Firm Pay and Consumption Inequality*

Sigurd Galaasen† Andreas R. Kostol‡ Michael Simmons§

February 28, 2025

Abstract

Despite extensive evidence on pay disparities, the impact of firms on consumption and worker well-being remains largely unexplored. We address this gap by combining, job ladder models with incomplete markets, and a novel dataset that links workers' consumption patterns to their employment history across the entire population of Norway. Under assumptions on the environment and the specification for consumption choice, we show that how average consumption varies across firms is related to, and in some cases equivalent to, the role of firms for welfare inequality. By estimating two-way fixed effects models on pay and consumption, we find that a one-standard deviation increase in firm consumption effects is associated with an 11% increase in consumption. The correlation between firm pay and consumption effects is positive, but only about a quarter. Yet, our firm consumption index yields a closer relationship to standard revealed preference-based rankings of firms. Importantly, the variance in firm consumption effects is roughly half of the variance of firm pay effects, suggesting that worker optimizing behavior (consumption smoothing) mitigates firms' role for inequality in well-being.

*This project received financial support from the Norwegian Research Council through grants 315437. Simmons thanks Handelsbankens forskningsstiftelser under project number Bv19-0010. This paper should not be reported as representing the views of Norges Bank. The views expressed are those of the authors and do not necessarily reflect those of Norges Bank. We would like to thank Simon Mongey and seminar participants at the Second Annual Oslo Workshop and UCLA for helpful comments and suggestions.

†Norges Bank. E-mail: sigurd-molster.galaasen@norges-bank.no

‡BI Norwegian Business School; Norges Bank; IZA. E-mail: andreas.r.kostol@bi.no

§Umea University. E-mail: michael.simmons@umu.se

1 Introduction

An extensive literature has shown that firms are an important source of wage inequality (e.g., Abowd et al., 1999, Card et al., 2013, Kline, 2024). The evidence has sparked renewed interest in models of wage setting, firms' market power, and whether public policy can mitigate disparities in well-being.[1] However, the extent to which contemporaneous pay differences affect well-being remains an open question, partly because pay is not equal to consumption. In contrast, consumption depends on net present values (e.g., Friedman, 1957) and expectations of working at lower- or higher-paying firms in the future (e.g., Lise, 2013). On top of future considerations, is also widely acknowledged that firms offer amenities beyond the wage that matter for welfare (Rosen, 1986), which may offset or even widen gaps in welfare.[2]

In this paper, we propose an alternative approach in which the firms' role in shaping the distribution of consumption provides a sufficient statistic for inequality in welfare. The core idea is that consumption choices depend on the current and future value of a job, as well as on its welfare-relevant attributes. Similar to methods for identifying the firms' role in shaping the distribution of pay, bringing this idea to the data requires a measure of consumption that can be linked to a history of employment relationships. Using existing econometric techniques, it is possible to recover the contribution of a firm for pay and consumption from workers who move between different firms, and to compute the firm's role in explaining the variance of pay and consumption.

We develop our empirical framework based on a standard model of search and consumption choices, following Burdett and Mortensen, in which risk-averse workers anticipate future transitions to higher- or lower-paying firms, as in Lise (2013). We demonstrate that the existence of a job ladder reduces the variance of firm components in log consumption relative to log wages, as estimated using a two-way fixed effects model in the framework of Abowd et al. (1999). When introducing additional heterogeneity, such as differential risks of job loss or injury, we find that variance of firm components in consumption can exceed that of wages, and the wage ranking no longer tracks the consumption ranking of firms. While consumption variance does not account for the role of additively separable amenities,

[1] For example, Card (2022) set forth a research agenda on wage setting and market power in his 2022 AEA presidential address. On July 9, 2021, President Biden signed an executive order on Promoting Competition in the American Economy, emphasizing that competitive labor markets can lead to better job opportunities and higher wages. The order directed the Treasury Department, along with the Departments of Justice and Labor and the Federal Trade Commission, to examine how reduced competition affects the U.S. labor market (see e.g., The Department of the Treasury, 2022)

[2] Recent evidence based on data-driven revealed-preference approaches suggests that high-paying firms are often also more desirable; see, e.g., Sorkin (2018), Bagger and Lentz (2019), Crane et al. (2023) for different approaches, and Kline (2024) for a review of the findings.

as proposed by Rosen (1986), and more recently by Card et al. (2018), we show that we can gauge the importance this type of amenities by correlating the consumption type of a worker and the pay type of a firm.[3]

Our approach builds on recent advancements in micro-data collection. While several countries have seen a surge in the availability of transaction data from large banks, we draw on the unique advantages of the Norwegian setting. The Norwegian context is particularly well-suited due to an extraordinarily long panel data set of debit card and electronic payments, which we can link to the history of employment relationships. This unique combination enables us to track spending responses to transitions between firms at different pay percentiles and to decompose wage and consumption expenditure variance into worker and firm contributions. Moreover, we estimate the differential spending responses associated with moving to higher-paying firms across different categories. As noted by Deaton (1992), the inability to smooth some goods, such as food expenditure has distinct welfare implications relative to others, such as consumption of services or durable goods.

Our analysis yields three key findings. First, we show that when a worker moves to a higher-paying firm, consumption expenditure moves in the same direction as pay. Specifically, moving to a higher-paying firm-quantified by the change in co-worker log wages-is associated with a 4.6 percent increase in the log of disposable income and a 1.8 percent increase in total consumption, yielding a consumption elasticity of income of approximately 40 percent. These consumption effects vary in magnitude across goods and services. The elasticity of spending on restaurants and furnishings approaches 1, whereas the income elasticity of food and car expenditures is nearly zero and statistically indistinguishable from zero.

Second, using a two-way fixed effects model for wages and consumption expenditure, we find that where a worker works matters for how much they consume: Moving to a one standard deviation higher consumption firm, yields an 11 percent increase in consumption, while the corresponding increase in wages is 6 percent. We perform a variance decomposition following the bias-correction method proposed by Kline et al. (2020), we find that the variance in firm consumption effects is roughly half of variance of firm pay effects.

Third, we show that the firm wage and consumption rank is positively correlated. However,the Spearmans rank correlation coefficient is just .27, suggesting consumption-relevant attributes lead to rank reversals in the firm's wage and consumption types. We explore whether consumption-relevant attributes are welfare-relevant following the revealed preference approach of Sorkin (2017). We find that the rank of firm values is better correlated

[3]To see this, note that a high-consumption worker with low pay, must have high consumption due to high wealth (or due to differences in preferences). Because the type of amenities cannot be purchased, high-wealth types will want to work for firms that produce amenities instead of paying high wages.

with firm consumption effects than wages, which supports the interpretation of the variance of firm's consumption effects as a sufficient statistic for inequality in welfare.

Our paper relates to a long and influential literature considering the role of firms for wage inequality using the seminal framework in Abowd et al. (1999). A robust finding, that has been refined with modern econometric techniques, is that the variance of firm pay premiums are substantial and there is positive sorting (Bonhomme et al., 2023). As well as this, firm pay effects are strongly related to different observable measures of firm quality, including value added, and firm size (Bloom et al., 2018; Card et al., 2016; Bagger and Lentz, 2019). We contribute to these findings along several important dimensions. First, our extraordinary dataset allows us to consider the framework in Abowd et al. (1999) but for expenditure. Our finding of a smaller contribution of firms to consumption than spending echoes insights from Lise (2013), where spending depends on a workers relative position on the job ladder.[4]

Our results also corroborate earlier work showing that firm pay is strongly correlated with desirable aspects of the job (e.g., Sorkin, 2018 and Bagger and Lentz, 2019). While higher paying firms are also higher spending firms, on average, the rank correlation coefficient is far from one, and in fact closer to zero. Thus, there are several firms who have high firm pay effects but low firm consumption effects, and vice versa. More broadly, the data suggests that non-pay job characteristics that affect consumption are an important determinant of consumption variation across firms.

Researchers have long been interested in the existence and role of sorting in labor market (Shimer and Smith, 2000; Eeckhout and Kircher, 2011). Often, empirical researchers consider the correlation between firm effects and person effects as a summary measure for the presence of sorting. Recent econometric techniques (Bonhomme et al., 2019; Kline et al., 2020), as a response to concerns over limited mobility bias (coined in Abowd et al. (2002)) has refined the conclusion that the correlation coefficient between worker and firm effects is consistently positive across countries. We show a similar results for consumption expenditure, both in the impact of mobility bias and the positive sorting. However, the correlation coefficient is substantially smaller for consumption.

Our paper also contributes to a long literature considering consumption inequality. Deaton (1992) and Attanasio and Pistaferri (2016) discuss the strengths and weaknesses of various data sets of consumption, in particular what is missing from income when thinking about inequality in welfare. More recent studies have used de-identified bank account data to study a variety of consumption-related questions, including Gelman et al. (2014), Baker (2018),

[4]Workers at lower rungs of the ladder foresee large wage gains as they climb the ladder and so dis-save. Workers at higher rungs of the ladder instead choose to save (or dis-save less) to attenuate the welfare consequences of unemployment.

Ganong and Noel (2019), Gerard and Naritomi (2019), and Andersen et al. (2023). We contribute to recent studies using transaction-level data by linking the universe of all debit card transactions and interbank wire-transfers with administrative information about employment relationships.

Our paper is related to the seminal theories of consumption, including Modigliani and Brumberg (1954) and Friedman (1957), who highlight workers' desire to transfer resources across states or nature or from the present to the future to maintain a stable consumption flow. The extent to which households can achieve this depends on the available consumption smoothing devices. This literature has also explored how consumption varies across individuals and evolves over time (see Attanasio and Pistaferri, 2016 for a comprehensive discussion), and how spending reacts to transitory and permanent shocks (e.g., Blundell et al., 2008), and consumption network effects from co-workers De Giorgi et al., 2020). We contribute to this body of work by investigating the extent that workplaces matter for consumption variance and levels.

The paper proceeds as follows. Section 2 considers the role of firms in shaping consumption using a job search framework and a two-way fixed effects regression framework. Section 3 describes the data and Section 4 offers evidence of wage and consumption changes due to changes in co-workers. Section 5 performs the main variance decomposition of pay and consumption. Section 6 concludes.

2 Considering the content of consumption effects

This section develops an extended off the shelf random search model with risk averse workers and two-sided heterogeneity and shows how the firm's role in shaping wages and consumption can be quantified using a standard two-way fixed effects regression framework.

2.1 A random search model with risk averse workers and two sided heterogeneity

Preliminaries. Time is continuous, and workers die at some rate d, and are replaced by identical workers with zero assets. Workers can either be unemployed or employed. Next, we describe the model in different sections: preferences and decision-making, heterogeneity, job finding and job losing, production and wages, and the timing of events.

Preferences and decision making. Workers are risk averse and derive utility from consumption, c, and amenities, ϑ, through the utility function, $u(c, \vartheta)$, where the function is strictly

increasing and concave in both arguments. The social discount rate is $\rho = \tilde{\rho} + d$, where each time period corresponds to a month. Individuals maximize the expected presented discounted utility written as

$$\mathbb{E}_0 \int_0^\infty e^{-\rho t} u(c_t, \vartheta_t) dt. \tag{1}$$

Each individual chooses to invest in a single risk-free asset, a, with rate of return equal to $r < \rho$, where assets evolve following

$$\dot{a} = ra + i - c,$$

where, i represents the wage, w, or the unemployment benefit received, b. The worker's objective is to pick their consumption c to maximize the expected present discounted value of working or not working. Throughout, a natural borrowing limit, $\underline{a}$, is imposed such that the worker maintains positive consumption at the lowest flow income level.[5]

Heterogeneity. Workers differ in their ability in production that we denote by $h \in [0, 1]$. Upon entry, worker ability is distributed uniformly. Firms differ in three ways, their productivity, $p \in [0, 1]$, the risk at which the job is destroyed, $\delta \in [0, 0.5]$, and amenity, $\vartheta \in [0, 1]$. These characteristics are all permanent.

Job finding and offers. Employed workers receive offers at rate λ^e and at rate λ^u when unemployed. There are three distributions characterizing firm heterogeneity: the marginal offer CDF of firm productivities, F_p, the marginal offer CDF of firm job destruction risk, F_δ, and the marginal offer CDF of firm job amenity, F_ϑ. We assume these are governed by a beta distribution and we assume independence such that the offer CDF is $F = F_p F_\delta F_\vartheta$.[6]

Production and wages. The purpose of the exercise is to consider the content of consumption effects when wages take the form considered in Abowd et al. (1999). While this appears reduced form at first, such a wage equation can be microfounded in this setting following the bargaining protocol of Elsby and Gottfries (2022) which we discuss in the Appendix. We assume workers receive benefits given as $b(h) = bh$.

[5] This implies $\underline{a} = -\underline{i}/r$.

[6] It is straight forward to incorporate some dependence using a copula but we abstract from this here.

2.2 The consumption function

Before we move to the empirical analysis, we next describe the consumption function in the model. A standard intertemporal optimality condition states that the marginal utility of consumption is equal to the marginal value of assets. The workers' optimal level of consumption whilst employed solves

$$u'(c, \vartheta) = V_a, \tag{2}$$

where V is a workers presented discounted utility. Consumption is equal to the inverse function of the marginal utility of consumption over the shadow price of assets. A worker's consumption is implicitly determined by the net-present value of the job, and how it varies with wealth. Note that the function is a highly non-linear object. Understanding the underlying properties, and the role of firms, thus requires simulation of the model.

2.3 Assessing pay and consumption effects

We must first solve and simulate the model to quantify the empirical content of pay and consumption. To do so, we set the utility function to the common and general CRRA representation and let the utility from amenities be in logs, written as as

$$u(c_t, \vartheta) = \frac{c_t^{1-\gamma} - 1}{1-\gamma} + \log(\vartheta). \tag{3}$$

Note that we impose an additively separable amenity. We do this since it is common formulation in the literature (Card et al., 2018; Sorkin, 2018; Lamadon et al., 2022), but we will also show that this matters for the ranking of firms later, since the consumption is not directly impacted by the availability of other amenities, or $u_{c\vartheta} = 0$. We let the distributions describing firm heterogeneity be symmetric beta distributions. We discretize the distributions and let a firm be a (p, δ, ϑ) combination. We assume a uniform distribution for worker skill levels, set $\gamma = 2$ which is a common value in the literature, and set $\rho = .0062$ to replicate a 5% annual discount rate and choose an interest rate, r, to be half of ρ. We choose b to be 0.2 which roughly reproduces a replacement rate of 0.4 in the steady-state economy (Shimer, 2005). Lastly, we must decide on job offer arrival rates in unemployment and employment, as well as the underlying offer distributions to construct. We set $\lambda_u = 0.15$, $\lambda_e = 0.05$ and assume that beta parameters of the offer distribution equal to 1.5. Given these parameters, we solve for the associated policy rules, and simulate the model at a monthly time interval for 50000 individuals for twenty years and 250 firms. We use the last ten years for estimation.

Next, we estimate a two-way fixed effect model for worker and firm wage effects for pay

and consumption. Following the notation in Card et al. (2013), the time interval is denoted t, there are N^* person-firm observations, with N workers and J firms. J(i, t) is a function mapping worker i to a particular firm at time t. We assume that the natural logarithm of wages, y^w_{it}, and consumption, y^c_{it}, can be expressed additively as the sum of a fixed worker component, α_i, a fixed firm component, $\psi_{J(i,t)}$, time varying observable characteristics, $x'_{it}\beta^o$, and an error term, $r_{i,t}$. For wages we will estimate the following equation

$$y^w_{it} = \alpha^w_i + \psi^w_{\mathbf{j}(i,t)} + r_{i,t}. \tag{4}$$

The worker effect (α^w) can be interpreted as skills that are rewarded equally across firms, and the firm effect (ψ^w) can be interpreted as a premium that the firm j pays to all workers. This interpretation is consistent with the model of wages we have described.

Similarly, for consumption we, will estimate

$$y^c_{it} = \alpha^w_i + \psi^c_{\mathbf{j}(i,t)} + r_{i,t}. \tag{5}$$

In contrast to the model for wages, the log-additive specification is not generally consistent with our model for consumption choices. Reassuringly, we will find that worker and firm effects explain virtually all the variation in consumption in the simulated data.

2.4 Variance Decompositions and Rank Correlations in the Models

The goal of this section is to perform decompositions of the total variance in wages and consumption under different configurations of the model. We provide further details of the model simulation and regressions in Appendix A and summarize the main predictions from different configurations of the model here.

1. The variance of firm consumption effects is lower than the variance of firm wage effects when firms only vary in productivity

2. The variance of firm consumption effects is wider than the variance of firm wage effects when the dispersion parameters of risks is above some threshold. The threshold is lower when the "cost of job loss" are larger.

3. The rank-rank correlation between a firms wage rank and its consumption rank is strictly below 1 if there is some dispersion in risk

4. High wealth individuals will tend to work for low-wage firms if additively separable amenities are welfare-relevant ($\vartheta > 0$)

The first result is that the firm's contribution to variance of consumption is strictly below that of its contribution to pay. This result comes directly from a workers' desire to smooth consumption. On the one hand, the worker wants to borrow against future income as the worker ascends the job ladder. On the other hand, as in the standard framework in Lise (2013), workers higher up in the ladder want to set aside income in anticipation of the costs of involuntary separations. This creates a lower gradient in consumption with wages than 1, and, therefore, lower variation in firm consumption effects than firm wage effects. As markets become more complete, the variance of consumption effects approaches 0. Whereas, if individuals are unable to save, and are more and more hand to mouth, the variance of consumption effects approaches the variance in wage effects. The strength of these conclusions naturally depends on the dispersion of productivity differences.

The second result is derived from introducing the heterogeneous job loss risk. We simulate the model by varying the distribution of risk relative to the distribution of productivity, and obtain the variance of firm effects in pay and consumption. The left panel of Figure 1 shows that the variance of firm's consumption effects exceeds the variance of the firm's pay effects when the dispersion of unemployment risk is several times as wide as the dispersion in productivity. It is important to note that whether the variance of firm's consumption effects exceeds the variance of the firm's pay effects depends the consequences of unemployment. If, for example, job finding is lower in unemployment, the potency of unemployment risk for consumption is exacerbated. We note that while the model includes risk of layoff, the key is that there are out-of-pocket expenditure risks associated with the job. These risks include the net earnings loss, after taxes and unemployment benefits, but potentially other salient risks such as injury or disability. We summarize the two first results in Panel A of Table 1.

Next, we explore the productivity and risks of a firm in shaping the covariance between a firm's pay and consumption rank. In the baseline environment, when jobs are uni-dimensional, the consumption and wage effects are essentially perfectly correlated. We illustrate how the rank-rank correlation between firm pay and consumption effects varies with the dispersion of the job destruction distribution (holding the productivity dispersion fixed) in the right panel of Figure 1. We see that, as heterogeneity in employer characteristics becomes more and more dispersed, the pay content of a job becomes less and less related to the associated consumption content of a job. We summarize these conclusions in Panel B of Table 1.

The fourth result concerns the role of the additively separable amenity. A limitation of our approach was first highlighted in section 2.2, in that consumption does not directly depend on this type of non-pay attribute. In one extreme, these amenities are orthogonal to

Figure 1: Changing the dispersion of productivity and risk.

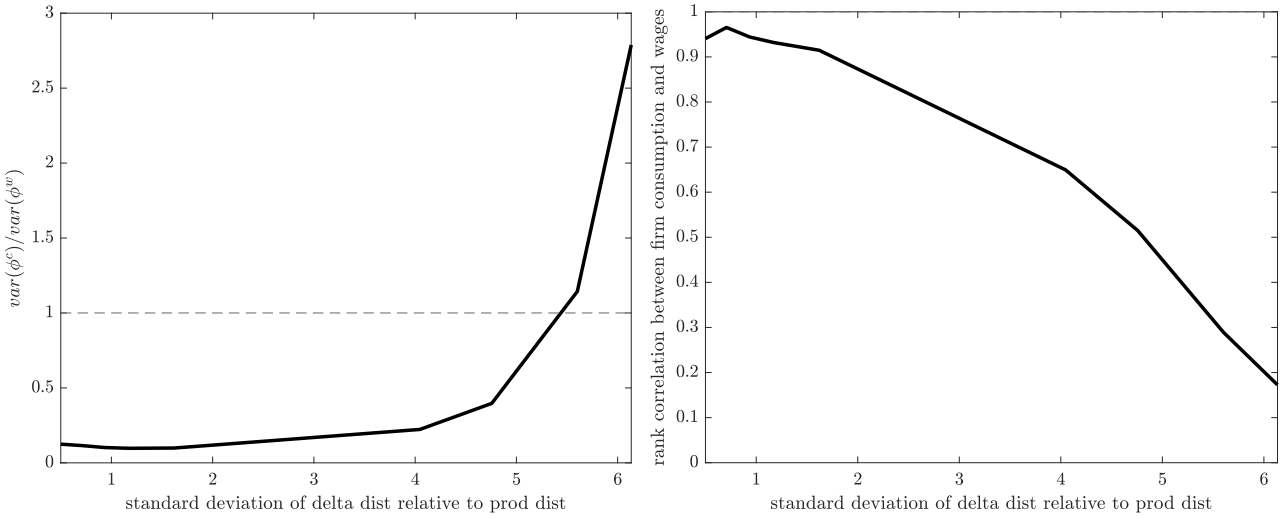

Notes: The figures compare the firms' contribution to the variance of consumption and wages in the left panel and the covariance between the firms' wage and consumption ranks in the right panel for different dispersion parameters of risk.

the consumption rank of a firm; while in another, these amenities may be correlated with the consumption rank, in which case our approach miss welfare-relevant aspects of firms. However, under the assumptions that the additively separable amenity (i) is a normal good that (ii) it cannot be purchased, workers with higher levels of lifetime wealth will sort to firms that offer this amenity instead of higher pay. This implies that we can assess the extent to which such welfare-relevant aspects of the job will alter our conclusion about welfare by regressing a measure of lifetime wealth on the likelihood of moving to a low-paying firm.

2.5 Relationship to firm ranks

Next, we relate the consumption rank of a firm to two well-established methods to rank firms: the Poaching Rank (Bagger and Lentz, 2019), and the Page Rank (Sorkin, 2017). Throughout we will sidestep the issue that workers may not perfectly agree on the ranking of firms because of differential asset levels, and assess quantitatively whether this changes the basic patterns substantially in the quantitative exercise. We describe these ranking methods carefully, and the representations through the lens of the model in the Appendix. The ranking measures are naturally essentially identical in simulated data. The key result here is that, as firms become progressively more heterogeneous and workers are allowed to save, the consumption firm effects provide a perfect ranking with firm values. While wages become progressively less correlated the more dispersion there is in the job destruction distribution. We present this result in Figure A3, and summarize it in Panel C of Table 1. As

Table 1: Firm wage and firm consumption effects through the lens of the model

Model	BM (1998)	Lise (2013)	Lise (2013) + Jarosch (2023)	Lise (2013) + Jarosch (2023) amenities
	(I)	(II)	(III)	(IV)
	A. Ratio of variance of fixed effects: $Var(\psi^c)/Var(\psi^w)$			
	1	<1	$\lessgtr$1	$\lessgtr$1
	B. Rank-rank correlation between the two firm effects			
	1	1	<1	<1
	C. Correlations between Sorkin (2017) firm values with ...			
consumption	1	1	1	<1
wages	1	1	<1	<1

Notes: Panel A shows the predictions for the ratio of the variance of consumption effects and the variance of wage effects. Panel B shows predictions for the rank-rank correlation between firm wage effects and firm consumption effects. Panel C shows predictions for whether the correlation between firm values and firm consumption is greater than the correlation between firms values and firm wages. = implies that the correlations are the same. > implies that the firm consumption is better correlated with values than firm wages. $R^2 = 1$ for all wage regressions and $R^2 \approx 1$ for all consumption regressions.

well as this, since consumption is not directly related to the separable amenity, adding this feature to the model results in a less than perfect correlation of firm values ranks and firm consumption ranks.

2.6 Welfare inequality across firms

Understanding the role of consumption across firms is particularly important because of its welfare relevance. We conclude this section by clarifying what we can say regarding *welfare* inequality. We consider welfare inequality using the consumption equivalent variation measure, which in general in our environment is written as solving

$$\mathbb{E}_0 \int_{t=0}^{\infty} e^{-\rho_i t} u(c_{ij}(1+\gamma_{ijk}), \vartheta_j) dt = \mathbb{E}_0 \int_{t=0}^{\infty} e^{-\rho_i t} u(c_{ik}, \vartheta_k) dt. \tag{6}$$

Let γ_{ijk} be the gain or loss in consumption required to make a move from firm j to firm k such that worker i is indifferent between the move.

Is the variance of AKM effects informative for this type of analysis? To test this, we proceed as follows. In the same sample used to estimated the two-way fixed effects regression, for each firm k, we look at changes in consumption as workers move from firm j to firm k. Then take a weighted average for each firm k, which then measures how much wages or consumption change as workers move from any firm to firm k. Note that the key point to this is that we do not impose a functional form assumption for consumption. If the AKM specification is accurate, we should reproduce the exact variance in firm effects. As a sanity check, when we conduct this exercise on wages, since the specification is correct we always produce the exact variation in firm effects. For consumption, the variance in firm consumption effects as estimated from the AKM is always very close to the variance in consumption effects as estimated not imposing a functional form assumption. Thus, the theory suggests that the estimated variance in consumption effects that we estimate has empirical content, in that it describes how consumption changes as workers move from one employer to the next, holding worker characteristics constant.

Given this result we may interpret a firm effect as an approximate homogenous treatment effect. The homogenous treatment effect implies that $\gamma_{ijk} = \gamma_{jk}$. When considering the role of firms for welfare inequality, we consider three cases, where the latter is most related to the job ladder model described earlier in this section. We treat each case as a planner having different preferences over present and future consumption, and a different understanding of the environment. Throughout this analysis, it should be made clear that we are concerned with the form of welfare that directly impacts the consumption choice, or determines the

workers future income stream. Alternatively, the following results can be interpreted as assuming that $\log(\vartheta)$ is zero across all firms.[7]

First, we consider a corner case where the planner is only concerned with the worker's contemporaneous consumption level. That is, the welfare measure solves

$$u(c_{ij}(1+\gamma_{jk}^1)) = u(c_{ik}) \implies c_{ij}(1+\gamma_{jk}^1) = c_{ik}.$$

Under this measure of welfare, it is possible to show that

$$\text{var}(\log(1+\gamma_{jk}^1)) = \text{var}(\psi^c). \tag{7}$$

Thus, the estimated variance of firms effects is equivalent to the role of firms for welfare inequality (the superscript refers to the case 1). This is powerful, but under strong assumptions that the planner does not care about the future consumption decisions of the worker. Nevertheless, this is a useful benchmark that we can then extend to consider when this equivalence goes.

In a second case, consider a planner that cares about the future consumption decisions of workers but assumes that worker's will continue to consume the same forever, or alternatively, that they will never switch firms or lose their job. It is straight forward to show that the equivalence shown in 7 remains, or $\gamma_{jk}^1 = \gamma_{jk}^2$. Thus, the variance firm effects is equivalent to the variance in welfare across firms when the planner assumes a constant consumption choice.

Finally, consider a more realistic third case where workers show a gradual tendency to move up and down the consumption job ladder, as is the case in job ladder models. Again, we assume that workers do not flow into unemployment. In such a case, we can show that

$$\text{var}(\log(1+\gamma_{jk}^3)) < \text{var}(\psi_j). \tag{8}$$

Thus, the variance in firm effects is an upper bound for the variance in welfare across firms. The reason for this is the following. Conditional on a worker's asset level, those working in the lowest consumption firm, will enjoy weekly greater consumption in the future through job switching up the ladder. Workers who are in the highest consumption firm, will enjoy weekly lower consumption in the future through falls down the ladder. Thus, dynamics attenuate differences in contemporaneous consumption, and therefore, welfare.

[7] It may be the case that an alternative is the assumption that ϑ is orthogonal to $w + \delta$ but we have yet to show this.

3 Data

This section describes the data, empirical setting and performs a descriptive analysis of consumption expenditure in the cross-section.

3.1 Employment Relationships and Labor Markets

Our empirical analysis draws on several attractive features of the Norwegian labor market. Norway has a population of approximately 5.5 million, with the capital Oslo accounting for almost 20 percent. We use administrative records to characterize the universe of employment relationships from 2006 to 2018.[8] We link this data to tax income data to measure annual wage income, taxes, residential municipality and other transfer payments, as well as population panel data, which allows us to characterize household members and measure some background characteristics, such as gender, age, and educational attainment. Commuting zones are defined using aggregate statistics between workers' residential and employers' workplace municipalities.

We will focus on single full-time employed aged 25 to 60 in our baseline analysis, and Table 2 reports key statistics for our population of interest.[9] We split the sample in two based on whether the individual is observed switching employer in the year, in which case the person belongs to a "mover" sample, where the remaining sample is denoted "stayers". The summary statistics show that movers are somewhat younger than the stayers, and that approximately 40 percent of the single population live and work in the commuting region of Oslo. Appendix Table C1 report the corresponding statistics for married and cohabiting, showing that the single are overrepresented in Oslo.

The Norwegian labor market is characterized by a combination of a generous unemployment insurance (UI) system and collective bargaining. The UI system covers almost two-thirds of lost earnings and can be extended to up to two years. The majority of workers are covered by agreements negotiated between trade unions and employers. About half of employees benefit from tariff agreements in the private sector, with a two-tier bargaining system. Tariff wages at first negotiated first at the industry-level are first set centrally, after which wages are supplemented by local adjustments, which is bargained over at the firm

[8] Before 2015, the data included employment relationships that had hours or income above a certain level while after 2014, the data was reported monthly for every type of employment except self-employment. The data provides start- and end dates, a measure of contracted hours, and the location of the workplace.

[9] Most individuals have completed a vocational program or college degree by age 25 and the majority of workers are not eligible for early retirement before after age 60. Our sample is further restricted to individuals who are residents of Norway throughout the period 2006 to 2018

Table 2: Summary Statistics for Single Full-Time Workers

Sample:	**(a) Stayers**		**(b) Movers**	
	Mean	St. Dev	Mean	St. Dev
Age	40.726	10.821	36.954	10.131
Male	0.657		0.649	
Lives in Oslo Region	0.393		0.424	
Household Size	1		1	
Household Wage Income	13.099	0.523	13.015	0.512
Household Income After Tax	12.873	0.371	12.806	0.387
Household Gross Wealth	13.420	2.391	12.976	2.513
Household Debt	13.171	1.944	13.143	1.907
Household Expenditure	12.341	0.829	12.313	0.793
Food	10.346	1.025	10.322	0.959
Clothing	8.476	1.288	8.547	1.280
Furnishing	8.509	1.464	8.488	1.461
Motor Vehicles	7.876	1.998	7.710	2.041
Electronics	7.225	1.484	7.233	1.501
Retaurants	8.898	1.377	9.084	1.343
Number of Observations	256,438		38,799	

Notes: The sample includes single fulltime workers aged 25 to 60 in 2016. The 2024 NOK/USD $\approx$ 10

level. The two-tier framework is considered a key reason for a relatively compressed wage structure in Norway.

Table 2 also reports the averages and standard deviation of workers' income and wealth. In the fifth row, we see that standard deviation of annual log wage income is 0.523 for stayers, with a variance of 0.274. Subtracting taxes and adding transfers gives us an individual's net income, available for consumption, reported in the sixth row. In the seventh and eight rows, we report total gross wealth and debt at the individual level. When calculating these outcomes, and in particular overall spending at the household level, we divide by the square root the household size.

3.2 Expenditure Data

The key to our empirical contribution is the combining the employment histories to a newly collected data set covering electronic payments for the universe of Norwegian residents. The electronic payments data are provided by the Norwegian retail clearing institution, Nets Branch Norway (henceforth referred to as Nets), covering two data sources:[10] (i) all debit card payments via BankAxept (ii) all online bank wire payments cleared via the Norwegian

[10]In addition the data covers all incoming transfers cleared via NICS. We do not make use of this data in this paper

Interbank Clearing System (NICS).[11] We received individual-level data aggregated at the weekly frequency, with information about total total expenditure, number of transactions, and cashback at a zip code and 25 broad consumption categories.

Overall, our data covers about 80 percent of the total value of electronic transactions made by Norwegian residents over the time period 2006-2018. Figure 2 shows that our expenditure measure tracks the quarterly level and growth rate of consumption from National Accounts well. The correlation between quarterly growth rates is 0.85.[12] The ninth row of Table 2 shows that the standard deviation of total log annual consumption expenditure is somewhat higher than the variance of annual wage income. Appendix B provides additional details on data cleaning procedures, such as trimming outliers, and comparisons to other measures of consumption expenditure.

In the following lines, we illustrate the log level of spending across cagetories. These categories are drawn from United Nations' 1999 COICOP system, which classifies consumption expenditures according to their purpose.[13] Using the industry code information, transactions to/from banks and the public sector are placed into two separate categories. As we explain in detail in Appendix B.3, we clean the bank payment category by removing payments that are likely related to servicing of mortgage (or other investment related) debt, and show that the remaining bank payments track aggregate statistics on credit card payments well. Hence, we interpret this category as capturing payments of credit card bills.

[11] BankAxept is the national payment system in Norway, owned by Norwegian banks. Typically, all debit card payments in domestic physical stores are BankAxept, whereas payments abroad, online or mobile payments are paid through VISA or Mastercard. NICS is the interbank clearing system for the Norwegian Krone (NOK). It is used by all banks operating in Norway, and that take part in the Norwegian banking community's infrastructure for payments. Transaction via NICS includes all invoices paid using a "KID-number", which includes all invoices paid via "Efaktura" and "Avtale Giro".

[12] National Accounts household consumption is net of imputed housing consumption, for which there is no corresponding transaction. Imputed housing consumption is only available at the annual frequency. We assume that imputed housing as a share of total housing-related consumption is fixed within the year to correct the quarterly series. Even if we do not correct for imputed housing, the Nets measure is highly correlated with National Accounts (0.83) while the R2 is 0.68.

[13] Card payments are mapped to COICOP based on the Merchant Category Codes of the card terminal (MCC). The COICOP classification was revised in 2018. Online wire transfers are classified based on the creditors' 5-digit industry codes (NACE) The crosswalks between MCC/NACE and COICOP are manually coded by Norges Bank and provided to Nets before aggregation. The correspondence tables are made available on the authors' websites. The top-level structure (first and second level), the level at which our data is aggregated on, remained mostly unchanged. The main change was the introduction of more granular categories (at a fourth level). Our categorization contains 24 COIOP groups and includes all 12 top-level codes which cover categories such as "Food and beverages","Restaurants", "hotels", "Clothing and footwear" and "Housing, water, electricity, gas". Some of the top-level codes are further divided into second-level COICOP groups, leaving us with 24 categories in total. Table B.1 in Appendix B.1 lists all our categories and provides details on how transactions are linked to a particular COICOP code.

Figure 2: Aggregated micro-level expenditure data compared with National Accounts data.

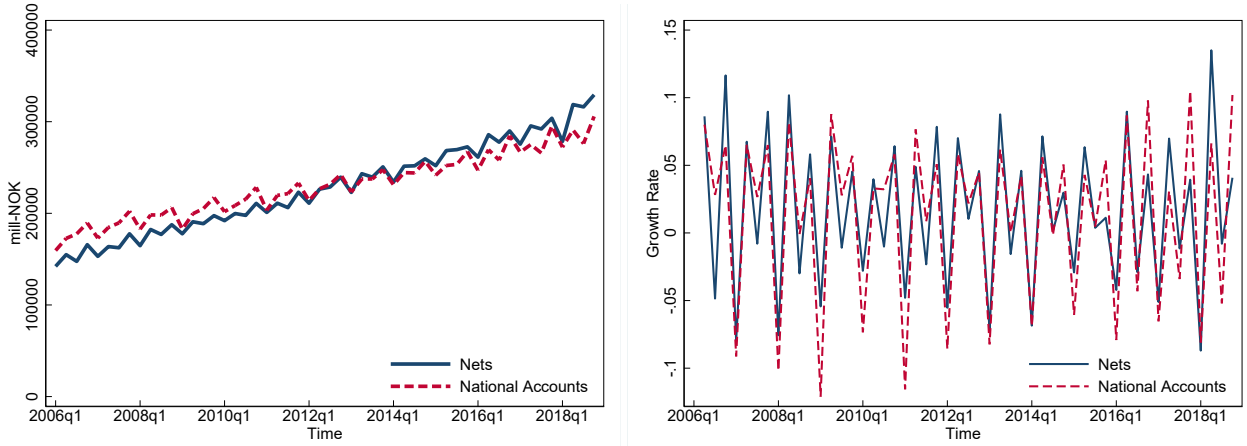

Notes: The figures compare the level of consumption from our data and the national accounts, and the quarterly growth rates from 2006 to 2018. Our data has attrition in bank account links before 2012.

3.3 Descriptive statistics

We first consider a simple variance decomposition of consumption, splitting consumption into between firm and within firm components. Panels A and B from Figure 3 show the proportion of total variation attributable to both components over the six years of available data. Around a quarter of variation in log wage income is explained by differences across firms. We observe a slightly smaller role for the between firm component for log expenditure at around a fifth.

While the overall measure of consumption is important, understanding the features of more granular categories on expenditures is also of interest because different categories may be more (food expenditure) or less (car expenditure) worrying from a welfare perspective. In the remaining panels of Figure 3 we show variance decompositions across expenditure on food, clothing, furnishing, cars, hospitality, and restaurants.

There are large differences in the proportion of overall variation explained by the between firm component across categories. For example, variation in the between firm component explains just 5 percent of the overall variation in car expenditure, whereas variation in the between firm component explains around 15 percent of the variation in food expenditure. While these results are descriptive in nature, these numbers suggest that the role of firms for consumption inequality may be multifaceted, where firms matter more for some consumption categories than others.

Figure 3: Variance decomposition plots: household wage and consumption categories.

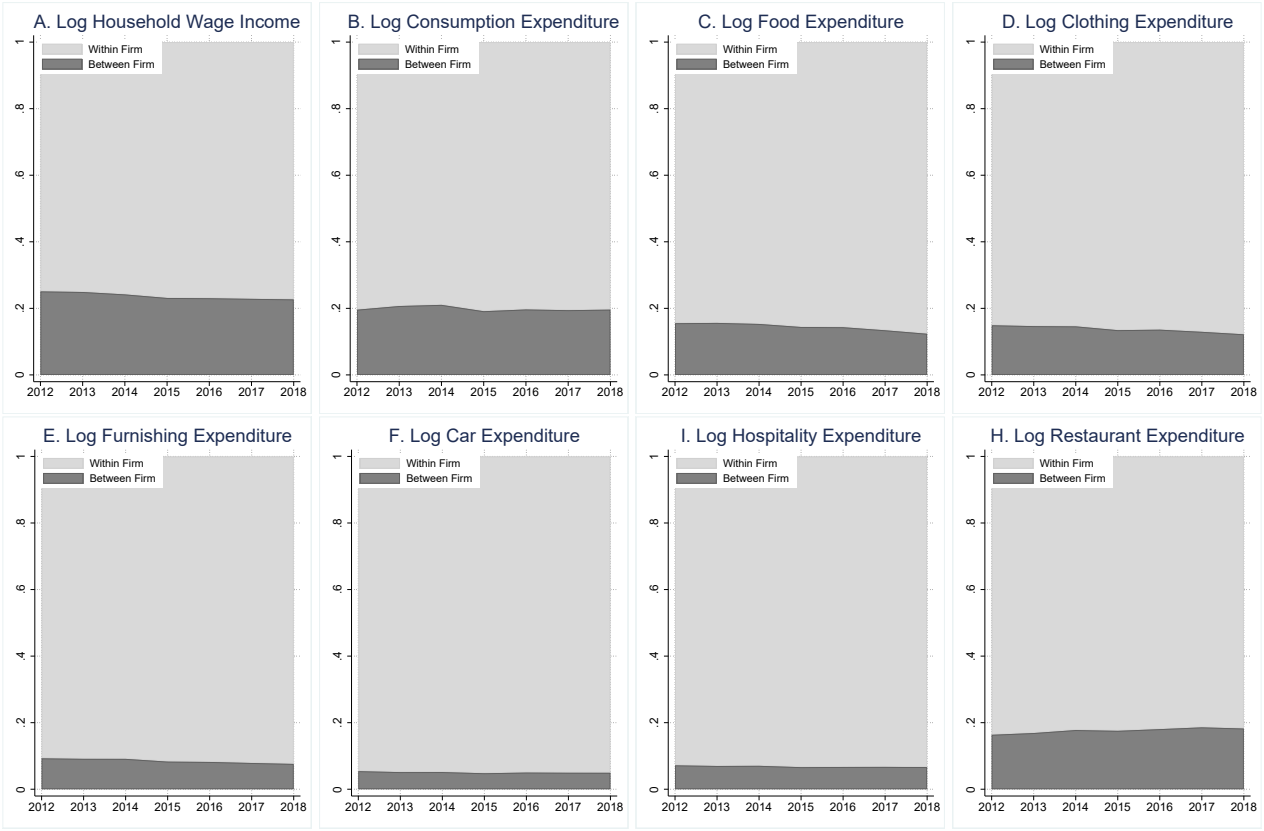

Notes: These figures report normalized variance decomposition of wages and components of consumption for singles.

4 Co-Worker Analysis

This section examines the movers design in the context of co-workers and the associated consumption changes.

4.1 Movers Design

Before turning to our statistical variance decomposition, we gauge the assumptions needed for a causal interpretation of the log-additive worker and firm effects in equation (5). To avoid a mechanical relationship between the firm effect, which is identified from movers, and the post-move wage change, we relate the wages of movers to the average log wage of coworkers. We measure the average wage of workers in origin and destination firm who do not move (i.e., the stayers).

We are interested in assessing the conditional orthogonality – or exogenous mobility assumption. This assumption relates the timing of a move to other shocks that may affect wages, such as health shocks or positive productivity shocks. A differential trend in earnings just prior to a lower- or higher-paying firm could reflect a potential violation of this assumption. We also rely on a parallel trend assumption; the potential wages of a mover in the departing firm does not change around the move date. We also note that such an interpretation does not follow directly if we were to perform a similar analysis around changes in co-worker consumption.

To perform the co-worker analysis, we restrict the sample of movers with at least three years of experience who move between two firms from 2012 to 2016. We include singles and married, but excludes households where more than one partner moves in a given year. Next, we divide the sample by quartiles of pay in the origin and destination firm, and report moves from the top and bottom. The left panel explores wage changes for movers, and the right panel explores the associated consumption changes for the same set of moves.

In the left panel of figure 4, we plot the evolution of wages around moves from the first quartile of co-worker wages to another firm in the first, second, third or top quartile of firm co-worker pay. Similarly, we plot the evolution of workers who depart a firm from the top quartile to other firms. Notably, we that the downward and upward moves, e.g., from "4 to 1" and from "1 to 4" appear symmetric and very similar in absolute change.

Figure 4: Wage and consumption change with job switching

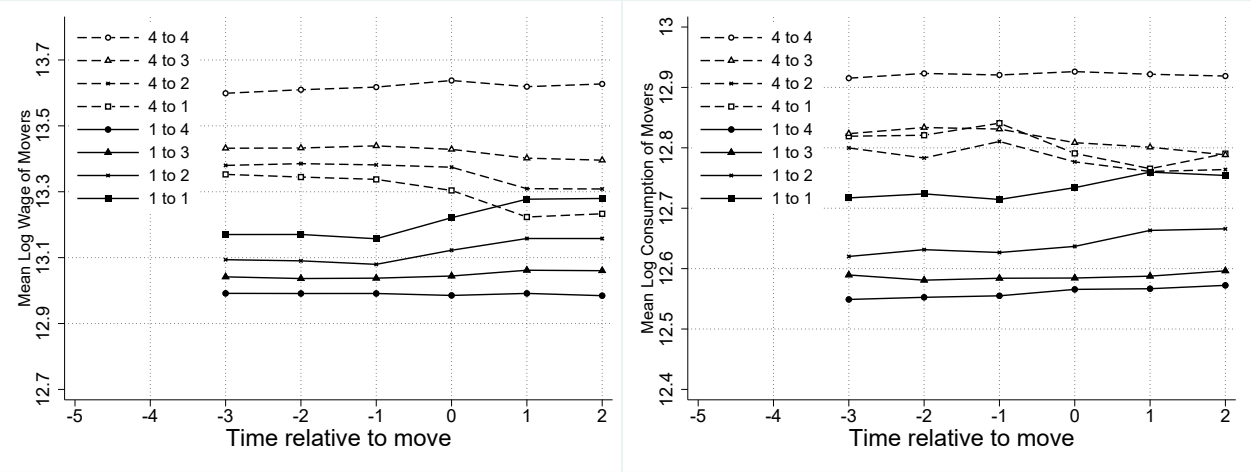

4.2 Consumption Elasticities

The moves cause changes in both wage income and consumption. Next, we perform a regression analysis to gauge the statistical precision and calculate implied consumption elasticities. To do so, we start by regressing log disposable income two years after the move on the change in log of average co-worker wages. Our measure of disposable income includes wage income, capital income, any government transfers and subtracts off taxes paid. We control for the average log disposable income before the move and dummies for each year, and cluster standard errors at the regional level. We find that the change in income is statistically significant and positive. Panel A of Table reports the coefficient estimate, standard error, implied t-statistic, dependent mean and the number of observations for each outcome. On average, the reduced-form regression shows that disposable income increases by 4.6 percent when moving to a firm with higher co-worker wages. We plot the coefficient estimates three years before and up to two years after in Appendix Figure D2.

Next, we regress the log of consumption expenditure on the change in co-worker wage, using the same specification as above. Our headline result is that total expenditure increases by 1.8, and that this effect is statistically singificant, confirming the graphical evidence from the event study graph. To calculate the implied elasticity of consumptiom, we divide the effect of coworkers on consumption by the effect of co-workers on disposable income. We find that the consumption elasticity respect to changes in income is almost 40%, indicating that workers consume a substantial fraction of their increase in income.

To better understand the nature of the consumption effects of moving to higher-paying firms, we regress the different components of log spending, starting with food and clothing. The first row of Panel B shows that food expenditure is non-responsive, with a point

Table 3: Reduced Form Evidence on Coworker Effects

	A. Overall Income and Spending				
	Estimate	(St. Error)	t-stat	Dep. Mean	N. obs
Log Net Income	0.046	0.007	6.930	13.570	82,441
Log Total Expenditure	0.018	0.007	2.632	13.316	82,525
	B. Categories of Spending				
	Estimate	(St. Error)	t-stat	Dep. Mean	N. obs
Log Food Expenditure	0.010	0.009	1.088	11.305	82,336
Log Clothing Expenditure	0.024	0.013	1.836	9.740	80,647
Log Furnishing Expenditure	0.038	0.014	2.799	9.666	81,169
Log Motor Vehicles	-0.007	0.026	-0.275	8.573	64,897
Log Restaurant Expenditure	0.048	0.008	5.890	9.668	81,691
Log Hotel Expenditure	0.036	0.023	1.572	7.411	66,801

Notes: The table reports estimates from regressions of log income or log consumption on the log change in coworker wages. Each regression controls for dummies for years and the average of the outcome variable before the year of the move. Standard errors are clustered at the regional level.

estimate that is close to one and not statistically significant different from zero. The second row reports a statistically significant impact spending on clothing, which is similar to the overall spending effect. In the third row, we find that the effect on spending on furnishing is statistically significant and quite elastic. In contrast, we find a smaller and not statistically significant different estimate on car purchases. Besides furnishing, the largest impact of coworker wages is on restaurant expenditure. The evidence is consistent with non-homothetic preferences. Lastly, we find an imprecise but similar effect size for hotel expenditure.

5 Variance- and Covariance Analysis

This Section present our main empirical variance decomposition results. Guided by our quantitative analysis in Section 2, we consider first the relative size of the variance of firm consumption effects and firm wage effects. Then we consider the relationship between person, and firm, consumption and wage ranks. Finally, we consider whether, consistent with models, the consumption firm effect is better correlated with value ranks as considered in Sorkin (2018).

5.1 Variance components

Table 4, columns 1 and 3, show the results from naive variance components not adjusting for limited mobility bias for log earnings and log consumption, respectively. Columns 2 and 4 show the results when adjusting for limited mobility bias following Kline et al. (2020).[14]

There are several takeaways. First, we find that, consistent with models of individual saving behavior, the variance of the firm effect in log earnings is larger (roughly double) than the variance of the firm effect in log consumption

However, the variance in firm effects on consumption is still substantial. A key empirical result, implied by these variance components, is that a one standard deviation increase in firm consumption effects is associated with approximately a 11 percent increase in consumption.[15] Thus, our analysis clearly shows that where a worker works appears to play an important role in how much they consume.

An obvious question is how much of this 11 percent increase in consumption is a result of increases in wages? To consider this we calculate the wage gain associated with a standard deviation increase in firm consumption effects. We do so by estimating

$$\psi_j^w = \alpha + \beta \frac{\psi_j^c}{\text{std}(\psi^c)} + \varepsilon \tag{9}$$

and retrieve a value for the estimated coefficient on the standardized firm consumption effect of $\hat{\beta} = 0.058$. This can be interpreted as wage gain of 6 percent from standard deviation increase in ψ_j^c. Hence, moving a worker up a standard deviation in the firm consumption effects results in a wage increase around half of that of the associated consumption increase.

Second, the correlation of the worker and firm effects for both earnings and consumption are positive in the bias corrected estimates. Thus, for earnings, our results corroborate a recent literature finding substantial sorting on worker and firm characteristics, suggesting high wage workers are more likely to work in high wage firms (Bonhomme et al., 2023). It is of equal interest whether high consumption workers also sort into high consumption firms. We find a slightly smaller correlation but similarly positive relationship between worker and firm effects for consumption. Thus, high consumption workers are also more likely to work in high consumption firms, but the strength of this sorting is weaker.

[14] The decomposition in Table 4 is performed on a sample of single individuals. In the Appendix we also report similar decompositions on the full sample of singles and couples (married and co-habitants) in Table C2, and on the same full sample but aggregating wages and consumption to the household level in Table C6. The results in the main Table 4 are shown to be robust to these alternative ways of sampling and measuring the data.

[15] This is calculated as $\exp(\text{Var}(\psi_{j(i,t)})^{1/2}) - 1$.

Finally, while we focus on considering the role of firms for consumption inequality in the present paper, we observe a strikingly large role for worker heterogeneity. Consistent with other work in the literature, the role of workers in earnings variation is substantial, suggesting substantial differences in the productive capacity of individuals. The variance of worker effects in consumption is much larger. Thus, a large part of the larger variance in log consumption compared to log wages is because workers appear to consume substantially differently. Again, while we don't focus on worker heterogeneity in the present paper, considering the implications of the extensive heterogeneity in consumption behavior as we estimate here is an exciting path for future research.[16]

Table 4: Variance contributions for the variance of log earnings and log expenditure

	log earnings		log expenditure	
	AKM	KSS	AKM	KSS
signal variance	0.176	0.154	0.497	0.454
variance of firm effects	0.037 (0.212)	0.020 (0.128)	0.032 (0.065)	0.011 (0.023)
variance of person effects	0.147 (0.835)	0.113 (0.731)	0.485 (0.977)	0.427 (0.941)
covariance of firm and person effects	-0.004 (-0.048)	0.011 (0.141)	-0.011 (-0.042)	0.008 (0.036)
correlation of worker and person effects	-0.057	0.230	-0.084	0.120
R squared	0.802	0.702	0.793	0.726

Notes: Variance decomposition of log earnings and log expenditures for the singles sample on the leave-one-out connected set. AKM refers to the standard AKM decomposition as in Abowd et al. (1999), while KSS is the bias-corrected estimates following Kline et al. (2020). Numbers in parenthesis refers to fraction of explained variance

[16]This dovetails with a large body of work considering the measurement and importance of heterogeneity in marginal propensities to consume (e.g., Johnson et al. (2006); Parker et al. (2013); Kaplan et al. (2018); Auclert et al. (2019)).

5.2 Firm wage and consumption effects

When firms are indexed by productivity and other characteristics that are relevant for consumption, the correlation between firm effects in consumption and pay will be lower than one. How much lower than one depends on how dispersed the underlying distributions of risk vs. productivity are. While our focus in the present paper is the firm, we note that the model implies a perfect correlation between wage and consumption effects of workers because they do not differ along other dimensions.[17]

We start by exploring the empirical relationship between firm pay and consumption effects. The left panel of Figure 5 shows the rank-rank correlations between firm log consumption and firm log earnings effects. The rank-rank correlation between firm effects is only 0.27; positive, yet substantially below 1. Hence, high wage firms are also more likely to be high consumption firms but the relationship is far from perfect. Through the lens of the model, the evidence suggests that firms differ along dimensions beyond just the wage, where these further dimensions matter for how much workers consume. Note that while we considered a model with job destruction heterogeneity is the key characteristic that matters for worker consumption beyond the wage, there could be many different facts that shape consumption risks associated with a firm, such as risk of work injury, disability or that some firms are associated with higher costs of job loss due to specific human capital.

The figu There is a stronger correlation between consumption and pay effects of workers than for firms. Through the context of the model discussed in Section 2, there is dispersion in permanent consumption across workers only because of differences in ability. The evidence suggests high wage workers are also more likely to be high consumption workers. Thus heterogeneity in other characteristics orthogonal to worker ability that impact how much workers consume, may be less important for workers than firms.

5.3 Firm values and firm effects

Motivated in part by research on non-pay characteristics of firms valued by workers, a recent literature has proposed several approaches to rank firms (Sorkin, 2018). These approaches use a revealed preference argument that captures unobserved characteristics of a job that is usually not observed in the data, but relevant for the overall value of a firm. We showed in Section 2 that firm consumption effects should be a better measure of the value of a job than firm wage effects, since consumption relates to the net present value of a job and captures factors of jobs that are relevant for consumption such as unemployment risk, injury risk

[17] Workers could differ, for example, in their marginal propensity to consume due to credit- or liquidity constraints, or they could differ in patience.

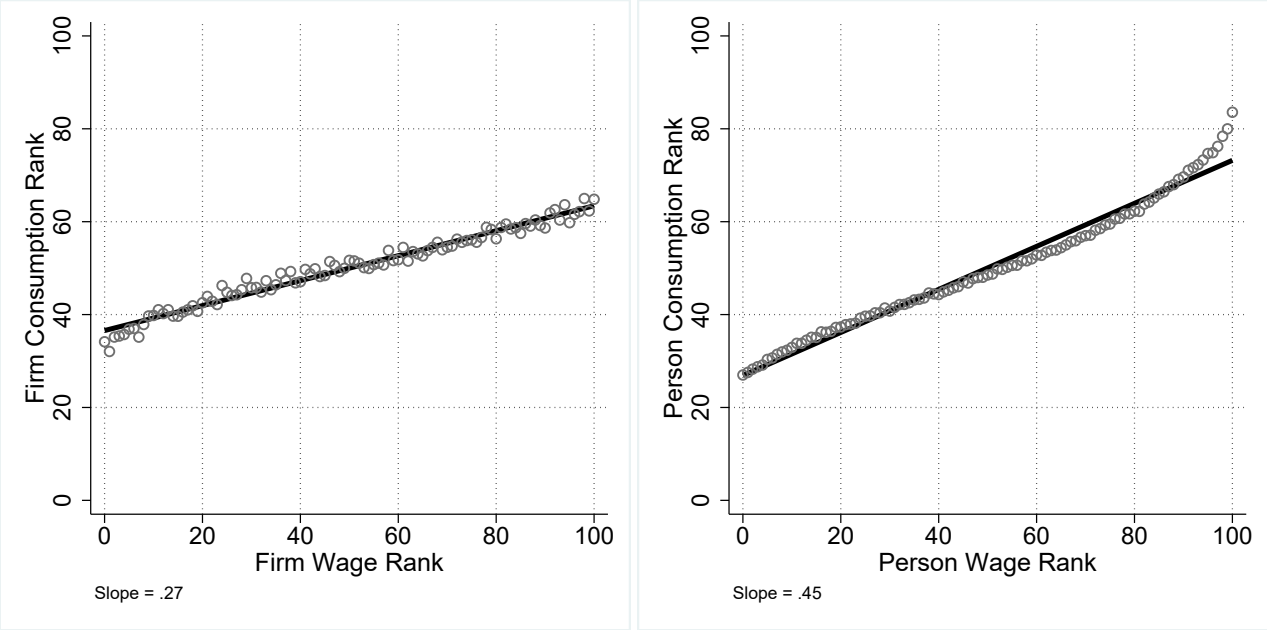

Figure 5: Rank-rank correlations between person log consumption and person log earnings effects (left), and firm log consumption and firm log earnings effects (right)

etc. We showed, however, that when jobs are also indexed by separable amenities, that the ranking between firm values and consumption effects becomes strictly less than one. Thus, the difference between the correlation of firm value ranks with firm consumption rank and firm value ranks with firm wage ranks, illuminates whether amenities in a job are relevant for a workers consumption or not.

To rank firms, we first use the revealed preference approach outlined by Sorkin (2017). Let the flow utility in a match be written as

$$v_j + e_{ij}, \tag{10}$$

where v_j is a common value of the firm and e_{ij} is a type-1 extreme value distributed random component, which may resemble preference heterogeneity. This setup, then allows for vertical and horizontal heterogeneity, such that some workers may travel from high to low value firms. To generate a ranking of firms, Sorkin (2017) iterates on the following equation to derive the common value component of the firm k based only on observed worker flows across firms,

$$\exp(v_k) = \sum_j \frac{\text{flows}_{jk} \exp(v_j)}{\text{flows}_{kj}}, \tag{11}$$

where flows_{jk} is the number of workers moving from firm j to firm k, and flows_{kj} is the number of workers moving from firm k to firm j. The numerator on the right is thus the

value weighted entry, with the number of exits in the denominator. Conceptually, high value firms as those that hire from other high value firms and where few workers leave. Results based on a simple implementation of the Sorkin (2018) page rank, assuming firms are of equal size, yields a higher rank-rank correlation of firm values on firm consumption effects. This is consistent with the models in Section 2 where firms matter for consumption choices beyond the wage. However, we are still in the process of estimating value ranks when we allow for firms to differ in size and offer intensity.

6 Conclusion

Despite extensive evidence on pay disparities, relatively little is understood about the extent to which firms influence worker consumption. In this paper, we aim to bridge this knowledge gap by constructing a novel dataset that links workers' consumption patterns to their employment histories. To guide our analysis, we develop a theoretical model featuring risk-averse workers and heterogeneity on both the worker and firm sides. Our model allows us to make several predictions for the characteristics of firm consumption effects. Using the dataset, we estimate two-way fixed effects for wages and consumption. Our findings suggest that firms matter for how much workers consume. Notably, a one-standard-deviation increase in firm consumption effects corresponds to an 11% increase in consumption but only a 6% increase in wages. We also observe that the correlation between firm pay and consumption effects is approximately 25%. Additionally, we show that a consumption-based index of firms aligns more closely with standard revealed preference-based firm rankings. These results shed light on the significant role firms play in shaping welfare inequality.

References

ABOWD, J., R. CREECY, AND F. KRAMARZ (2002): "Computing person and firm effects using linked longitudinal employer-employee data," Tech. rep., Center for Economic Studies, US Census Bureau.

ABOWD, J. M., F. KRAMARZ, AND D. N. MARGOLIS (1999): "High wage workers and high wage firms," *Econometrica*, 67, 251–333.

ANDERSEN, A. L., A. S. JENSEN, N. JOHANNESEN, C. T. KREINER, S. LETH-PETERSEN, AND A. SHERIDAN (2023): "How Do Households Respond to Job Loss? Lessons from Multiple High-Frequency Datasets," *American Economic Journal: Applied Economics*, 15, 1–29.

ATTANASIO, O. P. AND L. PISTAFERRI (2016): "Consumption inequality," *Journal of Economic Perspectives*, 30, 3–28.

AUCLERT, A., W. S. DOBBIE, AND P. GOLDSMITH-PINKHAM (2019): "Macroeconomic effects of debt relief: Consumer bankruptcy protections in the great recession," Tech. rep., National Bureau of Economic Research.

BAGGER, J. AND R. LENTZ (2019): "An empirical model of wage dispersion with sorting," *The Review of Economic Studies*, 86, 153–190.

BAKER, S. R. (2018): "Debt and the response to household income shocks: Validation and application of linked financial account data," *Journal of Political Economy*, 126, 1504–1557.

BLOOM, N., F. GUVENEN, B. S. SMITH, J. SONG, AND T. VON WACHTER (2018): "The disappearing large-firm wage premium," in *AEA Papers and Proceedings*, American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, vol. 108, 317–322.

BLUNDELL, R., L. PISTAFERRI, AND I. PRESTON (2008): "Consumption Inequality and Partial Insurance," *The American Economic Review*, 98, pp. 1887–1921.

BONHOMME, S., K. HOLZHEU, T. LAMADON, E. MANRESA, M. MOGSTAD, AND B. SETZLER (2023): "How much should we trust estimates of firm effects and worker sorting?" *Journal of Labor Economics*, 41, 291–322.

BONHOMME, S., T. LAMADON, AND E. MANRESA (2019): "A distributional framework for matched employer employee data," *Econometrica*, 87, 699–739.

BURDETT, K. (1978): "A theory of employee job search and quit rates," *The American Economic Review*, 68, 212–220.

CARD, D. (2022): "Who Set Your Wage?" *American Economic Review*, 112, 1075–90.

CARD, D., A. R. CARDOSO, J. HEINING, AND P. KLINE (2018): "Firms and labor market inequality: Evidence and some theory," *Journal of Labor Economics*, 36, S13–S70.

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., J. HEINING, AND P. KLINE (2013): "Workplace heterogeneity and the rise of West German wage inequality," *The Quarterly journal of economics*, 128, 967–1015.

CRANE, L. D., H. R. HYATT, AND S. M. MURRAY (2023): "Cyclical labor market sorting," *Journal of Econometrics*, 233, 524–543.

DE GIORGI, G., A. FREDERIKSEN, AND L. PISTAFERRI (2020): "Consumption network effects," *The Review of Economic Studies*, 87, 130–163.

DEATON, A. (1992): *Understanding consumption*, Oxford University Press.

EECKHOUT, J. AND P. KIRCHER (2011): "Identifying sortingâ€"in theory," *The Review of Economic Studies*, 78, 872–906.

ELSBY, M. W. AND A. GOTTFRIES (2022): "Firm dynamics, on-the-job search, and labor market fluctuations," *The Review of Economic Studies*, 89, 1370–1419.

FRIEDMAN, M. (1957): "A theory of the consumption function," *General series National Bureau of Economic Research (*.

GANONG, P. AND P. NOEL (2019): "Consumer spending during unemployment: Positive and normative implications," *American Economic Review*, 109, 2383–2424.

GELMAN, M., S. KARIV, M. D. SHAPIRO, D. SILVERMAN, AND S. TADELIS (2014): "Harnessing naturally occurring data to measure the response of spending to income," *Science*, 345, 212–215.

GERARD, F. AND J. NARITOMI (2019): "Job displacement insurance and (the lack of) consumption-smoothing," Tech. rep., National Bureau of Economic Research.

GOTTFRIES, A. (2018): "Bargaining with renegotiation in models with on-the-job search," *Available at SSRN 3058870*.

JAROSCH, G. (2023): "Searching for job security and the consequences of job loss," *Econometrica*, 91, 903–942.

JOHNSON, D. S., J. A. PARKER, AND N. S. SOULELES (2006): "Household expenditure and the income tax rebates of 2001," *American Economic Review*, 96, 1589–1610.

KAPLAN, G., B. MOLL, AND G. L. VIOLANTE (2018): "Monetary policy according to HANK," *American Economic Review*, 108, 697–743.

KLINE, P., R. SAGGIO, AND M. SØLVSTEN (2020): "Leave-out estimation of variance components," *Econometrica*, 88, 1859–1898.

KLINE, P. M. (2024): "Firm Wage Effects," *NBER Working Papers*.

LAMADON, T., M. MOGSTAD, AND B. SETZLER (2022): "Imperfect competition, compensating differentials, and rent sharing in the US labor market," *American Economic Review*, 112, 169–212.

LISE, J. (2013): "On-the-job search and precautionary savings," *Review of Economic Studies*, 80, 1086–1113.

MODIGLIANI, F. AND R. BRUMBERG (1954): "Utility Analysis and the Consumption Function: An Interpretation of Cross-Section Data," Post-Keynesian Economics/Rutgers University Press.

MORTENSEN, D. T. (1982): "Property Rights and Efficiency in Mating, Racing, and Related Games," *American Economic Review*, 72, 968–979.

PARKER, J. A., N. S. SOULELES, D. S. JOHNSON, AND R. MCCLELLAND (2013): "Consumer spending and the economic stimulus payments of 2008," *American Economic Review*, 103, 2530–2553.

ROSEN, S. (1986): "The theory of equalizing differences," *Handbook of Labor Economics*, 1.

SHIMER, R. (2005): "The cyclical behavior of equilibrium unemployment and vacancies," *American economic review*, 95, 25–49.

——— (2006): "On-the-job search and strategic bargaining," *European Economic Review*, 50, 811–830.

SHIMER, R. AND L. SMITH (2000): "Assortative matching and search," *Econometrica*, 68, 343–369.

SORKIN, I. (2017): "Ranking firms using revealed preference," *The Quarterly Journal of Economics*.

——— (2018): "Ranking firms using revealed preference," *The Quarterly Journal of Economics*.

THE DEPARTMENT OF THE TREASURY (2022): "The State of Labor Market Competition," Tech. rep., US Treasury.

Appendix

A Model appendix

A.1 Wages

Wage setting is challenging in a model with savings and on-the-job search. Firms and workers have to take into account the distribution of assets in the economy when bargaining. On top of this typical bargaining methods render non-convex bargaining sets (Shimer, 2006; Gottfries, 2018).

We follow the bargaining protocol of Elsby and Gottfries (2022). Bargaining takes place fully ex-post with continual renegotiation. This implies that worker future values no longer matter, and workers and firms bargain over the flow surplus. Let the production function, q, be written as

$$q(p, h) = ph. \tag{A1}$$

We describe this in more detail in the Appendix, but given this log-additively separable production function, the result of this bargaining protocol is a wage equation of the form

$$w(p, h) = \beta ph + \omega_0, \tag{A2}$$

where β is worker bargaining power, and ω_0 collects terms describing the cost of disruption of negotiation, which in principle can be arbitrarily small. Given a small value for ω_0 we reach a wage equation consistent with the AKM structure.[18]

A.2 Hamilton-Jacobi-Bellman (HJB) equations

The features of the model can be condensed into recursive form which we outline here. We first describe the HJB equations for the workers and then the firms. Let V^e and V^u be the present discounted value of employment and unemployment respectively. Let a firm-type vector, $\Delta^f = (p, \delta, \vartheta)$, where p is firm productivity, δ is firm job destruction risk, and ϑ is the amenity value of the job. Similarly, a worker-type vector, $\Delta^w = (a, h)$, where a is the asset level and h is worker productivity. Define the threshold firm vector, Δ^f_r, that solves

$$V^e(\Delta^f_r, \Delta^w) := V^e(\Delta^f, \Delta^w). \tag{A3}$$

[18] While we acknowledge that search models do not always render an AKM structure in wages, recent advances have argued that a log-additive specification in wages is a good approximation (Card et al., 2013).

The HJB equation for employment for a particular firm-type vector and worker-type vector is written as

$$\rho V^e(\Delta^f, \Delta^w) = \max_c u(c, \vartheta) + \lambda^e \int_{\Delta^f_r} V^e(x, \Delta^w) - V^e(\Delta^f, \Delta^w) dF(x) + \delta[V^u(\Delta^w) - V^e(\Delta^f, \Delta^w)] + \dot{a} V^e_a(\Delta^f, \Delta^w). \tag{A4}$$

An employed worker chooses their consumption to maximize the present discounted value. This is composed of a job finding component. A job destruction component. And a component attributable to change in assets.

Similarly for the unemployed, define the threshold firm vector, Δ^f_u, that solves

$$V^e(\Delta^f_u, \Delta^w) \coloneqq V^u(\Delta^w). \tag{A5}$$

The HJB equation for unemployment for a particular worker-type vector is written as

$$\rho V^u(\Delta^w) = \max_c u(c, \vartheta) + \lambda^u \int_{\Delta^f_u} V^e(x, \Delta^w) - V^u(\Delta^w) dF(x) + \dot{a} V^u_a(\Delta^w). \tag{A6}$$

An unemployed worker chooses their consumption to maximize the present discounted value. This is composed of a job finding component. And a component attributable to change in assets.

Firms make no choices other than wage determination. Firm values do not impact the wage bargaining game, but we outline the associated HJB equation here for completeness. Let V^f be the presented discounted value of a match to a firm, and assume the value of a vacancy is zero. The associated HJB equation is written as

$$(\rho + \lambda^e[1 - F(\Delta^f_r)]) V^f(\Delta^f, \Delta^w) = q(h, p) - w(p) + \dot{a} V^f_a(\Delta^f, \Delta^w). \tag{A7}$$

The firm has linear preferences, gaining returns $q(h, p) - w(p)$, and enjoys disutility from workers leaving to other firms at rate $\lambda^e[1 - F(\Delta^f_r)]$.

A.3 Ranking firms through the lens of the model

In this section we describe two prominent methods to rank firms in the data. Of course, if we simply wanted to rank firms in the model we could use the value functions when solving the model, eliminating any simulation noise. We are interested, however, in considering a theory consistent ranking of firms in the data. Thus we conduct these ranking methods on

the simulated data from the model. Borrowing notation from Bagger and Lentz (2019), the poaching rank for firm j, PR_j, is given as

$$PR_j = \frac{\text{entry from only employment}}{\text{entry from employment and unemployment}}, \tag{A8}$$

the proportion of hires that originate from employment. The intuition of this is straight forward. Consider firms ranked by their value. Firms higher up in the rank distribution are more likely to poach employers than workers in the lower part of the rank distribution. Indeed, the firm at the reservation value will only hire firms out of unemployment.From the model above, and conditional on an asset level, the poaching rank is given as

$$\pi_j = \frac{\left(G(\Delta_r^f)\lambda_e + \lambda_r\right)(1-u)}{\left(G(\Delta_r^f)\lambda_e + \lambda_r\right)(1-u) + \lambda_u u}, \tag{A9}$$

where $G(\Delta_r^f)$ is the proportion of the employed in firms with value less than Δ^f, and u is the unemployment rate. It is straightforward to see that as V increases, the numerator and left term in the denominator increases up to $\left(\lambda^e + \lambda^r\right)(1-u)$. While, assuming the job with value V is acceptable by the unemployed, the term on right hand side of denominator does not change as we move up or down the firm value rank. There is a slight issue with the poaching rank in that firms may disagree slightly on which firm is preferable based on their asset level.

An alternative and influential method to rank firms was posited by Sorkin (2017). Through a type 1 extreme value distribution assumption over "preference shocks", Sorkin relates worker values to the flow of workers into and out of a firm. Intuitively, better firms hire many workers, and have few leave. Since we do not incorporate preference shocks in the model, we focus on firm ranks and so consider the Sorkin rank for firm j, S_j, as

$$S_j = \frac{\text{entry from other firms}}{\text{exits to other firms}}. \tag{A10}$$

Through the lens of the model, with identical notation as for the poaching rank, this rank measure takes the following form

$$S_j = \frac{\left(G(\Delta_r^f)\lambda_e + \lambda_r\right)}{\left(\lambda_e[1 - F(\Delta_r^f)] + \lambda^r\right)} \frac{f(\Delta^f)}{g(\Delta^f)}. \tag{A11}$$

Analogous to Sorkin (2017), the rank measure includes the ratio of the inflow rate per offer density, and the outflow rate per firm size, multiplied by the ratio of the offer density and

firm size. Focusing on the first ratio, it is clear that this measure produces a perfect ranking of firms. As V increases, $G(\Delta_r^f)$ increases and $1-F(\Delta_r^f)$ declines. As in Sorkin (2017), in order to rank firms, we require measures of firm offers and firm size, which we can, of course, compute with the model.

A.4 Moving parameter values

This subsection considers how the measures of interest discussed in Section 2 depend on the parameter values. We focus on the parameter values most associated with a workers labor market outcomes: the job offer arrival rate in unemployment and employment, and the construction of the offer distributions.

While we will provide an exhaustive set of results over different parameter values, we cannot claim with certainty that some combination of parameter values would deliver results that would not be within our predictions in Section 2, however, we believe we consider parameter values in a sensible region of the data thus to provide reassurance for the predicted outcomes.

Lise (2013) with worker heterogeneity

Figure A1: Changing λ^u and λ^e in Lise (2013)

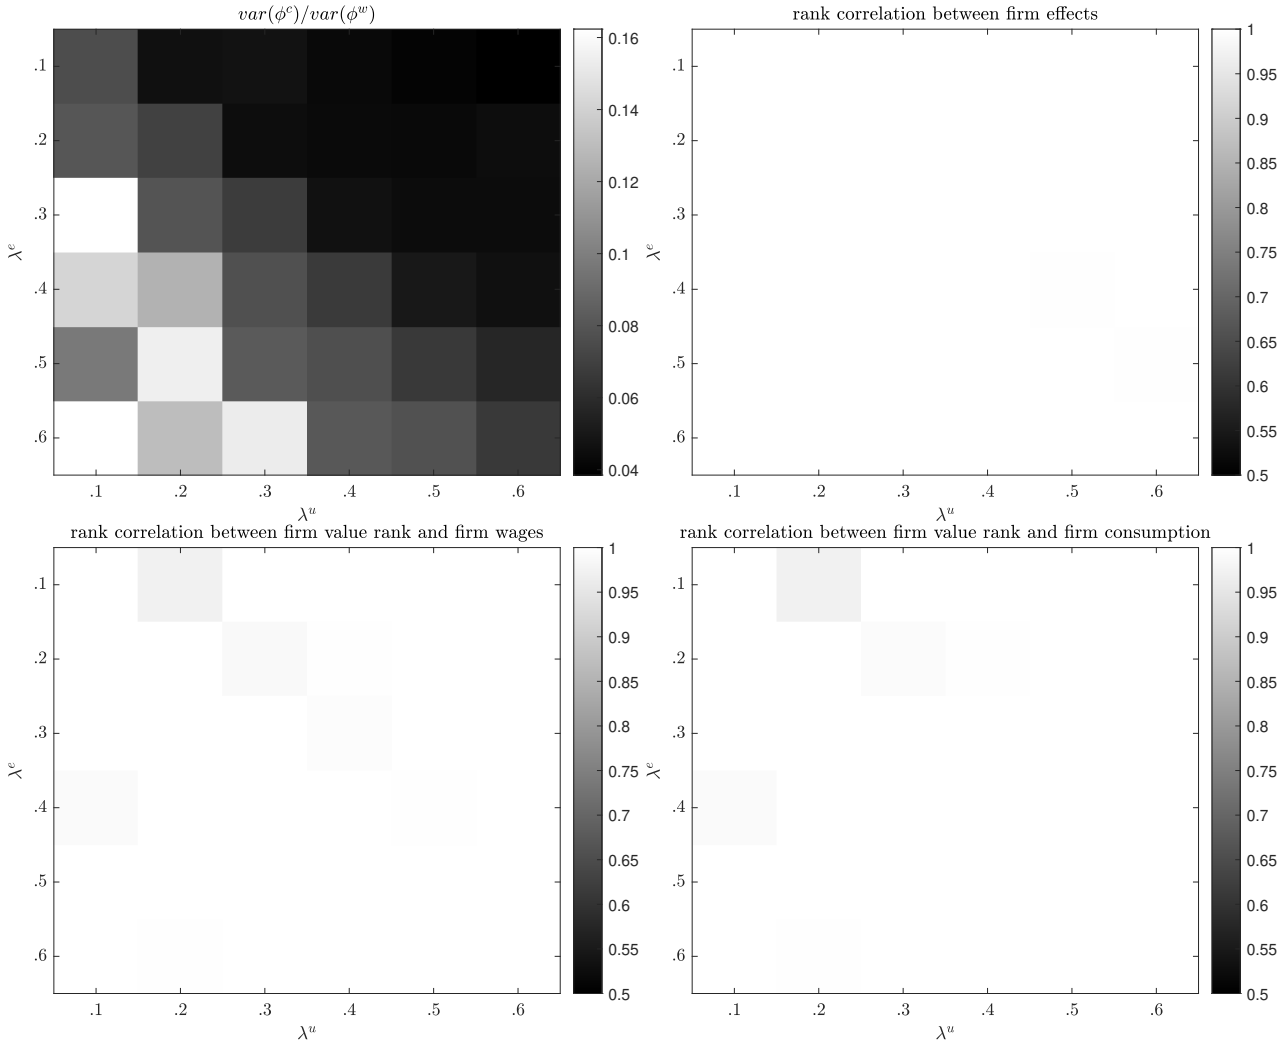

Lise (2013) + Jarosch (2023) with worker heterogeneity

Figure A2: Changing the productivity and job destruction distributions

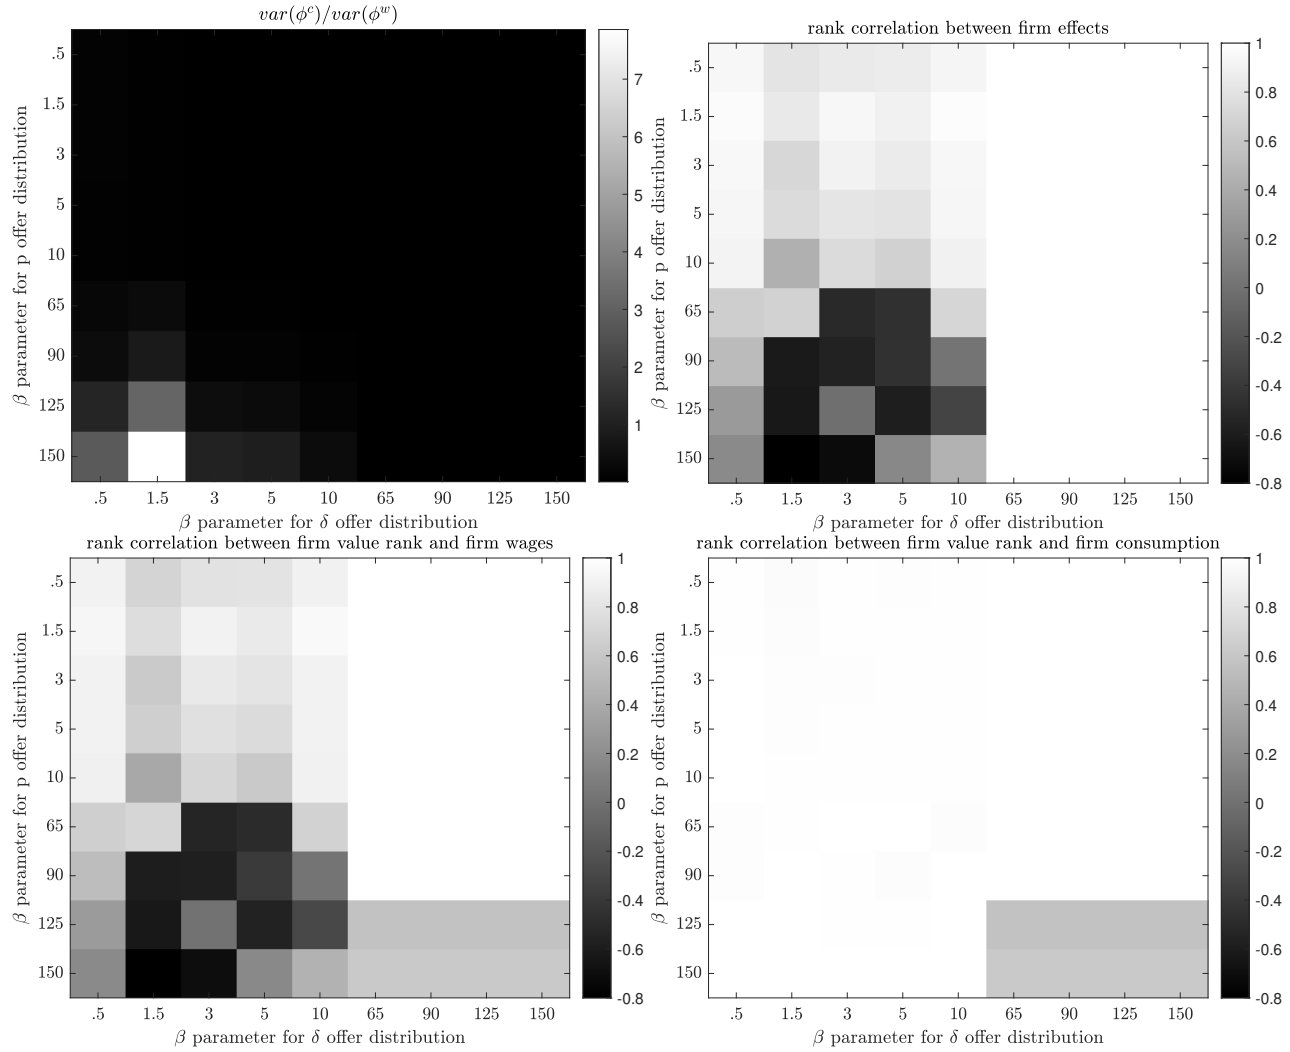

Ranking firms

Figure A3: Changing the dispersion of wages

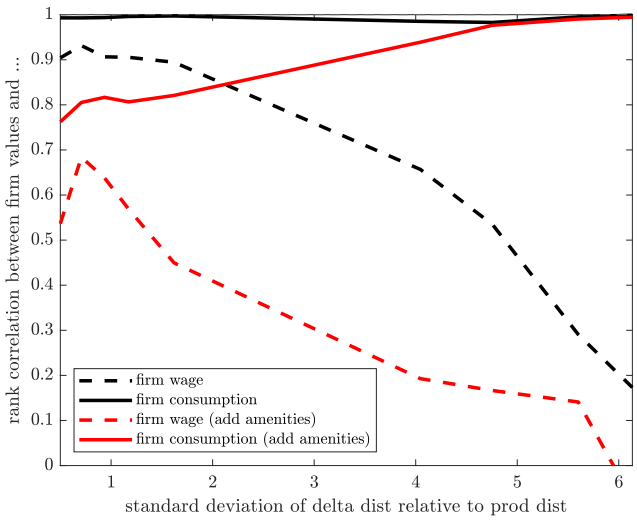

B Consumption Expenditure Data

B.1 Electronic Payments: Data Construction

The electronic payments data are provided by the Norwegian retail clearing institution, Nets Branch Norway. The data covers three data sources: (i) debit card payments via BankAxept made on physical card terminals in Norway (ii) online wire payments cleared via NICS, and (iii) all income payments cleared via NICS. The data includes all Norwegian individuals that have made or received an electronic payment using BankAxept or via the NICS clearing over the period 2006 to 2018. For each individual transaction, Nets Branch Norway stores a set of meta-data in addition to the transaction amounts. For debit card payments, the additional data is among other things the Merchant Category Code (MCC)[1], physical address (zip code) of the card terminal, and the business name and NACE Rev. 2 of the merchant. Payments via the NICS clearing system comes with information on the creditor/debtor's NACE Rev. 2 and location (zip code).

The data provided by Nets Branch Norway is aggregate to person-week-zip-consumer category level. The consumer categories are based on COICOP (UN classification system which the Norwegian National Accounts are based on), and there are 12 broad categories and 24 categories in total. In addition, there are two categories (13 and 14) which applies only to payments via NICS. Table B.1 describes all consumption categories.

Since COICOP classifications are not part of the meta-data stored by Nets Branch Norway, Norges Bank provided a crosswalk between each MCC and COICOP, and each 5 digit NACE and COICOP. The consumer categorization is based on these crosswalks. Among existing physical card terminals, 42,600 out of about 162,000 were classified with MCC "Miscellaneous and Specialty Retail Stores" (MCC = 5999) or had a missing MCC. These were classified using the registered information on the terminal, typically the name and NACE-code of the store.[2]

B.2 BankAxept Debit Card Payments

The debit card data contains the universe of individual's debit card payments via BankAxept from 2012 and onward. Prior to 2012 the coverage shrinks somewhat the further back we go. The reason is that Nets Branch Norway did not retain information on the account owner for accounts that were closed prior to 2012. Most of this is related to individuals who's accounts

[1] MCC is a four-digit number listed in ISO 18245 for retail financial services and is used to classify a business by the types of goods or services it provides.

[2] All cross-walks are available upon request.

Table B.1: Consumption Categories

	Category
01	Food and non-alcoholic beverages
02	Alcoholic beverages, tobacco and narcotics
03	Clothing and footwear
04	Utilities, electricity, gasoline, housing rent
05	Furnishings, household equip. and routine household maintenance
06	Health
07	Transport
	071 Purchase of vehicles
	072 Operation of personal transport equipment
	073 Transport services
08	Communications
09	Recreation and culture
	091 Audio-visual, photographic and information processing equipment
	092 Major durables for outdoor recreation
	093 Other recreational items and equipment, gardens and pets
	094 Recreational and cultural services
	095 Newspapers, books and stationery
10	Education
11	Restaurants and Hotels
	111 Restaurants
	112 Hotels
12	Miscellaneous services
	121 Personal care
	123 Personal effects
	124 Social protection
	125 Insurance
	126 Financial services
	127 Other services
13	Payments to banks (credit)
14	Payments to public institutions (public)
cash	cash withdrawal

Notes: Category 13 and 14 applies only to payments via NICS. Cash refers to withdrawals made when making a debit card payment.

are closed either because they are deceased or moved out of Norway. Figure B.1 displays the total number and value of transactions in our micro data versus aggregate values reported in official statistics.

Figure B.1: BankAxept transactions and total amount.

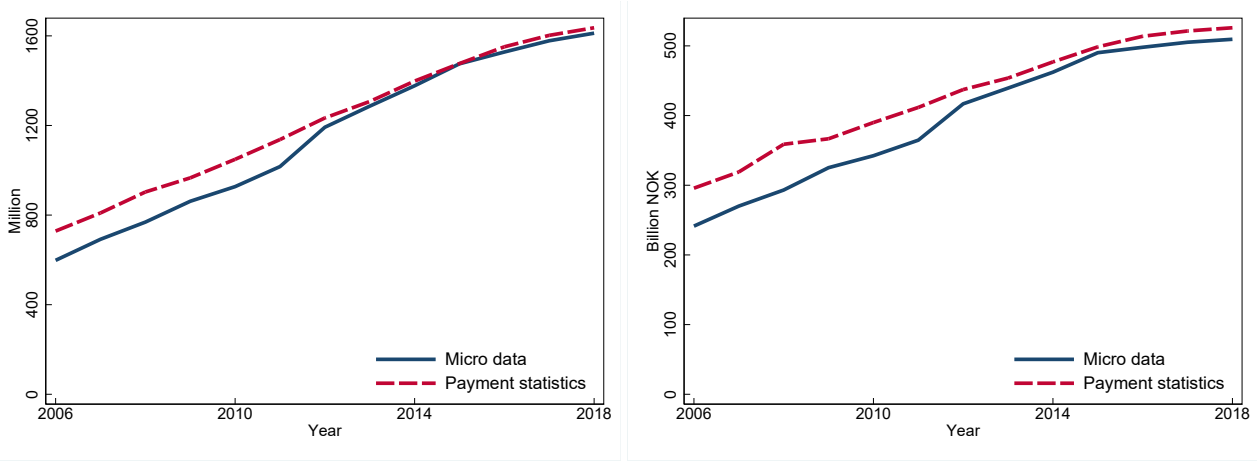

Overall BankAxept is the dominant means of card payment by Norwegians as shown in Figure B.2, panel (a). Over the time period 2006-2018 Bankaxept accounts for above 80 percent of domestic card payments made on physical terminals (in-store). Relative to all domestic card payments (physical terminals and online) BankAxept accounts for slightly less than 80 percent of total payments. The fraction is decreasing over time, starting at 82 percent in 2006 and ending at 71 percent in 2018. This reflects the fact that credit cards and online purchases (not covered by BankAxept) has become more important over the sample period. When including card payment abroad (physical and online) in total card payments, average BankAxept coverage drops to 69 percent and depicts a similar decline over the sample period. Cash as means of payment is not directly observed. However, using cash withdrawals from ATM and terminals as a proxy, Figure B.2, panel (b), shows that BankAxept card payments and withdrawals at terminals (which is what our micro data covers), the domestic (global) coverage rate is 67 (58) percent on average. Norway is among the least cash dependent countries in the world. Over the period 2006-2018 Cash in circulation has dropped from 6 to 2 percent of the M1 money supply. In contrast, cash accounted for more than 40 percent of the US M1 money supply in 2018.

Of the missing payments in Figure B.2, panel (b), credit card accounts for roughly 40 percent the payments. Despite not having the actual credit card purchases in our data, the granularity allows us to impute payments of credit card bills. In our wire transfer measure, payments to banks is a separate consumer category. Using an imputation algorithm (described below), we can further separate out payments of credit card bills. Our imputed

Figure B.2: BankAxept share of payments.

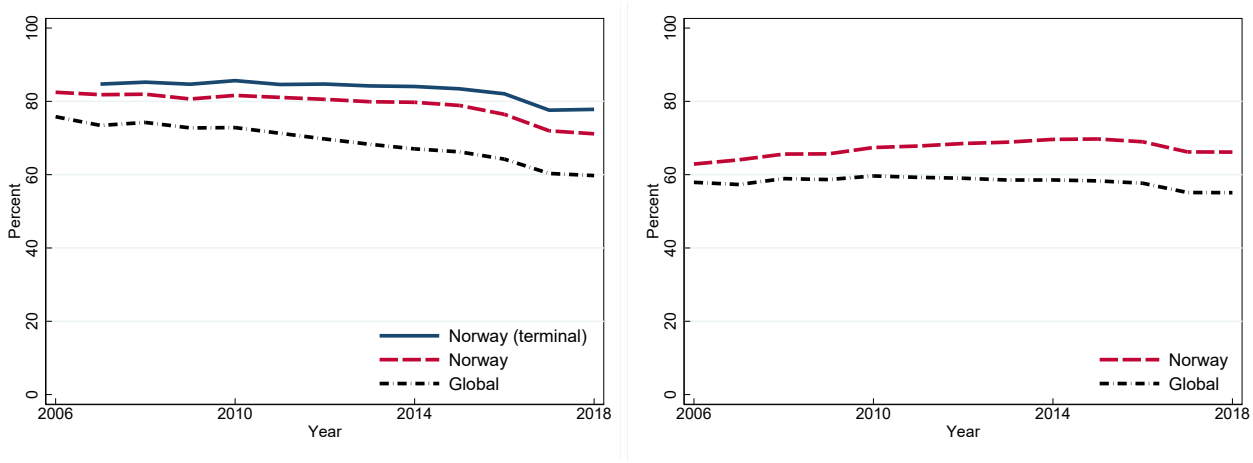

credit card payments covers on average 63 percent of total credit card payments from the official statistics. In figure B.3, when we add this imputed measure to total bankaxept payments, we cover around 70 percent of total card payments and cash withdrawals (ATM and terminal) by Norwegians over the sample period.

Figure B.3: Share of payments with imputed credit card.

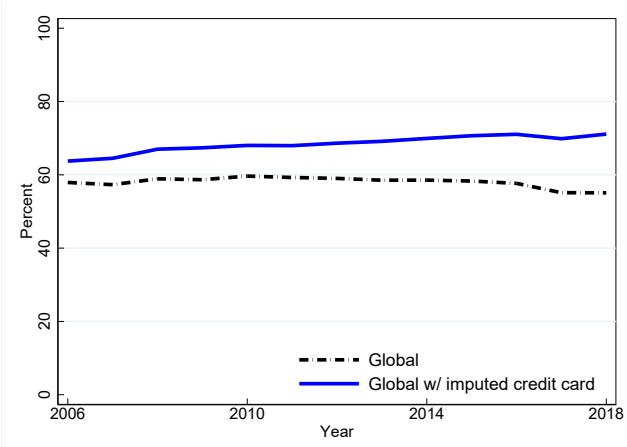

B.3 Bank Wire Transfers

The outgoing NICS payments contains wire transfers from individuals' bank accounts which are cleared via the Norwegian Interbank Clearing System (NICS). We remove large payments (above 500,000 NOK, 2015 CPI adjusted) as these are likely not capturing consumption, but rather saving/investment. If we do not remove such transactions our expenditure measure aggregates to more than 50 percent of National Accounts household consumption.

Large payments is located in certain categories, and these categories are typically associated with financial services (including housing transaction services).

Payments to banks (category 13) include debt service payments, in particular mortgage payments. We impute and remove these payments by splitting category 13 invoice payments in two groups. The grouping is based on the whether a payment is likely to be consumption (e.g. credit card) or investment (e.g. mortgage) related, and we clean the data by removing transactions classified as the latter. The imputation is based on the assumption that service of investment related debt is relatively stable over time, in contrast to other bank payments (e.g. credit card bills). The algorithm consist of 7 steps performed sequentially per individual. In each step, the payment is classified as imputed debt service if the condition is met and the transaction amount is above 2500 NOK (2018 CPI adjusted) .

1. Local stable payment: Compute the moving coefficient of variation (standard deviation over mean) over a rolling window of 8 payments within a zip-code, dropping the largest and smallest payment.for payments with multiple transactions within a week we first compute the average payment by dividing with number of transactions. 7.2 percent of invoice observations in category 13 involve multiple transactions. The payment is classified as imputed debt service if the coefficient of variation is below 0.12.

2. Typical payment within a year: Re-classify the remaining 25 percent of transactions as imputed debt service if the first step classifies more than 75 percent of transaction within the year as imputed debt service.

3. First and last year adjustment (I): Classify all payments in the individual's first (last) year of observation as imputed debt service if all payments in the following (preceding) year are classified as imputed debt service. This adjustment is done because the moving coefficient of variation is missing for the first and last 4 observations.

4. First and last year adjustment (II): Assume that the first (last) 4 observations are imputed debt service payments if the 5th (T-4) observation is a mortgage payment.

5. Temporary non-mortgage: Re-classify as imputed debt service if more than 75 percent of observations in a window of +/- 2 observations are classified as mortgage.

6. Typical payment in full sample: Re-classify as imputed debt service if more than 70 percent of all observations are imputed debt service.

7. Very large payment: Re-classify as imputed debt service is transaction amount is above 100,000 NOK (2018 CPI adjusted)

Some payments in category 13 we know are related to debt service in a known mortgage bank and student debt bank, and we separate out these two before applying the algorithm.[3] We nevertheless apply the algorithm on payments to these banks to check its performance. For the mortgage bank, our algorithm classifies 82.5 percent of the payments as debt service. For the student debt bank, it classifies 91.2 percent of the payments[4]. Moreover, when we compare to aggregate statistics on credit card payments, Figure B.4 shows that our cleaned debit category 13 tracks the official statistics on credit card payments well. We do not expect the algorithm to pick up all credit card bill payments, as internal bill payments (made from an account in the same bank that issued the credit card) does not show up in our transaction data.

Figure B.4: Imputed credit card spending.

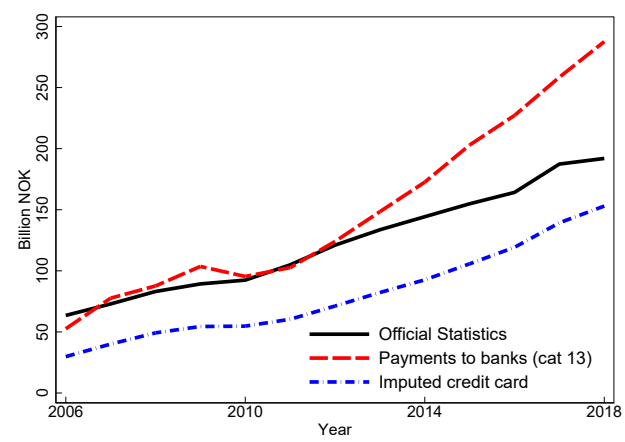

[3]We can precisely identify these payments because in certain zip codes there are only one bank. For two of these unique banks, individuals' bank service is solely related to either mortgage or student debt

[4]For the student debt bank, we do not apply the requirement that the transaction amount must be above 2500 NOK for it to be classified as imputed debt service.

C Additional Tables

Table C1: Summary Statistics for Married or Cohabiting Full-Time Workers

Sample:	**(a) Stayers**		**(b) movers**	
	Mean	St. Dev	Mean	St. Dev
Age	40.726	10.821	36.954	10.131
Male			0.649	
Lives in Oslo Region			0.424	
Household SIze	1		1	
Household Wage Income			13.015	0.512
Household Income After Tax			12.806	0.387
Household Gross Wealth			12.976	2.513
Household Debt			13.143	1.907
Household Expenditure			12.313	0.793
Food			10.322	0.959
Clothing			8.547	1.280
Furnishing			8.488	1.461
Motor Vehicles			7.710	2.041
Electronics			7.233	1.501
Retaurants			9.084	1.343
Number of Observations			38,799	

Notes: The sample includes married fulltime workers aged 25 to 60 in 2016. The 2024 NOK/USD $\approx$ 10

Table C2: Variance decompositions on singles and couples

	log earnings		log expenditure	
	AKM	KSS	AKM	KSS
variance	0.180	0.166	0.500	0.464
variance of firm effects	0.023 (0.129)	0.017 (0.102)	0.018 (0.036)	0.009 (0.019)
variance of person effects	0.144 (0.798)	0.125 (0.756)	0.486 (0.973)	0.445 (0.959)
covariance of firm and person effects	0.007 (0.074)	0.012 (0.142)	-0.002 (-0.009)	0.005 (0.022)
correlation of worker and person effects	0.115	0.255	-0.025	0.083
R squared	0.807	0.743	0.750	0.697

Notes: Variance decomposition of log earnings and log expenditures for singles and couples on the leave-one-out connected set. AKM refers to the standard AKM decomposition as in Abowd et al. (1999), while KSS is the bias-corrected estimates following Kline et al. (2020). Numbers in parenthesis refers to fraction of explained variance

Table C3: Variance decompositions on household level wage and consumption

	log earnings		log expenditure	
	AKM	KSS	AKM	KSS
variance	0.179	0.164	0.335	0.307
variance of firm effects	0.015 (0.082)	0.009 (0.057)	0.014 (0.042)	0.006 (0.021)
variance of person effects	0.162 (0.907)	0.144 (0.878)	0.330 (0.985)	0.297 (0.967)
covariance of firm and person effects	0.001 (0.010)	0.005 (0.065)	-0.004 (-0.027)	0.002 (0.012)
correlation of worker and person effects	0.019	0.145	-0.066	0.043
R squared	0.776	0.713	0.713	0.652

Table C4: Variance decompositions with household level wage and consumption

Notes: Variance decomposition of household level (equivalized) log earnings and log expenditures for singles and couples on the leave-one-out connected set. AKM refers to the standard AKM decomposition as in Abowd et al. (1999), while KSS is the bias-corrected estimates following Kline et al. (2020). Numbers in parenthesis refers to fraction of explained variance

Table C5: Variance decompositions on household level wage and consumption

	log earnings		log expenditure	
	AKM	KSS	AKM	KSS
variance	0.179	0.164	0.335	0.307
variance of firm effects	0.015 (0.082)	0.009 (0.057)	0.014 (0.042)	0.006 (0.021)
variance of person effects	0.162 (0.907)	0.144 (0.878)	0.330 (0.985)	0.297 (0.967)
covariance of firm and person effects	0.001 (0.010)	0.005 (0.065)	-0.004 (-0.027)	0.002 (0.012)
correlation of worker and person effects	0.019	0.145	-0.066	0.043
R squared	0.776	0.713	0.713	0.652

Table C6: Variance decompositions with household level wage and consumption

Notes: Variance decomposition of household level (equivalized) log earnings and log expenditures for singles and couples on the leave-one-out connected set. AKM refers to the standard AKM decomposition as in Abowd et al. (1999), while KSS is the bias-corrected estimates following Kline et al. (2020). Numbers in parenthesis refers to fraction of explained variance

D Additional Figures

Figure D1: Variance decomposition plots: household wage and consumption categories.

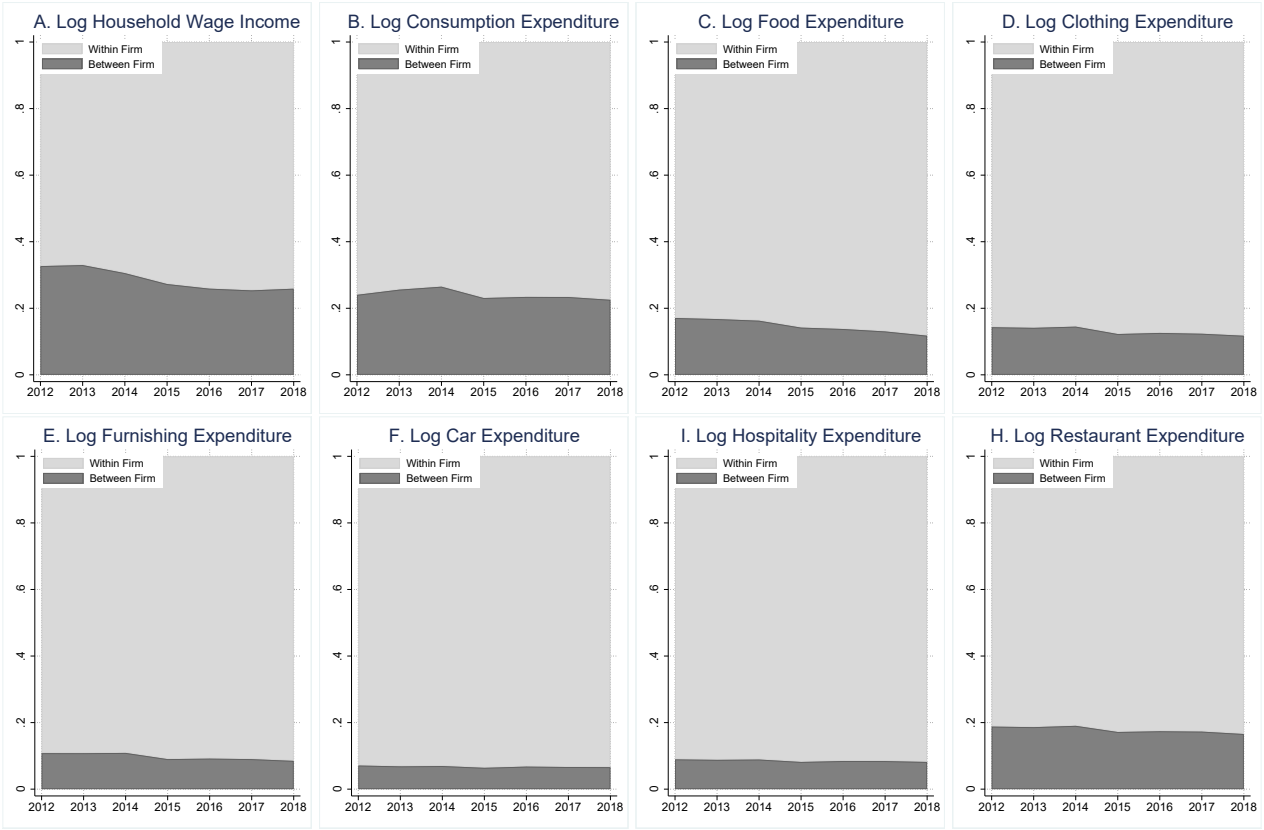

Notes: These figures report normalized variance decomposition of wages and components of consumption for married and cohabiting.

Figure D2: Effects of change in log co-worker wage

Notes: These figures report point estimates and 95% confidence intervals of the associated change in income and consumption from a change in log average co-worker wage. Standard errors are clustered at the region level.

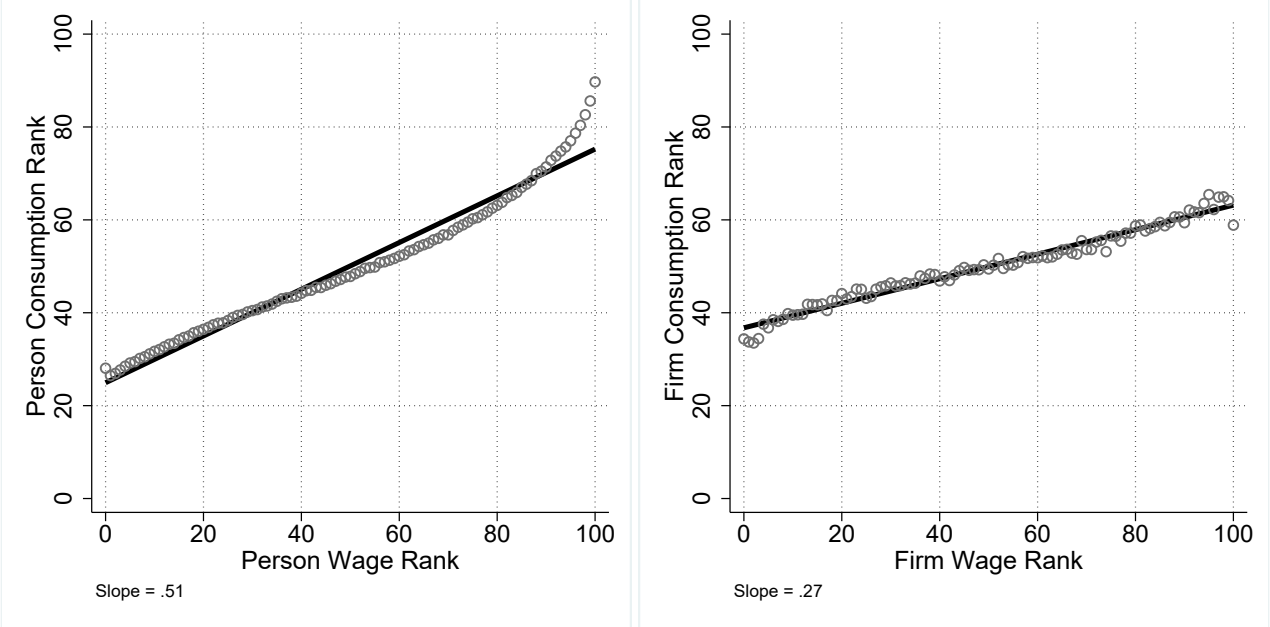

Figure D3: All individuals Rank-rank correlations between person log consumption and person log earnings effects (left), and firm log consumption and firm log earnings effects (right)