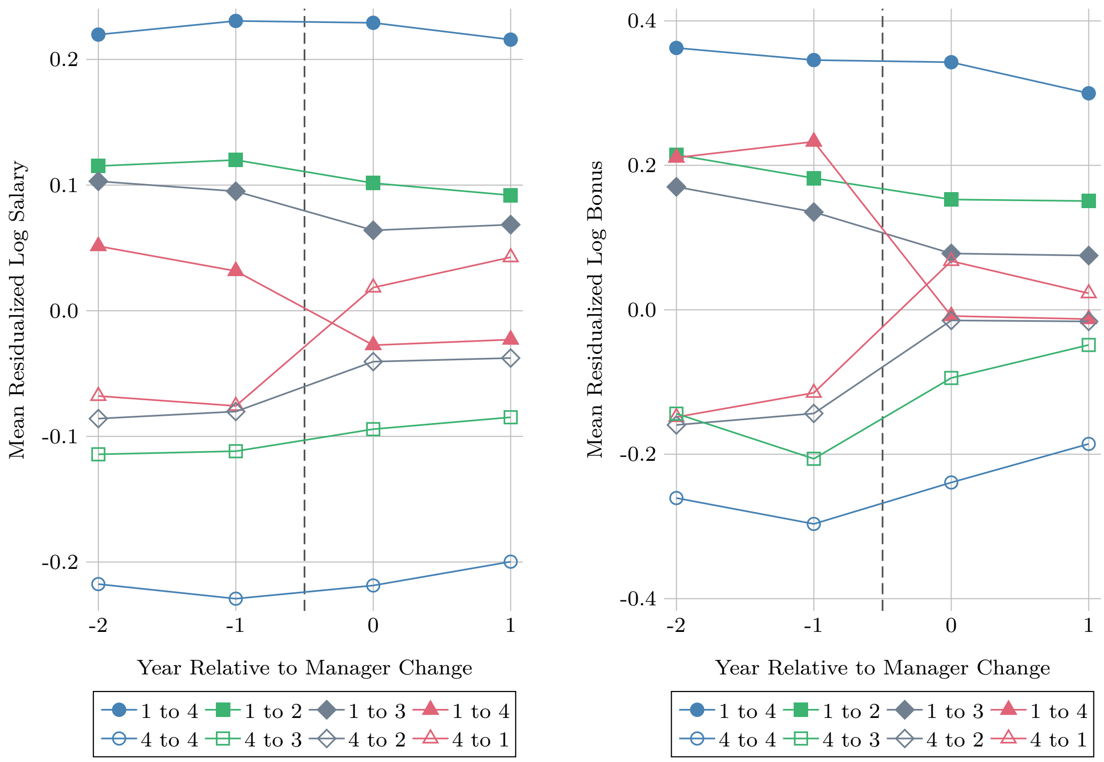

Notes: These figures plot average residual base salaries and bonus payments in the two years before and after manager changes. Based on Card et al. (2013), workers are classified into 16 groups, of which eight are displayed. Workers are grouped by the quartile of their coworkers' pay before and after the manager change. The label "1 to 4" denotes workers whose coworkers' average salary was in quartile 1 before and in quartile 4 after the change of managers.

4 Results

4.1 Decompositions

In the first set of results, we look at the impact of men and women doing different jobs and in particular working for different managers. The decomposition of gender pay gaps requires that male and female workers identify the same set of characteristics and full-rank matrices for each gender. Therefore, we reduce the data to a dual-connected set, as described in Section 3. We also impose that for a given observation we observe base salary, bonus payout, performance ratings, and targets. While this excludes workers who do not have performance pay in their contract, this has the advantage that different results for the outcomes are not driven by sample composition.

Table 3 displays the gender pay gap decomposition based on 59 813 observations. The raw gender pay gap in salary is 12 log points (13%). The raw gap in bonuses is even larger, with a difference between men's and women's payouts of 22 log points (24%). While the raw pay gaps are large, between 80 and 90% can be explained by different observed characteristics of men and women. The residual gender gap is 1.1 log points for base salaries and 3.8 log points for bonuses.

The Kitagawa-Oaxaca-Blinder decomposition in Table 3 shows the relative importance of the different characteristics explaining the gap. Age and tenure differences between men and women exist but are of small magnitude. Also, the fact that the gender distribution of the workforce might not be uniform over countries is not of primary importance. Job characteristics, i.e. the combination of hierarchical rank and occupation, explain between 55 and 60% of the gender gaps in salary and bonus. This means that if women worked in the same occupations and hierarchical levels as men, and earned the same returns from these jobs as men, the gender pay gap in salaries would be reduced by more than half.

The sorting of workers to managers matters. 25% of the gender pay gap in base pay and 19% of the gap in bonuses can be explained by the fact that men work for managers who have more positive impacts on pay. This is true conditional on workers doing the same job, i.e. working in the same occupation within the same hierarchical rank. While we cannot tell why women tend to work for bosses who are less generous or have less favorable impacts on productivity, we can see that managers are not perfectly substitutable but differ in terms of productivity or styles. This also means that firms seeking to reduce gender pay gaps need to carefully consider if and how to better match women to high-impact managers. Should women prefer to work for managers whose style implies a productivity reduction, for example due to flexible work hours, female workers might in spite of a financial cost prefer the current allocation to managers (Goldin and Katz, 2015).

The Role of Child Care A large part of gender pay gaps has been attributed to reduced working hours due to child care obligations (Kleven et al., 2019a,b). We neither

Table 3: Gender Gap Decomposition

	Salary		Bonus Payout	
	Log points	Share explained	Log points	Share explained
Age	0.75	6.1%	0.93	4.2%
Tenure	0.16	1.3%	0.25	1.1%
Manager	3.13	25.4%	4.16	18.7%
Job	6.59	53.5%	13.57	61.1%
Year	-0.23	-1.8%	-0.83	-3.7%
Country	0.84	6.8%	0.38	1.7%
Total explained	11.24	91.2%	18.45	83.1%
Total unexplained	1.08	8.8%	3.75	16.9%
Total gap	12.32	100.0%	22.20	100.0%

Notes: The table displays Kitagawa-Oaxaca-Blinder-decompositions of the gender gap in Log Salary and Log Bonus Pay based on 59 813 observations. Age and tenure are summarized in six-year bins. Jobs are the combination of hierarchical rank and occupation.

observe family status nor overtime work. By limiting the data to full-time workers only and controlling granularly for the nature of the job, much of the variation in actual hours worked is already taken into account. However, to examine the role of child care in our setting, we expand the decompositions in several dimensions.

First, we repeat the pay decomposition for workers aged 44 and younger and workers aged 45 and older. Here one would expect that older workers are less likely to be affected by small kids at home. While child birth may have a lasting effect on careers of older workers as well, our approach already takes into account that parents might climb up the hierarchy more slowly. Appendix Table A.2 shows that the adjusted salary gap among younger workers is actually zero, and smaller than among older workers. The adjusted bonus gap is around 5% for both groups. This suggests that differences in unobserved actual hours worked do not play a role for the gender gap in this sample. As the sample is limited to full-time workers only and due to the exact controlling for job fixed effects, most differences in work hours are probably already accounted for.

To examine this further, we recalculate pay decompositions while including all workers in the data and controlling for contracted working hours. As Appendix Table A.3 shows, a significant part of the gender pay gap can be explained by differences in working hours. 34% of the pay gap in base salary of 13.7 log points is attributed to this channel. Unexplained gender pay gaps are larger in this full sample, relative to full-time workers only. This finding suggests that gender differences in working hours do matter, but are largely taken into account already by limiting observations to full-time employees.

Full-time workers might differ in their accumulated working hours. In particular, workers with children might have worked fewer hours in the past. Workers can also have spells during which they did not work at all, for example because of child birth. We

Table 4: Gender Gap Decomposition

	High Performance		Bonus Target	
	Percentage points	Share explained	Log points	Share explained
Age	-0.93	46.7%	0.27	9.5%
Tenure	-0.00	0.0%	-0.13	-4.6%
Manager	2.12	-106.0%	-3.44	-121.9%
Job	0.71	-35.3%	6.61	234.6%
Year	-0.06	2.9%	-0.08	-2.8%
Country	-0.30	14.9%	-2.44	-86.7%
Total explained	1.54	-76.8%	0.79	28.1%
Total unexplained	-3.54	176.8%	2.03	71.9%
Total gap	-2.00	100.0%	2.82	100.0%

Notes: The table displays Kitagawa-Oaxaca-Blinder-decompositions of the gender gap in the probability of a high performance rating and contracted Log Bonus Targets based on 59 813 observations. Age and tenure are summarized in six-year bins. Jobs are the combination of hierarchical rank and occupation.

treat these spells as having worked zero hours. Having worked part-time in the past could be interpreted as a measure of experience, less flexibility, or reduced likelihood of working overtime because of child care obligations. Appendix Table A.4 decomposes full-time workers' pay while controlling in addition for accumulated full-time equivalent months. As we need to observe every worker's full employment history at the firm, the sample size is reduced. The unexplained gender pay gap is similar to our baseline estimate from Table 3. Differences in accumulated full-time equivalent months do not contribute to gender gaps. Once more, this could be because workers sort into jobs with different requirements of flexibility or because the firm requires certain experience for working in particular positions.

These extensions demonstrate that it is unlikely that the unexplained part of the gender pay gap documented in Table 3 can be attributed to child care obligations.

Sources of Gender Bonus Pay Gaps Bonuses depend on base salary, performance ratings, and contracted bonus targets. We have seen that a gap in base salaries exists and can be explained by sorting to jobs and managers. Here we study whether sorting with regard to contracted bonus target or performance rating matters as well.

Table 4 reports that the raw gender gap in performance is -2 percentage points. The penultimate line shows that the adjusted gender gap is even more negative (-3.5 percentage points). This is mainly driven by the fact that women tend to work for managers handing out worse ratings. This means that if women had the same characteristics as men and were to earn the same returns on these characteristics as men, their performance ratings would be even *better*.

It is a striking result that gender performance gaps are negative, i.e. favor women, while the unexplained gap in salaries and bonus payouts is positive. If performance ratings are a good proxy for actual performance, these residual gaps are difficult to reconcile with the explanation that men are more productive. Previous work by Azmat and Ferrer (2017) shows that female lawyers perform worse than their male colleagues. Our results demonstrate that the use of subjective ratings can make much of a difference when comparing the performance of male and female workers.

How can it be that women earn lower bonuses, despite the fact that their ratings are better? Bonuses depend on base salaries and contracted bonus targets. As these gaps benefit men—the unexplained gap in targets is 2.0 log points and the unexplained gap in base salaries is 1.1 log points—the eventual payout still favors men.

Women do not sort to more generous or high-impact managers as can be seen from the detailed results of the decomposition in Table 4. Instead, men work for managers who hand out better ratings. If women worked for the same managers and benefited from them in the same way as men, the probability to receive a high performance rating would go up by two percentage points.

Table 4 also indicates that women are sorted to jobs with significantly lower bonus targets. However, women actually tend to work for managers who negotiate higher bonus targets. This means that while women might shy away from competition (as found by Niederle and Vesterlund, 2007), it does not imply that they work for managers where performance pay plays a smaller role.

Are Women's Better Ratings Explained by Lower Costs? One explanation for negative performance gaps is that managers might be more willing to hand out good ratings if they are less costly. We observed that women earn lower salaries and negotiate lower bonus targets. Managers could therefore use performance ratings to compensate women for negotiating lower salaries and targets.

To check this, we once more decompose the gender gaps in performance ratings, taking salary and target as predetermined variables and including them in the controls. We find in Appendix Table A.5 that the unexplained gender gap opens even more in favor of women. This suggests that women do not simply receive better ratings because they cost the manager less.

Controlling for Additional Workplace Characteristics We control for job characteristics by including a set of indicators for each combination of occupation and hierarchical rank. However, it could also be the case that male and female workers who work in the same job are located in different units of the firm. For example, a software engineer could work in the compensation unit of the human resources department or in the financial analysis team – and these different workplaces could provide very different pay packages.

This means that by not taking into account the exact unit in which a worker is located, we could attribute such differences to the managers and thereby overestimate the role that the sorting of managers and workers plays for the gender pay gap.

In Appendix Table A.6, we repeat the pay gap decomposition, but now including indicators for the unit in which a worker is employed. Each worker is matched to one of more than 200 units. The results show that a differential sorting of men and women can explain only very little of the gender pay gap. If anything, women tend to work in teams with slightly higher bonus pay. While not substantially changing our previous findings, this exercise further reduces the sample size to 50 002 observations because finding a connected set in all categorical variables becomes more challenging.

4.2 Within-Manager Gender Gaps

Having documented that the sorting of women to managers contributes to gender pay gaps, we now ask whether managers affect residual gender pay gaps.

Based on the same data used in the Kitagawa-Oaxaca-Blinder decomposition we estimate equation (5) on salaries, bonus payouts, performance ratings, and bonus targets. The number of observations contributing to the estimates is slightly reduced relative to the sample in the decomposition due to the inclusion of worker fixed effects. If a worker is only observed for a single period she cannot contribute to the identification of ω, the coefficient on the term of interest, $male_i \times male_{M(i,t)}$.

In addition to this fully saturated specification, we estimate two less granular versions on the same observations. First, as specified in Equation (4), we do not account at all for unobserved heterogeneity at the worker and manager level. This allows us to include dummies for the worker and manager being male, respectively. In an intermediate step, we add manager fixed effects. A dummy indicating the gender of the worker can then still be estimated. Third, we estimate the fully saturated Equation (5).

Table 5 shows the results. The coefficients from the first row reports the adjusted gender gap when working under a female manager. The coefficient from the second row is the difference in outcomes of women when they work under a male instead of a female manager. The third row estimates the difference between the gender gap under male and female managers. Before turning to more granular estimations, we look at columns 1, 4, 7, and 10, which report the coefficients from estimating Equation (4).

The gender salary gap under female managers is 2.3 log points. Women who work for male managers earn 1.9 log points higher salaries. Men only earn 1.6 log points higher salaries, but the difference to women is insignificant. This means that the estimated gender pay gap at male managers is smaller than at female managers by 0.3 log points, but this is highly insignificant.

Turning to bonuses, we find that gender gaps are 4.1% larger when the manager is male. The other coefficients imply that there is no statistical gender bonus gap under

TABLE 5: Effects of Having a Same-Gender Superior

	log(Salary)			log(Bonus Payout)			High Performance			log(Bonus Target)		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Male	0.023	0.026		0.014	0.028		-0.037	-0.034		0.008	0.008	
	(0.006)	(0.005)		(0.012)	(0.010)		(0.008)	(0.009)		(0.008)	(0.006)	
Male Mngr.	0.019			0.018			-0.005			0.003		
	(0.007)			(0.013)			(0.008)			(0.010)		
Male × Male Mngr.	-0.003	-0.003	0.002	0.041	0.017	0.051	0.014	0.006	0.027	0.005	0.007	0.005
	(0.007)	(0.006)	(0.003)	(0.014)	(0.012)	(0.015)	(0.010)	(0.011)	(0.015)	(0.009)	(0.007)	(0.005)
Job FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Tenure FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Manager FE		Yes	Yes		Yes	Yes		Yes	Yes		Yes	Yes
Worker FE			Yes			Yes			Yes			Yes
N	55,464	55,464	55,464	55,464	55,464	55,464	55,464	55,464	55,464	55,464	55,464	55,464
R^2	0.903	0.954	0.995	0.816	0.875	0.939	0.073	0.191	0.563	0.861	0.934	0.988

Notes: Standard errors are in parentheses and clustered at the worker and manager level. Age and tenure are summarized in six-year bins. Jobs are the combination of hierarchical rank and occupation.

female managers and no advantage for women who work for male managers instead of female managers.

The increase in gender bonus gaps when the manager is male should be explained by higher salaries, performance ratings, or bonus targets. However, as we can see from the third row in columns 1, 7, and 10 this is not the case. The reason for this is that the high performance indicators are noise measures of actual performance ratings, which are more nuanced. We cannot reconstruct the exact mapping from more detailed performance ratings to bonuses. Also, omitted variable bias in unobserved worker or manager characteristics could invalidate the estimation approach.

In columns 2, 5, 8, and 11 we control for unobserved manager characteristics. Three observations can be made in comparison to the version without manager effects. First, the gender gap under female workers in salaries and ratings is almost unchanged. Second, the bonus gap under female workers now is higher and significant at 2.8 log points. Third, differences in gender gaps are all insignificant and smaller.

While the gender gap at female managers is an interesting statistic, we are looking for a causal interpretation of the effect of manager gender on gender gaps. To ensure that workers are comparable, we add worker fixed effects in the complete specification of Equation (5) in columns 3, 6, 9, and 12. This proves critical for the estimation of the coefficient of interest.

Working for a male manager increases bonuses by 5.1 log points relative to women when controlling for worker and manager fixed effects. Put differently, the within-manager gender gap increases by 5% if the manager is male. Due to symmetry, this also implies that the within-manager gender gap falls by 5%, i.e. moves in favor of women, if the manager is female. There is no evidence that the manager's gender matters for gaps in base salaries or bonus targets. Instead, the advantageous position of men at male managers is driven by changes in performance ratings. The performance gap increases by 2.7 percentage points, i.e. moves in favor of men, when the manager is male. This represents more than 10% relative to the observed probability of receiving a high rating. The finding that gender gaps open up when the manager is male relates to Hospido and Sanz (2019) and Mengel et al. (2019) who document differences in gender gaps for different decision-maker genders in academic settings and thereby challenges findings by Card et al. (2020) who cannot document such differences.

As the point estimate for high performance ratings is only statistically significant at 10%, we examine the robustness of our findings to alternative performance rating measures in Appendix Table A.7. The effect of male managers on the gender gap in the probability of receiving a very high rating (the highest possible grade) is not significant. However, Table A.7 also shows that the gender gap in the probability of receiving a bad rating (one of the two lowest possible grades) shrinks when the manager is male. Finally, we translate the six grades to a numeric scale from one to six, where six refers to the best grade and one to the worst. The gender gap in this measure significantly increases by 0.02 points when the manager is male. These alternative performance outcomes lend support to our finding that working for male managers benefits male workers' ratings more than those of female workers.

Heterogeneity in the Impact of Manager Gender on Gender Gaps We also allow the effect of male bosses on the gender gap to vary along other dimensions. Table 6 displays these results. Each panel and column refers to a separate regression. All estimates of manager gender impacts on base salaries are insignificant. The effects of manager gender on the bonus gap do not differ significantly by worker or manager characteristics. However, some interesting patterns emerge.

Younger workers seem to be more affected by the differential gender gaps. This could be because younger workers are found at lower ranks. But as panel B shows, if anything the opposite is true. It seems that at higher ranks the gender of the manager plays a larger role for gaps. In panel C it looks like the effects of manager gender are larger when the team is mostly male. This observation supports the interpretation that male bosses cater to the needs and preferences of male workers. Panels D to F look at characteristics of managers. The results for age and hierarchical rank follow a similar pattern as for workers. Because the cultural background of managers could impact their treatment of workers of different gender, we allow the effect to vary by the region of origin. We classify managers as "western" if they have a European or Anglo-Saxon nationality. However, the point estimates look very similar. Overall, there is little statistical variation in the effect of managers, suggesting that the effect of manager gender on the gender gap is an issue across the entire firm.

Other Manager Characteristics It could be that manager characteristics apart from gender are the actual fundamental drivers of within-manager gaps. Note that even if this was the case, it is already clear that managers do matter. Here we examine whether gender indeed drives this finding.

We repeat the estimation of Equation (5) for all four outcomes considered, but in addition interact the dummy $male_i$ with manager characteristics correlated with manager gender. We allow gender gaps to vary by managers' age, tenure, origin, and hierarchical rank. If managers can experience a change in their characteristics, for example in hierarchical rank, the manager fixed effect does not absorb it. In this case, in addition to $male_i \times rank_{M(i,t)}$ we include $rank_{M(i,t)}$ as a regressor, etc.

Comparing the results from Table 7 with the original results from Table 5 shows that even if we control for the interaction between worker gender and various manager characteristics, the estimated coefficient on $male_i \times male_{M(i,t)}$ in the first row is almost identical. This suggests that omitted variable bias is not driving our finding but that indeed manager gender affects within-manager gender gaps.[15]

Table 7 also reveals that older managers close gender gaps in base salaries and contracted targets, relative to younger managers. The estimation also suggests that gender gaps in performance rating are greater when the manager is "western", suggesting that coming from a culturally more progressive society does not guarantee a reduced impact of managers on gaps. The finding could also mean that managers from a minority group are more concerned about the interests of other minorities.

Allowing for Differential Returns for Men and Women In the estimation of Equation (5), returns from male managers can differ between men and women while all other characteristics are assumed to have equal effects on men's and women's outcomes. This simplification could lead to bias if the differential returns capture differential returns from other characteristics. For example, if male bosses are more likely to work in engineering, and men have, for whatever reason, higher returns than women when working in engineering, we would blame managers for these differential returns.

In Appendix Table A.8 we repeat the estimation of Equation (5) but interact all variables except worker and manager effects with a gender dummy. The effect of managers on gender gaps is unchanged. This adds to confidence that we are indeed estimating the gender-differential effect of bosses.

Controlling for Additional Workplace Characteristics We also examine the robustness of our findings when adding controls for work units. In the gender gap decomposition, this exercise did not affect the interpretation of the results. This is also the case here. In Appendix Table A.9 we report estimates which are highly similar to the effect of manager gender on gender gaps shown previously.

Identification from Workers not Switching Jobs In Equation (5), the coefficient of interest ω is identified by workers who work for different managers during the sample period. But identification could be compromised if a worker selects into a new job with a different boss based on unobserved time-varying performance of the worker. While our analysis of event studies (based on Card et al., 2013) suggests that assuming that the matches of workers and managers are as good as exogenous, we can of course not formally test whether this assumption is true. In Table A.10 we therefore replicate our analysis based on a sub-sample of workers who do not face a new manager because of job changes of

[15]With a p-value of 12%, the coefficient on $male_i \times male_{M(i,t)}$ in Table 7 is statistically insignificant when considering gaps in the probability of receiving a high performance rating. This again might be because the grading is more complex than just an indicator for high ratings. Furthermore, while effects on base salary and contracted bonus targets remain insignificant, these can of course contribute to the effect on bonus payments in column 2.

workers but because of new managers substituting the previous manager of a team. This means that we drop all observations where a worker who changes manager also changes her job. For this exercise we do not make use of the sample used in the decomposition where worker mobility was key in order to find a connected set and identify all fixed effects. Here, we want to restrict the mobility of workers. The results in Table A.10 look very similar to our main result from Table 5, providing further confidence that assuming as good as exogenous matching of workers and managers is a reasonable approximation.

4.3 Why Does Manager Gender Matter?

Three main mechanisms could explain why the gender pay gap moves in favor of men when managers are male. First, the sorting of more productive men to male managers could result in an observed gender gap. Second, within-gender complementarities might make male workers relatively more productive than women when working under a male manager. Third, (unconscious) discrimination of workers of the other gender could drive our findings. Whatever explanation holds true, as long as women are underrepresented in managerial positions, these explanations would imply structural disadvantages for female workers. We now evaluate which of these mechanisms is likely to explain the effect of manager gender.

Sorting Workers are not allocated randomly to managers. Managers might be better informed about the quality of a worker of the same gender. This could for example imply that a male manager's male workers are on average better than female workers. While such a mechanism could exist in the firm, it cannot explain our results. In specification (5), we control for unobserved heterogeneity of workers. Doing so fully takes into account time-invariant ability differences. This means that our results hold true conditional on the sorting of workers to managers.

Within-Gender Complementarities A second explanation could be complementarities within gender. While there is ample evidence that diverse teams are more fact-focused, process facts more carefully, and are more innovative (e.g. Díaz-García et al., 2013; Herring, 2009; Levine et al., 2014; Nathan and Lee, 2013; Phillips et al., 2009), one could think that homogeneity can also benefit employee performance. For example, a competitive worker might be more productive in a competitive environment, while a cooperative worker might be more productive in a cooperative environment. If the distribution of styles differs for men and women, men should be more productive when working with men and we would observe a wider gender gap due to productivity.

While we have no measure of productivity available, we can test for the plausibility of this mechanism. Under the assumption that within-gender complementarities also would exist with respect to coworkers of the same gender, we can test for complementarities by studying the effect of male coworkers on the gender pay gap. We therefore repeat the estimation of Equation (5) for the share of male coworkers working for the same manager. The interaction of male-dummies for worker and manager is replaced by the interaction of a male-dummy for the worker and the share of male coworkers. We include the share of male coworkers because in contrast to manager gender it is not absorbed by the manager fixed effect. We still include worker fixed effects, manager fixed effects, and all other controls. This differs from Panel C of Table 6 as here we directly estimate the effect of male coworkers instead, not the heterogeneous effect of a male managers leading majorly male teams.

Table 8 contains only one significant coefficient, indicating that bonuses are higher when working in predominantly male teams. But, if anything, bonus payments are lower for men compared to women when working with more male coworkers. Under the assumption that within-gender complementarities also need to exist with respect to coworkers, the results imply that such complementarities do not exist with respect to manager gender. Therefore, we take this as supporting evidence that complementarities within gender are highly unlikely to substantially explain the favorable outcomes of men (women) under men (women).

Discrimination and Biased Beliefs Managers who are uninformed about the productivity of their subordinates might be more likely to resort to (unconscious) biases when evaluating or negotiating with workers. Discrimination due to biased beliefs (Bohren et al., 2019) implies that decision-makers discriminate according to their biased priors if they are uninformed. If they are provided with previous evaluations of work quality by other decision-makers, discrimination is reduced and eventually flips if much previous information is available.

We suggest that a related mechanism can drive discrimination in our setting. Instead of collecting information from previous evaluators, one can easily think that decision-makers are less likely to discriminate according to their preexisting biased beliefs as they learn about the worker's true quality. So while Bohren et al. (2019) relate a reduction in discrimination to better information, one can also expect that a reduction in discrimination is related to better knowledge of the worker.

Besides learning about the true quality of workers, managers learning about their workers also could exhibit less discrimination for a separate reason. It could be the case that managers are less aware of the needs of workers of the other gender. For example, male managers could be less considerate of child care obligations of female workers. Over time, they could learn about their workers' needs and reshape the work environment such that all workers can show their best.

Note that this does not necessarily mean that only male managers discriminate. Female managers could be discriminating as well, so that we estimate the additional advantageous treatment of men. Theoretically, it could also be the case that male managers do not discriminate at all, and female managers favor female workers. We do not consider this a likely scenario, as previous research showed that women if anything tend to discriminate against women (e.g. Bagues and Esteve-Volart, 2010) and bias in evaluations has been found in other settings to be driven by men (Hospido and Sanz, 2019; Mengel et al., 2019)

Whether discrimination takes more direct forms, e.g. biased ratings, or indirect forms, e.g. biased work environments, we would expect in both cases that discriminating is reduced when managers are more knowledgeable about their workers. To test this, we group observations by a number of variables capturing aspects of information.

First, we split managers into a group with less than ten and a group with more than ten years of tenure at the firm. One would expect that these managers are more experienced and therefore less likely to fall back to preexisting biases, or less aware about the needs of workers of the other gender. Second, we examine whether managers of larger teams (five or more subordinates) show stronger effects of manager gender on gaps. Managers of larger teams might find it harder to observe the true effort of each worker and to cater to specific needs. Similarly, a manager who works at a different location than the subordinate and who only communicates remotely might find it harder to evaluate workers or build a connection with workers who are less alike themselves. Finally, we allow effects to change with the time a worker and manager have worked together. We do so by grouping relationships into a group with two or less and a group with more than two years of joint work. Over time, one would expect that superiors learn more and build connections with workers. Note that this measure could be biased if workers who feel discriminated against are more likely to leave their team.

Table 9 shows the results. Each column of each panel refers to a separate regression. Sample sizes vary because some variables are not observed for all workers and managers. For salary, the effect of manager gender on pay gaps remains statistically insignificant. For bonus payments, one can observe a pattern that could be taken as support of the findings of Bohren et al. (2019), i.e. that better information leads to less discrimination. The coefficients are not statistically distinguishable at conventional levels. However, in all cases it seems that gender effects are larger in magnitude when information is harder to obtain. In particular, manager gender effects on pay gaps are larger when managers are less experienced, manage larger teams, work at different locations and know workers for a shorter amount of time. While none of these differences are statistically significant, all coefficients are negative and, in the case of location and joint time, close to significance with respective p-values of 12.0 and 15.8%, respectively. Proximity could also have led to the opposite effect on gender gaps if closer ties between managers and workers facilitate favoritism, an observation made by Cullen and Perez-Truglia (2019). Of course, such an effect might exist but seems to be more than netted out by managers accumulating additional information about workers.

The observation that overall the effect of manager gender on pay gaps seems to fall when the manager should be better informed supports the interpretation that some form of (unconscious) discrimination is an important driver of the results. While the observations here are in favor of discrimination by initially biased managers, they also contradict the mechanism discussed previously. If gender-specific complementarities would play a role, it is unclear why they should fall with better informed managers.

Table 6: Effects of Having a Same-Gender Superior: Heterogeneity

	log(Salary)	log(Bonus Payout)
Panel A: Age		
Male × Male Mngr.	0.003	0.062
	(0.004)	(0.021)
Male × Male Mngr. × > 44 years	−0.003	−0.027
	(0.005)	(0.028)
N	55,464	55,464
Panel B: Hierarchical Rank		
Male × Male Mngr.	0.003	0.052
	(0.004)	(0.037)
Male × Male Mngr. × Medium Rank	−0.002	−0.011
	(0.006)	(0.040)
Male × Male Mngr. × High Rank	0.005	0.070
	(0.011)	(0.058)
N	55,464	55,464
Panel C: Team Composition		
Male × Male Mngr.	−0.003	0.045
	(0.004)	(0.020)
Male × Male Mngr. × Majority Male	0.007	0.014
	(0.005)	(0.024)
N	55,016	55,016
Panel D: Manager Age		
Male × Male Mngr.	0.003	0.049
	(0.004)	(0.019)
Male × Male Mngr. × > 44 years	−0.002	0.006
	(0.006)	(0.024)
N	55,464	55,464
Panel E: Manager Hierarchical Rank		
Male × Male Mngr.	−0.011	−0.177
	(0.020)	(0.167)
Male × Male Mngr. × Medium Rank	0.012	0.215
	(0.020)	(0.168)
Male × Male Mngr. × High Rank	0.013	0.256
	(0.021)	(0.168)
N	54,654	54,654
Panel F: Manager Region		
Male × Male Mngr.	0.002	0.065
	(0.007)	(0.034)
Male × Male Mngr. × Western	−0.001	−0.018
	(0.008)	(0.038)
N	55,464	55,464

Notes: Standard errors are in parentheses and clustered at the worker and manager level.

Table 7: Effects of Having a Same-Gender Superior: Alternative Channels

	log(Salary)	log(Bonus Payout)	High Performance	log(Bonus Target)
	(1)	(2)	(3)	(4)
Male × Male Mngr.	0.001	0.055	0.024	0.004
	(0.003)	(0.016)	(0.015)	(0.006)
Male × Mngr. Age $\geq$ 45	-0.006	-0.020	0.012	-0.009
	(0.003)	(0.013)	(0.013)	(0.004)
Male × Mngr. Tenure $\geq$ 10	-0.001	0.003	-0.000	-0.002
	(0.003)	(0.014)	(0.013)	(0.005)
Male × Mngr. Western	0.011	-0.012	0.066	0.010
	(0.009)	(0.036)	(0.034)	(0.010)
Male × Mngr. High Rank	0.002	0.014	0.015	-0.005
	(0.005)	(0.020)	(0.020)	(0.008)
Job FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes
Tenure FE	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes
Manager FE	Yes	Yes	Yes	Yes
Worker FE	Yes	Yes	Yes	Yes
N	54,784	54,784	54,784	54,784
R^2	0.995	0.937	0.565	0.988

Notes: Standard errors are in parentheses and clustered at the worker and manager level. Age and tenure are summarized in six-year bins. Jobs are the combination of hierarchical rank and occupation.

Table 8: Effects of Having more Same-Gender Coworkers

	log(Salary)	log(Bonus Payout)	High Performance	log(Bonus Target)
	(1)	(2)	(3)	(4)
Male × Share of Male Coll.	0.001	-0.040	0.006	-0.001
	(0.006)	(0.028)	(0.027)	(0.011)
Share of Male Coll.	0.004	0.045	-0.004	0.006
	(0.006)	(0.023)	(0.022)	(0.009)
Job FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes
Tenure FE	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes
Manager FE	Yes	Yes	Yes	Yes
Worker FE	Yes	Yes	Yes	Yes
N	55,016	55,016	55,016	55,016
R^2	0.995	0.939	0.564	0.988

Notes: Standard errors are in parentheses and clustered at the worker and manager level. *Share of Male Coll.* is the share of colleagues, defined as workers with the same superior, who are male. Age and tenure are summarized in six-year bins. Jobs are the combination of hierarchical rank and occupation.

Table 9: Effects of Information

	log(Salary)	log(Bonus Payout)
Panel A: Manager Tenure		
Male × Male Mngr.	−0.001	0.058
	(0.004)	(0.021)
Male × Male Mngr. × $\geq$ 10 yrs Mngr. Tenure	0.006	−0.013
	(0.005)	(0.027)
N	55,450	55,450
Panel B: Team Size		
Male × Male Mngr.	0.002	0.057
	(0.003)	(0.017)
Male × Male Mngr. × Small Team	−0.002	−0.022
	(0.005)	(0.023)
N	55,464	55,464
Panel C: Location		
Male × Male Mngr.	−0.000	0.096
	(0.007)	(0.032)
Male × Male Mngr. × Same Location	0.001	−0.054
	(0.008)	(0.034)
N	54,326	54,326
Panel D: Joint Time		
Male × Male Mngr.	0.001	0.057
	(0.003)	(0.017)
Male × Male Mngr. × > 2 Years Joint	0.001	−0.040
	(0.005)	(0.028)
N	49,338	49,338

Notes: Standard errors are in parentheses and clustered at the worker and manager level. Age and tenure are summarized in six-year bins. Jobs are the combination of hierarchical rank and occupation.

5 Conclusion

Using novel personnel data from a large company we examine gender gaps in wages and performance, and we show that they are affected by managers. Our analysis yielded three main findings.

First, a significant part of the gender gap can be explained by the sorting of male and female workers to different types of managers. Men are more likely to work for "better" managers, i.e. managers whose workers receive higher salaries, bonuses, and performance evaluations. The observation that women tend to work for "worse" managers means that firms seeking to improve gender equity need to find out what makes a "good" manager and why women tend to work for managers with a lower impact. While our research cannot determine which underlying manager characteristics drive wage inequality or why workers are sorted as they are, our results imply that firms should not only foster the occupational upgrading of female employees, but also consider how female workers can work for "better" bosses. However, women might actually prefer to work for "worse" managers if these offer more family-friendly environments (Goldin and Katz, 2015).

Second, we show that the gender pay gap cannot be explained by the notion that men outperform women. On the contrary, women actually receive significantly better ratings than men. Yet, on average, they earn less. This is a striking result as it implies that adjusted gender pay gaps should be even larger. It also challenges the notion that women's performance is worse due to being less ambitious or competitive (Azmat and Ferrer, 2017; Niederle and Vesterlund, 2007). In spite of the positive impact of ratings on bonus pay, bonus pay gaps still favor men because the performance effect is outweighed by differences in salaries and targets. Firms often resort to performance ratings determined by superiors if actual output cannot be quantified, as is typical in complex organizations characterized by division of labor. Future research should examine whether male and female superiors differ in what aspects of worker performance they value most when determining ratings. If firms interpret performance ratings as good proxies of actual performance, gender equity is not achieved when women earn the same as men. If anything, a negative gender performance gap means that women should earn more than men.

Third, we show that manager gender matters as male managers cause within-team bonus gaps to increase. This is driven by the fact that performance ratings are relatively more favorable towards men when the manager is male. As manager gender affects pay gaps, the over-representation of men in management position puts women at a disadvantage. Therefore, our research has important implications for the discussion of gender quotas. The basic requirement for such quotas to work is that having more female managers indeed improves gender equality. In contrast to quotas applying at the executive level only (Bertrand et al., 2019; Maida and Weber, 2019), our findings imply that quotas across all hierarchical ranks can be effective. Future research would need to consider other requirements that need to be fulfilled such that gender quotas across all ranks are indeed a suitable policy.

Digging deeper, we find suggestive evidence that discrimination due to biased beliefs could drive the findings, as manager gender tends to matter less with more knowledge about the workers. Alternatively, managers might learn about workers' needs and improve upon their initially biased work environments. This observation can inform alternative pathways for promoting gender equity. Organizing employees into smaller and more stable teams in closer physical proximity could be a feasible measure to reduce gender gaps. In addition, many firms train their staff to be more aware of gender-related biases. While the success of such diversity programs is found to vary significantly (Chang et al., 2019), they seem to be a necessity to make managers more aware of their gender-related biases.

Bibliography

Abendroth, A.-K., S. Melzer, A. Kalev, and D. Tomaskovic-Devey (2017): "Women at Work: Women's Access to Power and the Gender Earnings Gap," *ILR Review*, 70, 190–222.

Abowd, J. M., R. H. Creecy, and F. Kramarz (2002): "Computing Person and Firm Effects Using Linked Longitudinal Employer-Employee Data," Tech. Rep. 2002-06, Center for Economic Studies, U.S. Census Bureau.

Abowd, J. M., F. Kramarz, and D. N. Margolis (1999): "High Wage Workers and High Wage Firms," *Econometrica*, 67, 251–333.

Altonji, J. G. and R. M. Blank (1999): "Chapter 48 Race and Gender in the Labor Market," in *Handbook of Labor Economics*, Elsevier, vol. 3, 3143–3259.

Azmat, G. and R. Ferrer (2017): "Gender Gaps in Performance: Evidence from Young Lawyers," *Journal of Political Economy*, 125, 1306–1355.

Babcock, L., S. Laschever, M. Gelfand, and D. Small (2003): "Nice Girls Don't Ask," *Harvard Business Review*.

Bagues, M. F. and B. Esteve-Volart (2010): "Can Gender Parity Break the Glass Ceiling? Evidence from a Repeated Randomized Experiment," *The Review of Economic Studies*, 77, 1301–1328.

Barnet-Verzat, C. and F.-C. Wolff (2008): "Gender Wage Gap and the Glass Ceiling Effect: A Firm-level Investigation," *International Journal of Manpower*, 29, 486–502.

Bertrand, M. (2011): "Chapter 17 - New Perspectives on Gender," in *Handbook of Labor Economics*, ed. by D. Card and O. Ashenfelter, Elsevier, vol. 4, 1543–1590.

Bertrand, M., S. E. Black, S. Jensen, and A. Lleras-Muney (2019): "Breaking the Glass Ceiling? The Effect of Board Quotas on Female Labour Market Outcomes in Norway," *The Review of Economic Studies*, 86, 191–239.

Bertrand, M., C. Goldin, and L. F. Katz (2010): "Dynamics of the Gender Gap for Young Professionals in the Financial and Corporate Sectors," *American Economic Journal: Applied Economics*, 2, 228–255.

Biasi, B. and H. Sarsons (2020): "Flexible Wages, Bargaining, and the Gender Gap," Working Paper 27894, National Bureau of Economic Research.

Blau, F. D. and L. M. Kahn (2017): "The Gender Wage Gap: Extent, Trends, and Explanations," *Journal of Economic Literature*, 55, 789–865.

Blinder, A. S. (1973): "Wage Discrimination: Reduced Form and Structural Estimates," *The Journal of Human Resources*, 8, 436–455.

BLS (2019): "Report 1083," *U.S. Bureau of Labor Statistics*.

Bohren, J. A., A. Imas, and M. Rosenberg (2019): "The Dynamics of Discrimination: Theory and Evidence," *American Economic Review*, 109, 3395–3436.

Bordalo, P., K. Coffman, N. Gennaioli, and A. Shleifer (2019): "Beliefs about Gender," *American Economic Review*, 109, 739–773.

Card, D., A. R. Cardoso, and P. Kline (2016): "Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women," *The Quarterly Journal of Economics*, 131, 633–686.

Card, D., S. DellaVigna, P. Funk, and N. Iriberri (2020): "Are Referees and Editors in Economics Gender Neutral?" *The Quarterly Journal of Economics*, 135, 269–327.

Card, D., J. Heining, and P. Kline (2013): "Workplace Heterogeneity and the Rise of West German Wage Inequality," *The Quarterly Journal of Economics*, 128, 967–1015.

Chang, E. H., K. L. Milkman, L. J. Zarrow, K. Brabaw, D. M. Gromet, R. Rebele, C. Massey, A. L. Duckworth, and A. Grant (2019): "Does Diversity Training Work the Way It's Supposed To?" *Harvard Business Review*.

Cullen, Z. B. and R. Perez-Truglia (2019): "The Old Boys' Club: Schmoozing and the Gender Gap," Working Paper 26530, National Bureau of Economic Research.

Díaz-García, C., A. González-Moreno, and F. J. Sáez-Martínez (2013): "Gender Diversity within R&D Teams: Its Impact on Radicalness of Innovation," *Innovation*, 15, 149–160.

Dohmen, T. J., H. Lehmann, and A. Zaiceva (2008): "The Gender Earnings Gap inside a Russian Firm: First Evidence from Personnel Data - 1997 to 2002," *Zeitschrift für ArbeitsmarktForschung - Journal for Labour Market Research*, 41, 157–179.

Egan, M. L., G. Matvos, and A. Seru (2017): "When Harry Fired Sally: The Double Standard in Punishing Misconduct," Working Paper 23242, National Bureau of Economic Research.

Frederiksen, A., L. B. Kahn, and F. Lange (2019): "Supervisors and Performance Management Systems," *Journal of Political Economy*, 128, 2123–2187.

Goldin, C. (2014): "A Grand Gender Convergence: Its Last Chapter," *American Economic Review*, 104, 1091–1119.

Goldin, C. and L. F. Katz (2015): "A Most Egalitarian Profession: Pharmacy and the Evolution of a Family-Friendly Occupation," *Journal of Labor Economics*, 34, 705–746.

Halldén, K., J. Säve-Söderbergh, and Å. Rosén (2018): "Gender of the Immediate Manager and Women's Wages: The Importance of Managerial Position," *Social Science Research*, 72, 115–133.

Herring, C. (2009): "Does Diversity Pay?: Race, Gender, and the Business Case for Diversity," *American Sociological Review*, 74, 208–224.

Hoffman, M. and S. Tadelis (forthcoming): "People Management Skills, Employee Attrition, and Manager Rewards: An Empirical Analysis," *Journal of Political Economy*.

Hospido, L. and C. Sanz (2019): "Gender Gaps in the Evaluation of Research: Evidence from Submissions to Economics Conferences," Working Paper 1918, Banco de Espana.

Ichino, A. and E. Moretti (2009): "Biological Gender Differences, Absenteeism, and the Earnings Gap," *American Economic Journal: Applied Economics*, 1, 183–218.

Juhn, C. and K. McCue (2017): "Specialization Then and Now: Marriage, Children, and the Gender Earnings Gap across Cohorts," *Journal of Economic Perspectives*, 31, 183–204.

Kahn, S. (1992): "Economic Implications of Public-Sector Comparable Worth: The Case of San Jose, California," *Industrial Relations: A Journal of Economy and Society*, 31, 270–291.

Kitagawa, E. M. (1955): "Components of a Difference Between Two Rates," *Journal of the American Statistical Association*, 50, 1168–1194.

Kleven, H., C. Landais, J. Posch, A. Steinhauer, and J. Zweimüller (2019a): "Child Penalties across Countries: Evidence and Explanations," *AEA Papers and Proceedings*, 109, 122–126.

Kleven, H., C. Landais, and J. E. Søgaard (2019b): "Children and Gender Inequality: Evidence from Denmark," *American Economic Journal: Applied Economics*, 11, 181–209.

Kunze, A. and A. R. Miller (2017): "Women Helping Women? Evidence from Private Sector Data on Workplace Hierarchies," *The Review of Economics and Statistics*, 99, 769–775.

Kurtulus, F. A. and D. Tomaskovic-Devey (2011): "Do Female Top Managers Help Women to Advance? A Panel Study Using EEO-1 Records:," *The Annals of the American Academy of Political and Social Science*.

Lazear, E. P., K. L. Shaw, and C. T. Stanton (2015): "The Value of Bosses," *Journal of Labor Economics*, 33, 823–861.

Levine, S. S., E. P. Apfelbaum, M. Bernard, V. L. Bartelt, E. J. Zajac, and D. Stark (2014): "Ethnic Diversity Deflates Price Bubbles," *Proceedings of the National Academy of Sciences*, 111, 18524–18529.

Maida, A. and A. Weber (2019): "Female Leadership and Gender Gap within Firms: Evidence from an Italian Board Reform," Discussion Paper 13476, CEPR.

Malkiel, B. G. and J. A. Malkiel (1973): "Male-Female Pay Differentials in Professional Employment," *The American Economic Review*, 63, 693–705.

Mengel, F., J. Sauermann, and U. Zölitz (2019): "Gender Bias in Teaching Evaluations," *Journal of the European Economic Association*, 17, 535–566.

Nathan, M. and N. Lee (2013): "Cultural Diversity, Innovation, and Entrepreneurship: Firm-Level Evidence from London," *Economic Geography*, 89, 367–394.

Niederle, M. and L. Vesterlund (2007): "Do Women Shy Away From Competition? Do Men Compete Too Much?" *The Quarterly Journal of Economics*, 122, 1067–1101.

Oaxaca, R. (1973): "Male-Female Wage Differentials in Urban Labor Markets," *International Economic Review*, 14, 693.

Pekkarinen, T. and J. Vartiainen (2016): "Gender Differences in Promotion on a Job Ladder: Evidence from Finnish Metalworkers:," *ILR Review*.

Pema, E. and S. Mehay (2010): "The Role of Job Assignment and Human Capital Endowments in Explaining Gender Differences in Job Performance and Promotion," *Labour Economics*, 17, 998–1009.

Phillips, K. W., K. A. Liljenquist, and M. A. Neale (2009): "Is the Pain Worth the Gain? The Advantages and Liabilities of Agreeing With Socially Distinct Newcomers," *Personality and Social Psychology Bulletin*, 35, 336–350.

Ragan, J. F. and C. H. Tremblay (1988): "Testing for Employee Discrimination by Race and Sex," *The Journal of Human Resources*, 23, 123.

Ransom, M. and R. L. Oaxaca (2005): "Intrafirm Mobility and Sex Differences in Pay," *ILR Review*, 58, 219–237.

Rothstein, D. S. (1997): "Early Career Supervisor Gender and the Labor Market Outcomes of Young Worker," in *Gender and Family Issues in the Workplace*, ed. by F. D. Blau and R. G. Ehrenberg, Russell Sage Foundation.

Sarsons, H. (2018): "Interpreting Signals in the Labor Market: Evidence from Medical Referrals," *mimeo.*

Sarsons, H., K. Gërxhani, E. Reuben, and A. Schram (forthcoming): "Gender Differences in Recognition for Group Work," *Journal of Political Economy.*

Sorensen, E. (1986): "Implementing Comparable Worth: A Survey of Recent Job Evaluation Studies," *The American Economic Review*, 76, 364–367.

A Additional Tables

TABLE A.1: Descriptive Statistics

	Mean	SD	10th Perc.	Median	90th Perc.	Observations
Salary [€]	64 315.32	75 438.37	23 323.65	56 252.3	114 753.28	59 813
Bonus [€]	12 293.32	21 707.14	1393.49	4816.59	28 924.78	59 813
High Performance	0.25	0.43	0.0	0.0	1.0	59 813
Low Performance	0.1	0.3	0.0	0.0	1.0	59 813
Bonus Target [%]	12.36	8.68	4.0	11.0	25.0	59 813
Male	0.58	0.49	0.0	1.0	1.0	59 813
Age	42.65	9.04	30.0	43.0	55.0	59 813
Tenure	10.73	9.24	2.0	8.0	25.0	59 813
Span of Control	1.64	4.55	0.0	0.0	6.0	59 813
Coworkers	8.33	9.6	2.0	5.0	18.0	59 813
New Manager in Same Job	0.24	0.43	0.0	0.0	1.0	46 534
New Manager in New Job	0.07	0.25	0.0	0.0	0.0	46 534
Male Manager	0.72	0.45	0.0	1.0	1.0	59 813
Male & Male Manager	0.45	0.5	0.0	0.0	1.0	59 813
Male & Female Manager	0.13	0.34	0.0	0.0	1.0	59 813
Female & Male Manager	0.27	0.44	0.0	0.0	1.0	59 813
Female & Female Manager	0.15	0.36	0.0	0.0	1.0	59 813
Age of Manager	45.92	7.93	35.0	46.0	56.0	59 813

Notes: Extreme values omitted for confidentiality reasons. Unbalanced panel based on 20 048 workers and the years 2014-2019. Monetary variables normalized to € in 2010.

Table A.2: Gender Gap Decomposition for Younger and Older Workers

(a) 44 and younger

	Salary		Bonus Payout	
	Log points	Share explained	Log points	Share explained
Age	0.40	6.2%	0.60	4.7%
Tenure	-0.01	-0.2%	-0.04	-0.3%
Manager	1.50	23.2%	4.39	34.9%
Job	3.87	59.9%	1.93	15.3%
Year	-0.13	-2.0%	-0.44	-3.5%
Country	0.95	14.8%	1.06	8.4%
Total explained	6.58	101.9%	7.50	59.6%
Total unexplained	-0.12	-1.9%	5.09	40.4%
Total gap	6.46	100.0%	12.59	100.0%

(b) 45 and older

	Salary		Bonus Payout	
	Log points	Share explained	Log points	Share explained
Age	0.11	0.6%	0.13	0.3%
Tenure	-0.00	-0.0%	0.16	0.4%
Manager	3.93	22.7%	4.85	13.1%
Job	13.21	76.4%	28.98	78.0%
Year	-0.29	-1.7%	-0.91	-2.5%
Country	-0.93	-5.4%	-1.42	-3.8%
Total explained	16.03	92.7%	31.78	85.5%
Total unexplained	1.27	7.3%	5.37	14.5%
Total gap	17.30	100.0%	37.15	100.0%

Notes: The tables display Kitagawa-Oaxaca-Blinder-decompositions of the gender gap in Log Salary and Log Bonus Payout based on 48 537 observations. Age and tenure are summarized in six-year bins. Jobs are the combination of hierarchical rank and occupation.

Table A.3: Gender Gap Decomposition Including Part-Time Workers

	Salary		Bonus Payout	
	Log points	Share explained	Log points	Share explained
Working hours	4.67	34.1%	1.93	7.5%
Age	0.56	4.1%	0.62	2.4%
Tenure	0.04	0.3%	-0.04	-0.2%
Manager	1.45	10.6%	4.52	17.5%
Job	5.72	41.8%	13.08	50.7%
Year	-0.25	-1.8%	-0.88	-3.4%
Country	-0.18	-1.3%	-0.54	-2.1%
Total explained	12.01	87.7%	18.69	72.5%
Total unexplained	1.68	12.3%	7.10	27.5%
Total gap	13.69	100.0%	25.79	100.0%

Notes: The table displays Kitagawa-Oaxaca-Blinder-decompositions of the gender gap in Log Salary and Log Bonus Payout based on 66 040 observations. Age and tenure are summarized in six-year bins. Jobs are the combination of hierarchical rank and occupation.

Table A.4: Gender Gap Decomposition Controlling for Accumulated Full-Time-Equivalent Months

	Salary		Bonus Payout	
	Log points	Share explained	Log points	Share explained
FTE months	-0.00	-0.0%	0.12	0.6%
Age	0.50	4.1%	0.51	2.4%
Tenure	-0.03	-0.2%	0.41	2.0%
Manager	1.37	11.2%	2.63	12.5%
Job	7.38	60.0%	12.21	58.1%
Year	-0.18	-1.5%	-0.25	-1.2%
Country	1.37	11.2%	0.42	2.0%
Total explained	10.42	84.7%	16.06	76.3%
Total unexplained	1.88	15.3%	4.98	23.7%
Total gap	12.30	100.0%	21.04	100.0%

Notes: The table displays Kitagawa-Oaxaca-Blinder-decompositions of the gender gap in Log Salary and Log Bonus Payout based on 24 389 observations. Age and tenure are summarized in six-year bins. Jobs are the combination of hierarchical rank and occupation.

Table A.5: Gender Gap Decomposition Controlling for Salary and Target

	High Performance	
	Percentage points	Share explained
Salary	1.62	-88.6%
Bonus Target	0.04	-2.1%
Age	-0.81	44.4%
Tenure	-0.03	1.5%
Manager	1.80	-98.7%
Job	-0.56	30.9%
Year	-0.04	1.9%
Country	-0.20	10.9%
Total explained	1.82	-99.7%
Total unexplained	-3.64	199.7%
Total gap	-1.82	100.0%

Notes: The table displays Kitagawa-Oaxaca-Blinder-decompositions of the gender gap in the probability of a high performance rating based on 59 813 observations. Age and tenure are summarized in six-year bins. Jobs are the combination of hierarchical rank and occupation. Salary and bonus target are measured in logs.

Table A.6: Gender Gap Decomposition Controlling for Unit

	Salary		Bonus Payout	
	Log points	Share explained	Log points	Share explained
Age	0.68	6.4%	0.90	4.3%
Tenure	0.13	1.2%	0.16	0.8%
Manager	2.29	21.4%	4.85	23.4%
Job	6.07	56.6%	11.72	56.5%
Year	-0.25	-2.3%	-0.83	-4.0%
Country	0.64	5.9%	0.24	1.2%
Unit	-0.15	-1.4%	-1.17	-5.6%
Total explained	9.41	87.8%	15.87	76.5%
Total unexplained	1.31	12.2%	4.89	23.5%
Total gap	10.72	100.0%	20.75	100.0%

Notes: The table displays Kitagawa-Oaxaca-Blinder-decompositions of the gender gap in the probability of a high performance rating based on 50 002 observations. Age and tenure are summarized in six-year bins. Jobs are the combination of hierarchical rank and occupation. Salary and bonus target are measured in logs.

Table A.7: Effects of Having a Same-Gender Superior: Additional Performance Measures

	High Performance	Very High Performance	Low Performance	Linearized Performance
	(1)	(2)	(3)	(4)
Male × Male Mngr.	0.027	0.009	-0.022	0.058
	(0.015)	(0.007)	(0.013)	(0.025)
Job FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes
Tenure FE	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes
Manager FE	Yes	Yes	Yes	Yes
Worker FE	Yes	Yes	Yes	Yes
N	55,464	55,464	55,464	55,464
R^2	0.563	0.469	0.533	0.607

Notes: Standard errors are in parentheses and clustered at the worker and manager level. Age and tenure are summarized in six-year bins. Jobs are the combination of hierarchical rank and occupation.

Table A.8: Effects of Having a Same-Gender Superior: Gender-Interacted Controls

	log(Salary)	log(Bonus Payout)	High Performance	log(Bonus Target)
	(1)	(2)	(3)	(4)
Male × Male Mngr.	-0.001	0.053	0.026	0.003
	(0.003)	(0.016)	(0.015)	(0.005)
Job × Gender FE	Yes	Yes	Yes	Yes
Year × Gender FE	Yes	Yes	Yes	Yes
Age × Gender FE	Yes	Yes	Yes	Yes
Tenure × Gender FE	Yes	Yes	Yes	Yes
Country × Gender FE	Yes	Yes	Yes	Yes
Manager FE	Yes	Yes	Yes	Yes
Worker FE	Yes	Yes	Yes	Yes
N	55,412	55,412	55,412	55,412
R^2	0.995	0.940	0.569	0.989

Notes: Standard errors are in parentheses and clustered at the worker and manager level. Age and tenure are summarized in six-year bins. Jobs are the combination of hierarchical rank and occupation.

Table A.9: Effects of Having a Same-Gender Superior: Controlling for Unit

	log(Salary)	log(Bonus Payout)	High Performance	log(Bonus Target)
	(1)	(2)	(3)	(4)
Male × Male Mngr.	0.002	0.047	0.030	0.004
	(0.003)	(0.017)	(0.017)	(0.005)
Job FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes
Tenure FE	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes
Manager FE	Yes	Yes	Yes	Yes
Worker FE	Yes	Yes	Yes	Yes
Unit FE	Yes	Yes	Yes	Yes
N	48,138	48,138	48,138	48,138
R^2	0.995	0.939	0.560	0.990

Notes: Standard errors are in parentheses and clustered at the worker and manager level. Age and tenure are summarized in six-year bins. Jobs are the combination of hierarchical rank and occupation.

Table A.10: Effects of Having a Same-Gender Superior: Keeping workers who do not switch managers due to job changes

	log(Salary)	log(Bonus Payout)	High Performance	log(Bonus Target)
	(1)	(2)	(3)	(4)
Male × Male Mngr.	0.001	0.052	0.044	-0.002
	(0.002)	(0.018)	(0.017)	(0.004)
Job FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes
Tenure FE	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes
Manager FE	Yes	Yes	Yes	Yes
Worker FE	Yes	Yes	Yes	Yes
N	68,719	68,719	68,719	68,719
R^2	0.996	0.939	0.574	0.991

Notes: Standard errors are in parentheses and clustered at the worker and manager level. Age and tenure are summarized in six-year bins. Jobs are the combination of hierarchical rank and occupation.