Firm Pay and Consumption Inequality*

Sigurd Galaasen† Andreas R. Kostøl‡ Michael Simmons§

September 12, 2025

Abstract

Despite extensive evidence documenting that firms affect wage inequality, less is known about the role of firms in shaping workers' welfare. We address this gap by combining expenditure data with administrative employer-employee records from Norway and a movers-based instrumental variables strategy. We find that when workers move to a new firm, a $1 pay differential leads to a 45-cent differential in spending. There is substantial heterogeneity in the consumption response, with no spending cuts for downward pay adjustments, with an 80% spending response to upward adjustments. Using a variance decomposition, we find that firms explain roughly half the variance of consumption inequality relative to wages. Finally, we demonstrate that a search-and-savings model quantitatively reproduces these results and provides support for a simple revealed preference approach that connects the empirical firm consumption effects to welfare. Taken together, our results indicate that worker spending decisions compress the pass-through from firm pay inequality to welfare inequality by roughly a half.

*This project received financial support from the Norwegian Research Council through grant 345236. 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 want to thank Axel Gottfries, Daniel Haanwinckel, Rasmuz Lentz, Luigi Pistaferri, and seminar participants at ASU, the 2nd Oslo Labor Economics Workshop, 2025 SED, 2025 EEA, UCLA and Strathclyde for helpful comments and suggestions.

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

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

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

1 Introduction

An extensive body of research has shown that firms play an important role in determining pay disparities (e.g., Abowd et al., 1999, Card et al., 2013, Kline, 2024). These findings have fueled growing interest in identifying the underlying drivers of pay differentials.[1] Despite substantial progress in understanding, little is understood about the extent to which these disparities translate into differences in well-being. Assessing the welfare implications of firm-level pay differentials is challenging due to several conceptual and measurement issues. For example, they depend on workers' expectations over future pay (Friedman, 1957), job search behavior (Lise, 2013), and liquidity (Kaplan and Violante, 2014), factors that are typically not observed in existing datasets.

This paper circumvents the measurement and conceptual issues by combining a revealed preference approach with newly collected data on consumption expenditure and employment from Norway. We adopt a movers-instrumental variables (IV) approach to estimate the marginal propensity to consume (MPC) out of pay premiums. To do so, we estimate the firm pay premium using workers who switch between employers, and use this estimate as an instrument to identify the marginal propensity to consume of income. The advantage of this approach is that it addresses endogeneity, and by estimating the MPC, we can remain agnostic about how and why pay differentials affect average consumption choices.

We find that the MPC out-of-pay premiums are sizable. On average, a $1 pay differential leads to a 45-cent differential in spending. These effects are persistent, displaying only a modest decline three years after the move. However, the average pay differential masks important heterogeneity. While there are no spending cuts for downward pay adjustments, there is an 80% spending response to upward adjustments. Using details of the spending data, we further decompose the spending effects by category. We find that pay differentials mainly cause non-essential spending differentials, such as restaurants and hotels, but find no statistically significant effect on essential spending, such as groceries and communication.

Next, we assess the role of firms for inequality in consumption and welfare by applying a consumption equivalent variation (CEV) approach. We rely on a log-linear approximation of both pay and spending functions, allowing for a reduced-form estimate of the importance of the role of firms for inequality in pay and welfare. Performing a variance decomposition of

[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 titled "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).

wage and consumption expenditure, we find that the variance of firm effects for consumption is approximately half the variance of firm effects for wages – and is consistent with important consumption smoothing behavior. We show that a model of job search, two-sided heterogeneity, and savings can quantitatively rationalize our estimated MPC and variance decomposition results. Moreover, we use the model to relax the log-linearity assumption in consumption when considering welfare inequality. We find that the model-based variance of welfare due to firms is strikingly close to the variance of firm consumption effects, lending support to the reduced form analysis.

In Section 2, we describe the consumption expenditure and earnings data and document their cross-sectional properties. The key to our empirical contribution is the ability to link newly collected data on expenditure with two key data sets. First, we link it to employment histories, which allows us to characterize pay differentials between employers and measure how consumption expenditure evolves over a career. Second, we link expenditure to information on households, which allows us to account for a key demographic variable when studying spending; whether a person is making purchasing decisions based on the needs of one or many persons in a household. We show that the level of consumption expenditure is lower than disposable income, consistent with savings behavior, and that the standard deviation of consumption is lower than that of gross wages, which is consistent with consumption smoothing behavior.

In Section 3, we discuss the empirical challenges associated with assessing how firms affect consumption and our approach to circumvent these challenges. Our approach addresses reverse causality and endogeneity concerns using a movers-IV design, where we leave out half of the sample when estimating firm pay effects, and use firm pay effects as an instrument for disposable income. Our measure of income accounts for all taxable- and non-taxable income after tax, and we apply our empirical strategy to a sample of employees who switch employers between 2012 and 2015. This allows us to track their spending three years before and after the move, visual and formal tests of pre-trends and the dynamics of treatment effects. We find no evidence of differential pre-trends in income or consumption of workers who move to higher- or lower-paying firms.

We present four main findings on MPCs for job movers in Section 4. First, we find that the average consumption response to the pay premium differentials is 0.45 among singles. This effect manifests itself fully one year after the move and only decreases slightly three years after. Second, we find that the MPC is quantitatively similar for married and cohabiting individuals, where the spending effect is entirely driven by the individual who moves. Third, we document important effect heterogeneity along the job ladder. There is no evidence of a drop in spending for negative changes, while the MPC estimate for positive

changes is not statistically different from one. Moreover, the level of the wage ladder matters: the effects are higher when moving from the lowest-paying firms and when moving to the highest-paying firms. Fourth, we find significant differences in the elasticity of spending across consumption categories. While there is no statistically significant effect across essential spending categories, such as groceries, transportation, and communication, we find the largest relative change in spending on hotels, restaurants, and financial services.

In Section 5, we assess the extent to which firms matter for *welfare inequality*. We build on the large literature showing that the (bias-corrected) variance of firm fixed effects explains a substantial share of overall wage inequality. To extend this framework to consumption, we adopt a log-linear specification of spending. While restrictive, this assumption provides a natural benchmark given the widespread use of two-way fixed effects in studies of pay inequality. Under this restriction, and assuming utility is increasing in consumption, we show that the variance of firm consumption effects is equivalent to the variance of welfare as measured by CEV gains or losses from moving between firms. Empirically, we find that the variance of firm effects in consumption is about half the size of that for pay. This ratio is remarkably stable across data subsamples and robust to alternative treatments of outliers, cash use, and government transfers. Taken together, the evidence suggests that pay inequality is informative about welfare inequality, but firms' role is attenuated: their contribution to welfare inequality is about half their contribution to wage inequality.

In Section 6, we explore whether an off-the-shelf search model with savings and two-sided heterogeneity can account for these new facts, and we use the model to relax the log-additive assumption when studying firms' role in welfare. We show that the model quantitatively reproduces both the estimated MPC and the variance decomposition under a reasonable degree of risk aversion.[2] The key mechanism is search frictions: workers do not know when job-switching opportunities will arise, which prevents perfect smoothing of consumption across firms. As a result, they experience discrete jumps in consumption and welfare at job switches. Although consumption in the model is not log-additive, since responses depend on workers' asset holdings, we show that the variance of firm consumption effects remains close to the variance of firms' contribution to welfare inequality as measured by average CEVs. Overall, firm pay heterogeneity is a significant source of wage inequality; however, individuals' consumption choices mitigate the pass-through to welfare inequality by approximately 50%.

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

[2]The coefficient of relative risk aversion is tightly linked to the MPC and the variance of firm effects on consumption; we estimate a value just above 2.

been refined with modern econometric techniques, is that the variance of firm pay premiums are important (see, e.g., Kline, 2024 for an extensive review). Moreover, firm pay effects have been shown to be strongly related to different observable measures of firm quality, including value added, and firm size (e.g., Bloom et al., 2018 and Card et al., 2016), and other desirable aspects of the job (e.g., Sorkin, 2018 and Bagger and Lentz, 2019). We contribute to this literature in important ways. We document that firm pay premiums are not only reflected in wages, but also translate into systematic differences in household consumption, and we provide quasi-experimental evidence of this link using job movers. Moreover, we show that the variance of firm effects in consumption is smaller than in wages, highlighting that while firms matter for welfare inequality, their role is attenuated compared to wage inequality.

Our paper also contributes to a long literature considering consumption inequality and the marginal propensity to consume. Deaton (1992) and Attanasio and Pistaferri (2016) discuss the conceptual differences between inequality in income and welfare, and Cutler and Katz (1992) documents a widening of the distribution of both income and consumption in the US during the eighties. Blundell and Preston (1998) propose using the covariance between income and consumption to assess changes in household welfare over time, and Blundell et al. (2008) document that consumption responses to income shocks depend on whether they are transitory or persistent. More recent studies have used de-identified bank account data to study consumption responses to income and unemployment shocks.[3]. In two related studies, Ganong et al. (2020) show that low-asset individuals also cut down on spending to transitory income shocks, and Fagereng et al. (2021) estimate an MPC of 0.45 using imputed consumption and large lottery winnings. We contribute by combining expenditure data with administrative employment records to estimate how firm pay premiums pass through to household spending. In doing so, we provide the first evidence demonstrating how pay disparities between firms translate into consumption and welfare inequality.

Finally, we complement empirical and structural work motivated by models of job search (e.g., Lentz et al., 2023), and job ladders (Burdett and Mortensen, 1998; Postel-Vinay and Robin, 2002), to assess the underlying sources of wage inequality (e.g., Morchio and Moser, 2024, Lamadon et al., 2024 and Borovicková and Shimer, 2024). In related work, Conlon et al. (2018) document large differences between expected pay and realized wage changes, often leading to significant upward adjustments of income expectations. Our evidence is consistent with the presence of search frictions as workers climb the job ladder. Similar to our quantitative exercise, these moves generate welfare gains as workers climb the ladder. The model-generated MPCs reflect that the expected future path of income changes, since

[3]See, for example, Baker (2018) and Ganong and Noel (2019) for evidence from the US, Gerard and Naritomi (2019) for Brazil, and Andersen et al. (2023) for evidence from Denmark.

workers randomly, and infrequently, receive opportunities to climb the rungs of the ladder.

Our quantitative framework builds on Lise (2013), where workers in lower-paying jobs consume proportionally more of their income than workers in higher-paying jobs. Low-wage workers want to increase spending as they expect to climb the ladder, while workers in higher-paying jobs save to ensure against the higher income risk of falling off the ladder. This pattern flattens the gradient in consumption along the job ladder than wages. Our evidence on highly asymmetric MPC out of positive and negative firm pay effects is consistent with these mechanisms. Moreover, the large MPC resulting from positive wage changes suggests that important search frictions blur workers' expectations of future income. Using an extended version of Lise (2013), we show that the model is able to quantitatively reproduce the estimated MPC along the job ladder, and the contribution of firms for consumption inequality relative to wages, which we measure empirically in this paper.

The paper proceeds as follows. In Section 2, we describe the data we use for our analyses. In Section 3 we describe our method for estimating MPCs, and we implement the method in Section 4. In Section 5 we describe the revealed preference approach and estimate the variance contributions for wages and consumption. In Section 6, we consider the model, and Section 7 concludes.

2 Data

What makes our empirical assessment possible is the unique combination of data employment histories and electronic payments for the universe of Norwegian residents. This section describes the data and illustrate their cross-sectional properties.

2.1 Linked Expenditure and Employment Data

The main data source is provided by the Norwegian retail clearing institution, Nets Branch Norway (henceforth referred to as Nets). The data set contains information on all debit card payments and all online bank wire payments cleared via the Norwegian Interbank Clearing System (NICS).[4] We received individual-level data aggregated at the weekly frequency, with information about total total expenditure, number of transactions, and cashback at a

[4]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". In addition, NETS covers all incoming transfers cleared via NICS. We do not make use of this data in this paper

zip code and 25 broad consumption categories. Moreover, the electronic payments are classified by purpose derived from United Nations' 1999 COICOP system.[5] Using the industry code information, transactions to/from banks and the public sector are placed into two separate categories. We further refine the bank payment category by removing payments that are likely related to servicing of mortgage (or other investment-related) debt, and demonstrate that the remaining bank payments align well with aggregate statistics on credit card payments. Appendix B provides additional details on data cleaning procedures.

What makes our analysis unique is that we can link the expenditure data to administrative records using unique anonymized personal identifiers. The first key linkage is to the universe of employment relationships from 2006 to 2018.[6] The data provides complete spell data of all employers, with start- and end dates, a measure of contracted hours, allowing for a definition of full-time employees, in addition to the workplace location. The second key linkage is to population registers, allowing us to define consumption expenditure at the household level and not only at the individual level. In addition to married couples, we can also define unmarried couples who reside at the same address, as well as the number of children living in the same household. The data set also contains information on age and gender. Finally, we measure annual labor income from tax data, taxes, residential municipality, and other transfer payments.

2.2 Measurement

This section assesses the quality of the expenditure data by providing a set of facts about their cross-sectional properties. To this end, we restrict the sample to workers used in the subsequent analysis, namely full-time employees (working at least 30 hours per week) aged 25-60. Next, we deflate wages and expenditures to 2018 values using the consumer price index, and then winsorize each variable at the 1st and 99th percentiles at the household level, which is used throughout the analysis. To illustrate what these variables measure in

[5]Card payments are mapped to COICOP based on the Merchant Category Codes of the card terminal (MCC). This 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, resulting in a total of 24 categories. Table C.1 in Appendix B.1 lists all our categories and provides details on how transactions are linked to a particular COICOP code.

[6]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 cross-section, we distinguish between single-person and multi-person households.

Table 1 reports averages and standard deviations of wages and consumption separately for these two groups of workers. As expected from a savings motive, households consume less than their wages. This holds for the level of consumption and wages for single-person households in Panel A and the individual- and household-level in Panel B. Next, we observe that the standard deviation of consumption is lower than that of individual wages, consistent with consumption smoothing behavior. Appendix Figure A1 further shows that this relative dispersion pattern is relatively stable over our observation period. An exception is for observed individual spending in multi-person households, as evident from the last two rows in Table 1. This reflects the fact that spending and consumption are often shared within households, highlighting the need to measure consumption at the household level.

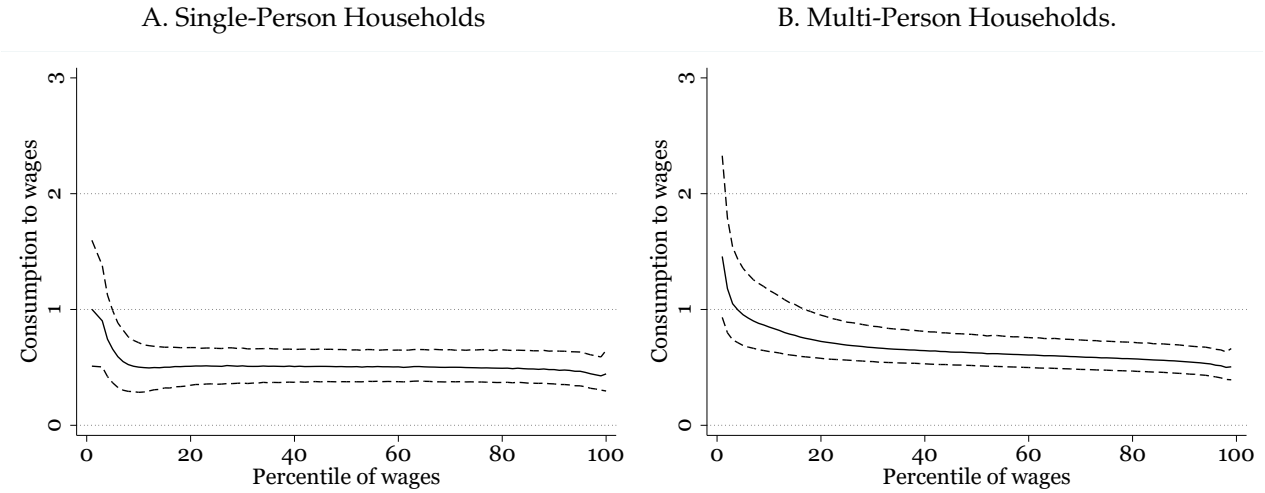
Table 1: Summary Statistics: Consumption and Wage Income (NOK1,000)

	A. Single-Person Household		**B. Multi-Person Household**	
	Mean	SD	Mean	SD
Consumption (household)	318.27	249.10	789.75	453.48
Wages (household)	574.50	259.35	1139.37	503.19
Consumption (own)			396.69	295.21
Wages (own)			628.85	285.00
Observations	2,570,293		10,280,952	

Notes: Variables denominated in thousands of real 2018 NOK. Single-person refers to a household with only one member, while multi-person refers to those with more than one member. Unit of observation is the worker. First two lines report the workers' total household spending and wage income, while the last two report the workers' own spending and wage income. We remove worker-year observations when the workers' total household consumption is missing. We then winsorize variables at the 1st and 99th percentile of their full distribution (pooling all workers).

Next, we examine the ratio of expenditure over wages along the wage distribution. To do so, we order households by percentiles of their household labor income, and report the first, second, and third quartiles within each percentile group. Figure 1 illustrates that the ratio of consumption to wage income is generally stable and below one across the wage distribution. At the same time, the 75-25th percentile range is wide, suggesting important heterogeneity in consumption along the wage distribution. We also see a noticeable difference at the lowest 10-20 percent of the distribution, where the consumption-to-wage ratio is higher, and spending is likely to be financed through other sources, such as government transfers. This pattern is more pronounced in multi-person households, where only one household member needs to be fully employed.

Figure 1: IQ Range Consumption-Wage Income Ratio



Notes: This figure plots the median and interquartile range of consumption to wage income ratios, within percentiles of wage income, for single-person (the left-hand-side panels) and multi-person (the right-hand-side panels) households. Both wages and consumption are aggregated to the household level, and the unit of observation is the worker.

Our final exercise is to compare the measurement of consumption expenditure using NETs with aggregate statistics from the National Accounts, which is measured using a combination of different data sources. In Appendix Figure A2 we show that the Nets expenditure measure tracks the quarterly level and growth rate of consumption from the National Accounts reasonably well. The correlation between quarterly growth rates of national account household consumption and our expenditure measure is 0.85.

3 Empirical Approach

This section describes our empirical approach for estimating the effects of firms on consumption. To understand how firms' pay policies affect household consumption, we must address two main empirical challenges. The first concerns simultaneity, where a firm's pay policies and households' consumption needs may be determined jointly by common shocks. The second is related to reverse causality. For instance, households may face higher consumption demand due to a growing family size, which induces greater effort to remain employed or to seek higher wages.

3.1 A Movers' Design

We address these challenges in two steps. In the first step, we apply a two-way fixed effects model for gross wages, following Abowd et al. (1999), where the basic idea is that workers typically experience wage gains as they move between firms with different pay policies. We begin by specifying a standard Abowd, Kramarz, and Margolis (AKM) two-way fixed effects model, in which the log annual labor income of individual i in year t, is denoted $y_{i,t}$,

$$y_{i,t} = \alpha_i + \psi_{J(i,t)} + X'_{i,t}\beta + \varepsilon_{i,t}, \tag{1}$$

where α_i is a worker fixed effect capturing time-invariant individual heterogeneity, $\psi_{j(i,t)}$ is the firm component of pay, or the fixed effect associated with the employer j. $X_{i,t}$ is a vector of time-varying observable characteristics, β is a vector of coefficients, and $\varepsilon_{i,t}$ is an idiosyncratic error term.

The two-way fixed effects approach enables us to disentangle variation in pay that is unrelated to individual earnings capacity under an exogenous mobility assumption. Card et al. (2013) shows that a firm fixed effect can be interpreted as a firm premium under an exogenous movers assumption. This assumption requires that individuals do not systematically sort into firms based on unobserved pay characteristics, such as changes in health or due to match-specific components in pay, and is discussed extensively in Kline (2024).

To estimate pay premiums, we use the total sum of taxable labor earnings, as determined by the tax data. We estimate the log wage model using all workers aged 20-66 and ten years of data, from 2009 to 2018. In addition to person and firm dummies, we control for dummies for each age and calendar year. We estimate ψ using the full sample and a split-sample design, in which we estimate it on only half of the sample. We report summary statistics of key demographic variables, annual wages and the estimated firm-fixed effects for the 2013 sample of movers and stayers in Table 2. We see that movers are younger and more likely to be male than stayers. Importantly, we see that the standard deviation of the overall estimate of ψ is substantially smaller than the split sample estimate, which is true for both samples.

3.2 An Instrumental Variables Approach

In the second step, we use the change in estimated firm-fixed effects as an instrumental variable (IV) to estimate the marginal propensity to consume (MPC) out of income. To implement this design, we must address the fact that the firm component in pay is identified from the sample of movers, which creates a mechanical correlation between individual (i.e.,

Table 2: Summary Statistics for Population of Movers and Stayers in 2013

Sample:	**A. Movers**		**B. Stayers**	
	Mean	St. Dev	Mean	St. Dev
Age	38.840	9.304	42.391	9.777
Male	0.644		0.594	
Single	0.694		0.719	
Number of Children	0.894	1.046	0.863	1.041
Wage Income (NOK1,000)	557.776	279.045	566.325	272.800
Overall ψ	0.062	0.218	0.049	0.216
Split-sample $\tilde{\psi}$	0.060	0.247	0.048	0.245
Number of Observations	116,152		1,115,813	

Notes: The table reports summary statistics for the sample of full-time employed workers in 2013.

the mover) pay and the firm component (i.e., the "firm premium"). We address this using the split-sample estimate of $\psi_{j(i,t)}$, assigning to each worker the firm premium estimated from the opposite subsample, $\tilde{\psi}_{j(i,t)}$. Specifically, the first stage of this movers-IV model captures the change in net after-tax income induced by the difference in estimated pay policies between firm k and j,

$$\underbrace{y_{i,t} - \bar{y}_{i,s-1}}_{\textbf{Change in net income}} = \underbrace{\gamma_t^y [\tilde{\psi}_{j(i),s} - \tilde{\psi}_{k(i),s-1}]}_{\textbf{Change in Firm Fixed-Effect}} + \eta_{i,t} \tag{2}$$

where a worker moves to the new firn k in event year s. Following Yagan (2019), we difference out permanent characteristics of households using the pre-event average net income $\bar{y}_{i,s-1}$, so that outcomes are determined by the change in firm fixed effects (i.e., the firm pay premium) and a residual $\eta_{i,t}$. The estimate of γ_t^y gives us the first stage coefficient, i.e., how much net income changes in period t relative to the year before the move, $s-1$.

Similarly, a job switch will induce a change in household spending from year s to t,

$$\underbrace{c_{i,t} - \bar{c}_{i,s-1}}_{\textbf{Change in Spending}} = \underbrace{\gamma_t^c [\tilde{\psi}_{j(i),s} - \tilde{\psi}_{k(i),s-1}]}_{\textbf{Change in Firm Fixed-Effect}} + \varepsilon_{i,t} \tag{3}$$

where γ_t^c captures the reduced form effect of the difference in firm fixed-effects on household spending. The estimates of γ_t^y and γ_t^c are both attenuated by noise in the estimates of the $\tilde{\psi}_{j(i)} - \tilde{\psi}_{k(i)}$, where the attenuation bias depends on the variance of the true difference relative

to the variance of the noise.[7] However, applying indirect least squares, the noise from the denominator and numerator cancels out, yielding an unbiased estimate of **MPC** = $\frac{\gamma^c}{\gamma^y}$.

3.3 Identifying Assumptions

To interpret this ratio as the causal effect of the change in average firm pay premia on household consumption expenditure, we must invoke the following assumptions. First, we impose a stronger exogenous mobility assumption. By strong exogeneity, we mean that job mobility is not only independent of unobserved shocks to earnings, but also unobserved changes in preferences and constraints. This assumption is violated by match-specific components in pay, as discussed by Card et al. (2013), and from unobserved increases in spending needs resulting from changes in household size, or adjustments in constraints due to relocation to a new labor market or the dissolution of marriage. Second, we must invoke the standard exclusion restriction, meaning that the firm-fixed-effect only affects household spending through the net income channel. This is violated, for example, if higher pay premiums are associated with other aspects of a job, such as higher unemployment risk, where the pay premium also affects spending through other components. We will address this potential violation in section 4.

It is important to note that we do not require that changes in pay premia are unanticipated. If a job switch and the associated wage gain are anticipated, workers would have priced the expected pay gain in their consumption changes before the job switch. Thus, we should observe muted changes in consumption at time $t \geq s$, whereas moving to a higher-paying firm might lead to an increase in spending the year before the move.

3.4 Sample Construction and Key Variables

We implement our approach by estimating the MPC using two-stage least squares (2SLS) for a movers sample. We construct the sample by first restricting the data to all full-time workers aged 25-60 who change employers between 2012 and 2015. These restrictions ensure that we have a balanced sample of workers and households for which we can measure consumption and income for at least three years prior to and three years after a move, i.e., $t - s \in [-3, 3]$. Next, we retain workers with at least 36 months of tenure at their original firm (j), but impose no restrictions on tenure in the destination firm (k), except that the move is not associated with a change in the indicator for full-time work. We also exclude

[7] Omitting individual subscripts and assuming time-invariant firm effects, $\tilde{\psi}_j = \psi_j + u_j$ and $\tilde{\psi}_k = \psi_k + u_k$, the measured firm-premium change is $\tilde{X} = X + (u_j - u_k)$. Under classical measurement error assumptions, regressing $c_t - c_{s-1}$ on $\tilde{X}$ recovers γ^c attenuated by the signal to noise ratio $\text{Var}(X)/[\text{Var}(X) + \text{Var}(u_j) + \text{Var}(u_k)]$.

a small number of individuals in which the spouse also moves to a new firm in the year s. These additional restrictions alleviate the concern that other unobserved factors drive both the change in firm pay premiums and changes in household consumption.

We measure net after-tax income using tax data. This includes labor and capital income, as well as transfers from the government and taxes paid. The measure of consumption is the sum of all spending categories for all household members. Both income and consumption are computed as the long difference between time t and $s-1$, and normalized using the level of disposable income in period $s-1$. Finally, we winsorize the normalized difference in wage income and consumption at the top and bottom 1% of the distribution.

Throughout our analysis, we estimate the average MPC out of pay premiums separately for single individuals and those in cohabiting relationships. This allows us to zoom in on decisions made without regard to a partner for one sample, and then to decompose the total spending effect into the amount spent by each spouse and the pay premium spent by the other spouse. We provide summary statistics for our sample of movers in Table 3. We see that singles have fewer children than cohabiting (or married individuals). As anticipated, singles are younger and more likely to work and live in Oslo, Norway's capital. The level of pay and disposable income is slightly more than half that of the cohabiting, and they have a consumption-to-net-income ratio of 0.75, which implies a 10 percentage point higher savings rate than those cohabiting, whose spending ratio is 0.85. As singles have fewer children, we see that the spending share on food and clothing is lower than that of couples, whereas spending shares on the selected set of other categories are broadly similar.

Table 3: Key Descriptives for Single and Married/Cohabiting Movers

Sample:	**A. Single**		**B. Cohabiting & Married**	
	Mean	St. Dev	Mean	St. Dev
Age	41.110	9.859	42.897	9.226
Male	0.654		0.731	
Lives in Oslo Region	0.374		0.325	
Number of kids	0.226	0.540	1.163	1.072
Household Gross Wealth	2,831.029	2,803.709	5,275.457	4,011.693
Household Debt	1,407.801	1,318.743	2,525.667	1,821.583
Household Wage Income	662.245	296.124	1,221.495	503.204
Household Income After Tax	514.215	217.026	983.304	340.254
Household Expenditure	401.098	292.501	838.032	446.924
Food	49.175	34.324	113.705	52.748
Clothing	12.233		32.500	21.971
Furnishing	15.094	22.208	32.738	32.909
Motor Vehicles	12.802	48.076	28.614	74.236
Retaurants	14.394	15.366	22.717	18.362
Hotels				
Number of Observations	200,730		552,677	

Notes: The table reports summary statistics for the estimation sample of movers in 2012-2015 (excluding movers with a change in full-time employment and multiple movers in a household the same year in Panel B), aged 25-60, with at least 36 months of tenure.

4 The Marginal Propensity to Consume Pay Premiums

This section presents evidence on marginal propensity to consume (MPC) out of income using a quasi-experimental approach. Suppose contemporaneous pay differences reflect luck in the timing of job offers, while workers are otherwise fully informed about the wage distribution (i.e., the wage ladder in Burdett and Mortensen, 1998). Since workers derive utility from expected lifetime pay, their consumption response to a pay differential depends on two factors. First, if they expect to climb the ladder and receive higher wages in the future, the current pay differential may have little impact on their consumption. Second, if workers face liquidity constraints and cannot smooth their income, they may consume a large share of a pay increase. More generally, any mechanism that prevents workers from knowing or accessing their lifetime wealth implies that firm pay policies can affect household consumption. In this case, the MPC provides a measure of the firm's relevance for welfare, even if we cannot define the firm pay premium as a permanent or a transitory income shock. We expand on this revealed-preference approach in the next section.

4.1 Observational Evidence

We begin by presenting evidence using the sample of stayers to examine consumption responses do differential pay across individuals. As we illustrated in Table 2, stayers tend to work for firms with slightly lower average pay premiums, on average, and are older than workers who move between employers (i.e., the movers). Although the firm premium is estimated based on the movers' data, the regression analysis does not use a clear source of variation in pay, except that it is identified from movers, while controlling for the quality of workers (through the connected job mobility set). To examine the firm premium and the marginal propensity to consume, we regress net income and household spending on the estimated firm-pay premium $\tilde{\psi}_{j(i)}$, and calculate clustered robust standard errors at the firm level. With these caveats in mind, our thought experiment involves assigning a worker a firm premium that would double the level of pay (i.e., moving $\tilde{\psi}$ from 0 to 1).

We report our estimates in Table 4. In the first column, we regress income on the firm pay premium but has otherwise no controls. The result indicates that doubling the average wage would increase the average net income by NOK 227,000 (approximately USD 27,000). The estimated change in consumption is NOK 146,000, which implies an MPC of 0.64. Next, we add dummies for age controls and only see a minimal change in the estimated coefficient on income, whereas the coefficient on spending increases slightly. In columns (3) and (4), we add more control variables and see that the coefficient on consumption rises to 164,000,

implying an MPC that exceeds 0.7.

This first evidence shows a strong positive relationship between higher-paying employers and higher spending, which is consistent with the idea that heterogeneity in firm pay policies translate into differences in welfare. However, this approach relies on across-individual variation, meaning that the estimated marginal propensity to consume reflects not only the effect of higher pay, but also systematic differences between workers who sort into different firms. We next move to the quasi-experimental approach that relies on within-individual variation using a sample of job movers.

Table 4: Income, Consumption, and Firm Pay Premia

	(1)	(2)	(3)	(4)
A. Dep. variable: Net Income (NOK1,000)	226.932 (22.407)	230.945 (21.839)	229.251 (22.084)	231.727 (21.950)
B. Dep. variable: Consumption (NOK1,000)	145.583 (13.498)	150.725 (12.842)	160.994 (12.747)	164.278 (12.560)
Age control		Yes	Yes	Yes
Female control			Yes	Yes
N Children control				Yes
Observations	154987	154987	154987	154986

Notes: The table reports coefficient estimates from regressions of (A) total net income and (B) total expenditure on the split-sample firm-fixed effect. The sample is full-time employed singles with and without children aged 20–66 in 2013. Robust standard errors (in parentheses) are clustered at the firm level.

4.2 Quasi-Experimental Evidence

Turning to the mover-IV approach, we remind the reader that we impose (1) a strong exogenous mobility assumption and (2) exclusion restriction. These assumptions allow us to identify the impact of firm pay policies on consumption within households. We begin by examining the impact of the split-sample pay premium income on normalized net income changes among singles, year by year, from three years prior to the job change (from firm j to k) up to three years after. We plot the first-stage coefficient estimates from equation 2 as the dashed line in Panel A of Figure 2. The horizontal axis refers to the year relative to the move (e.g., normalized event time t – s), and the vertical axis denotes the change in normalized income and consumption expenditure. Reassuringly, we find no statistically significant changes before the move, visualized by the 95 percent confidence intervals, which supports

the exogenous mobility assumption. In the year of the job switch, there is an immediate effect that is both economically significant at seven percent and highly statistically significant. Note that, on average, the move occurs in the summer in year 0, so the full effect of the pay gain is not experienced until year 1, where we see that the effect has doubled. The effect size persists over the following two years.

Turning to the impacts on consumption expenditure, we plot the reduced form coefficient estimates from equation 3, as the solid line in Panel A of Figure 2. The horizontal axis refers to the year relative to the move. Recall that the job switch occurs between -1 and 0 on the horizontal axis, so that the estimates for year 0 should be attenuated by the fact that the worker moves between firms within the year and that earnings and consumption are measured annually. The vertical axis denotes the change in disposable income (solid line) and total expenditure (dashed line) normalized by the level of disposable income the year prior to the move (i.e., t=-1). The vertical lines around each point estimate reflect 95 percent confidence intervals. The first column of Panel A of Table 5 reports the first stage, reduced form coefficient of the change in firm pay on consumption, and the 2SLS estimate of the marginal propensity to consume. The first stage is strong, with a F-statistic far exceeding 100, and the MPC estimate is highly statistically significant,

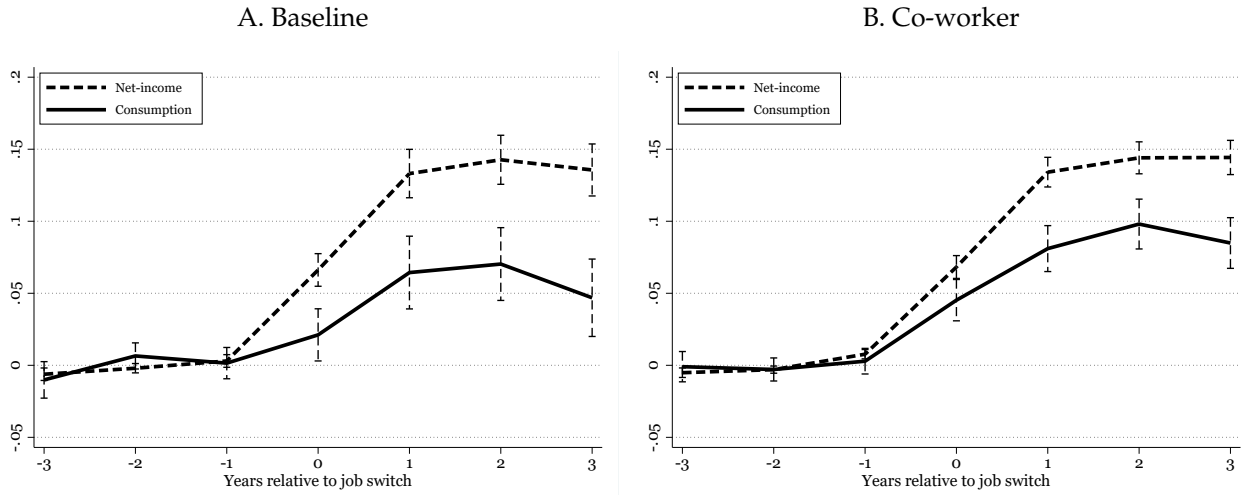
Table 5: MPC out of firm pay premia

	(1)	(2)	(1)	(2)	(3)
	Panel A: Singles		*Panel B*: Couples		
	Split sample	Co-worker	Household	Own	Spouse
Net-income	0.137	0.115	0.219		
(First stage)	(0.008)	(0.006)	(0.062)		
Consumption	0.061	0.058	0.091	0.088	0.003
(Reduced form)	(0.010)	(0.007)	(0.012)	(0.007)	(0.010)
MPC (2sls)	0.443	0.505	0.415	0.403	0.012
	(0.068)	(0.056)	(0.126)	(0.117)	(0.046)
Observations	74976	77605	231706	231706	231706

Notes: The table reports pooled (over years 1, 2, and 3) coefficients from the first stage (2), reduced form (3), and two-stage least squares regressions for singles in Panel A and couples in Panel B. In Panel A, column (1) uses firm effects as our preferred measure of firm pay, measured in log wages from an AKM regression shown in (1). Column (2) uses co-worker average wages, while column (3) uses firm effects from a split-sample AKM. Panel B presents the analysis for couples: column (4) reports total household results, while columns (5) and (6) report results split by own consumption change and spousal consumption change, respectively. The sample consists of full-time workers aged 25-60 who moved between firms between 2012 and 2015. Robust standard errors in parentheses are clustered at the firm level.

While the split-sample estimate of firm fixed effects accounts for small sample bias due to mechanical correlations between the firm fixed effect (identified from movers) and consumption changes of movers, we contrast this design with differences in the average wage of coworkers as the exogenous regressor (and instrument). The second column of Panel A of Table 5 shows that the estimates on income and consumption are very similar to the baseline specification. Moreover, the 2SLS estimate of the MPC using the co-worker wage as an instrument is statistically indistinguishable from the estimate in column (1). Panel B of Figure 2 shows that the time pattern is also mirrored.[8]

Figure 2: Net-income and consumption responses to changes in firm pay



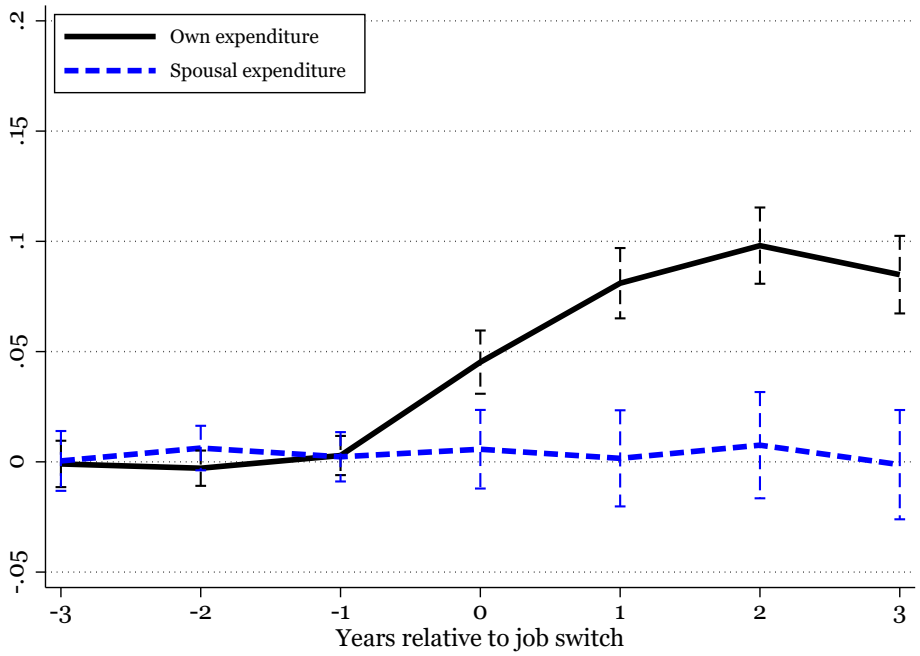
Notes: The figure reports coefficient estimates from the first stage (2) and reduced form (3) equations for the baseline using split-sample firm effects in Panel A, and co-worker wages in Panel B.

Two-Adult Household Consumption. So far, we have focused on a single worker's own consumption following their job switch. The extent to which the person shares the pay premium with other adults in a household, and whether the MPC among couples is different from that of singles, are questions we address next. To do so, we use the sample of individuals who move and live with another partner. This includes both married and cohabiting couples. Our measure of consumption is first aggregated at the household level and then disaggregated to examine the extent to which pay increases are shared among households.

Figure 3 reports coefficient estimates of the effects of a change in firm-pay on the mover's own income, own spending, and the spouse's consumption. We see that the change in income and expenditure of the mover is very similar to that of singles. In contrast, there is

[8]We also perform the exact same regressions on a sample of single-person household, which removes singles with children from the sample. These estimates are reported in Panel A of Appendix Table A1, and are virtually the same.

Figure 3: Household decomposition



Notes: The figure reports coefficient estimates from the reduced form (3) equation for couples split into the mover response and spouse response.

little, if any, change in the spending of the spouse. We report the point estimates, averaged over the years 1-3 in Panel B of Table 5. These confirm that the 2SLS estimates are similar, and not statistically significantly different from the estimate for singles. Moreover, it shows that the point estimate for the spouse allows us to rule out an effect size that is larger than a ten percent increase in spending growth.[9]

Robustness to Other Characteristics of Workers and Firms. This subsection examines the stability of these estimates to the inclusion of worker control variables and challenges the exclusion restriction. So far, our estimating equation has been the normalized change in income and consumption on the left-hand side, with the change in estimated firm-fixed effects on the right-hand side. In Table 6, we show that including individual-level controls does not change the results. The baseline with no control is reported in Column (1), and the individual-level control is in Column (2). Next, we show that adding calendar year fixed effects barely moves the estimated MPC in Column (3).

Finally, we consider whether the firm pay fixed effect estimate may be related to other

[9] Again, we perform the exact same regressions on a sample of multi-person household, which adds singles with children to the sample. These estimates are reported in Panel B of Appendix Table A1, and are statistically indistinguishable from the baseline sample of married and cohabiting individuals.

characteristics of the firm that are important for a workers' expected future income and, therefore, consumption choice. Whereas workers may consume out of changes in pay when switching jobs, workers may also change their spending due to changes in job stability (Pinheiro and Visschers, 2015; Jarosch, 2023). This would invalidate the IV design as the exclusion restriction would not hold if the risk (or disamenity) channel also influences spending. This section examines the relationship between pay differences and differences in unemployment risks, and whether the results for pay differences hold when we directly control for these risk differences. To do so, we calculate firm-wide measures of unemployment risk and assess their direct relationship to consumption, as well as how controlling for it affects our estimates of γ^c. Following Lachowska et al. (2025), we measure unemployment risk at the firm by regressing an indicator for whether an employee of firm j is receiving any unemployment benefits year t or t + 1. Next, we control for permanent worker-level differences in unemployment risk while identifying the firm-level component from workers who experience job loss from more than one employer. Again, the estimated MPC changes only slightly.[10]

Table 6: MPC Robustness Checks

	(1)	(2)	(3)	(4)
MPC 2SLS Estimate	0.443	0.431	0.449	0.447
(standard error)	(0.068)	(0.069)	(0.072)	(0.073)
Worker controls		Yes	Yes	Yes
Year controls			Yes	Yes
Unem. risk controls				Yes
Observations	77605	77605	77605	77504

Notes: The table copies in the baseline specification for 2SLS estimates in Column (1), add age, gender, and number of children indicator controls in Column (2), calendar year indicators in Column (3), and the change in firm-specific unemployment risk component following Lachowska et al. (2025) in Column (4). The sample consists of single full-time workers aged 25-60 who moved between firms between 2012 and 2015. Robust standard errors are clustered at the firm level.

[10] In the Appendix Table A2, we estimate an augmented version of the reduced form equation as in (3), where we include the difference in firm unemployment risk as an explanatory variable $\Delta_j - \Delta_k$ written as

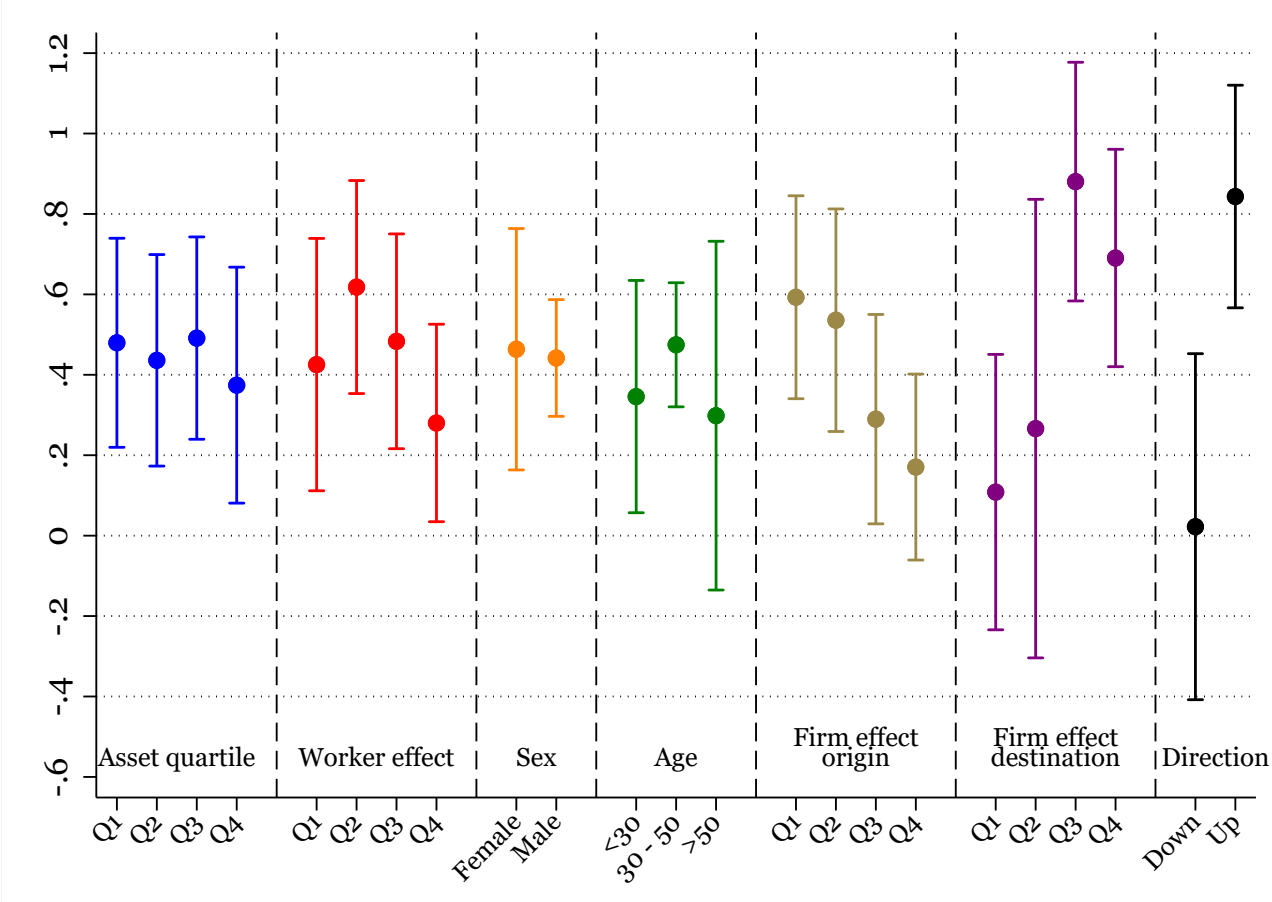
$$\underbrace{c_{i,t} - \bar{c}_{i,s-1}}_{\text{Change in Spending}} = \underbrace{\gamma_t^c[\tilde{\psi}_{j(i),s} - \tilde{\psi}_{k(i),s-1}]}_{\text{Change in Firm Pay Fixed-Effect}} + \underbrace{\gamma_t^{\Delta}[\Delta_{j(i),s} - \Delta_{k(i),s-1}]}_{\text{Change in Firm UI Fixed-Effect}} + \varepsilon_{i,t}. \quad (4)$$

We show that while the coefficient on firm pay does not change, moving to a firm with lower estimated firm fixed effects barely moves, if anything, it slightly enhances the consumption response.

4.3 Heterogeneity

In this subsection, we perform a series of subsample regressions to understand the role of worker- and firm heterogeneity. In Figure 4, we report estimates of MPCs across various splits of the sample, such as by household asset quartiles, worker skill quartiles as measured by the worker wage effect, whether the worker moves down or up the firm pay ladder, the origin quartile of the firm pay distribution, the destination quartile of the firm pay distribution, age, and gender.

Figure 4: MPC heterogeneity across different sample splits for singles



Notes: The figure reports estimates of the MPC out of firm pay following job switches across different sample splits. There are no control variables included. Worker effects are taken from the estimated wage equation. Direction refers to movements up or down in firm pay effects.

Starting with household assets, we note that a standard consumption-savings models imply a higher marginal utility of consumption for the less affluent. We find some indication that low asset households exhibit higher MPCs than those with the highest level of assets. Moreover, we see that, even among workers with assets that likely put them beyond

borrowing constraints, they also display sizable consumption responses. If firm pay policies solely matter through borrowing constraints, we would expect near-zero MPCs among wealthy workers. Instead, the data points to an additional role for job search frictions, where workers cannot perfectly foresee job switching and the subsequent wage changes. We report point estimates, standard errors, level of consumption, and number of observations in thousands in Table 7.

Next, we explore worker heterogeneity by worker-pay types, age, and gender. We find some evidence that workers of the highest earnings types have a somewhat smaller consumption response, although we cannot reject that the estimates are the same as for low-wage types. Similarly, we find no evidence suggesting that women or men have a different response, nor do we find evidence of differential effects across age groups.

Finally, we explore heterogeneity along the job ladder. Starting with upwards and downwards moves, we find that consumption responses are zero when experiencing a negative change in pay premiums. Even though the estimates for downward moves are imprecise, we can reject that the estimate is the same as for workers move to higher-paying firms. Moving up the ladder is associated with an MPC of around 0.8. We also observe that movements from lower-paying firms are associated with greater MPCs as opposed to moving from higher-paying firms, even though some of this effect may reflect that the likelihood of a downward move is higher at the top of the distribution. Concurrently, moving to higher-paying firms is associated with higher MPCs than moving to lower-paying firms.

4.4 Consumption categories

Is the average MPC response broad-based or or is it concentrated in specific consumption categories? The answer to this question has relevance for welfare consequences of pay differentials. For example, consumption changes in categories, such as necessities, like clothing and food, or spending on goods associated with large adjustment costs, are more directly tied to welfare than others (see, e.g., Chetty and Szeidl, 2007). By contrast, spending increases on luxury items and other non-essential goods, such as hotels and restaurants may be less important for workers overall well-being.

We address this question using the richness of our spending data. To this end, we define normalized spending growth in 19 different spending categories. Next, we implement the movers-2SLS strategy for each category and divide the estimated MPC by the budget share to obtain an elasticity, allowing for straightforward comparison across goods. We report the 2SLS estimates, along with 95% confidence intervals, in Figure 5. The results show a broad-based increase in consumption, with a few noteable excepÂ´tions. While spending ef-

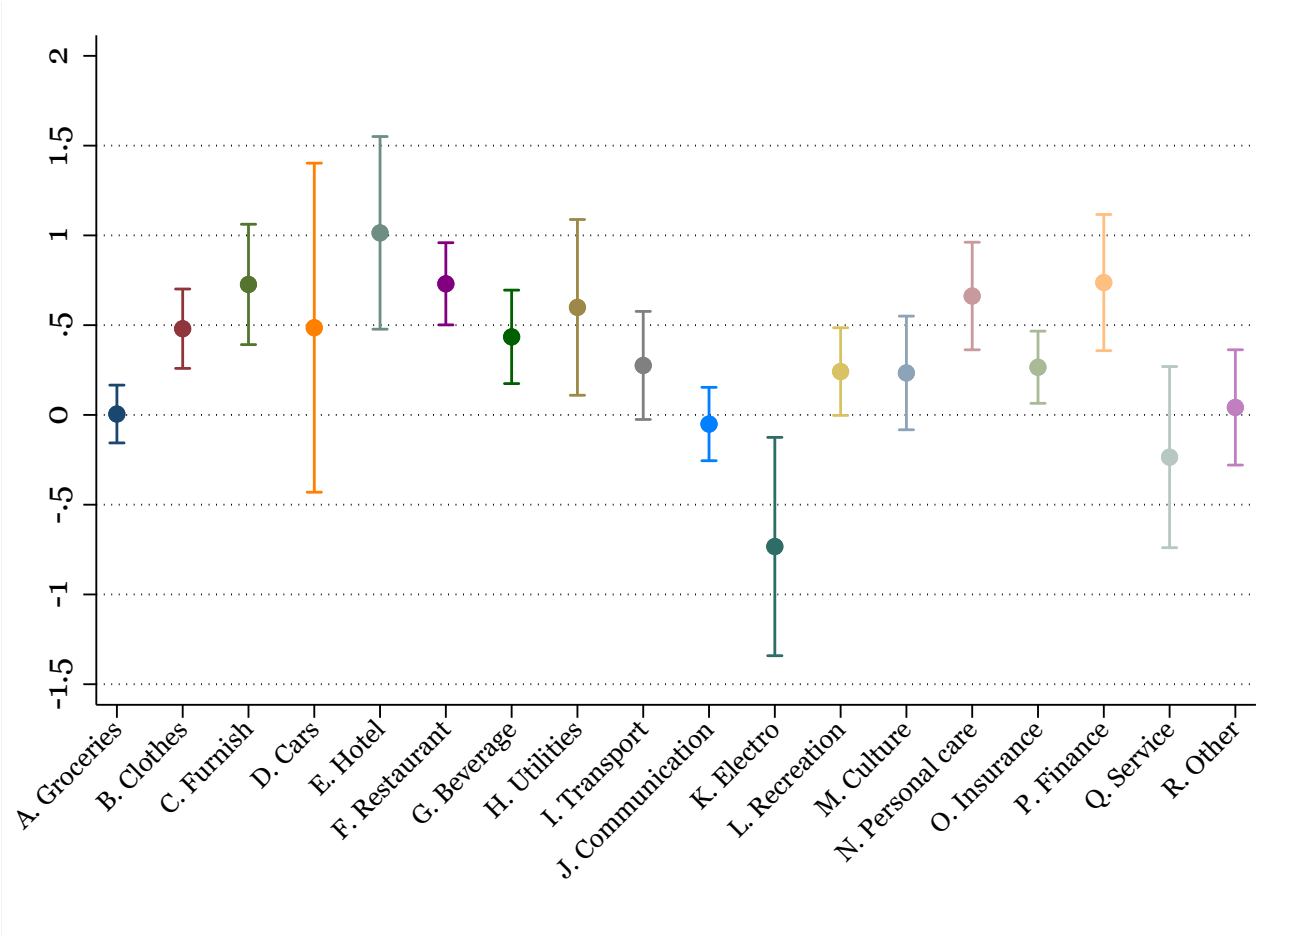
Table 7: MPC Heterogeneity

	(1)	(2)	(3)	(4)		(1)	(2)	(3)	(4)
Panel A: worker types					*Panel B*: firm types				
Subgroup	MPC	SE	E[c]	N	Subgroup	MPC	SE	E[c]	N
Assets					Firm origin				
Q1	0.480	(0.133)	281.4	18.4	Q1	0.593	(0.129)	304.4	19.5
Q2	0.436	(0.134)	333.1	18.9	Q2	0.536	(0.141)	326.8	19.5
Q3	0.491	(0.128)	360.9	19.4	Q3	0.290	(0.133)	355.9	19.4
Q4	0.374	(0.150)	413.8	20.8	Q4	0.171	(0.118)	392.3	19.1
Worker effect					Firm destination				
Q1	0.425	(0.160)	259.3	20.0	Q1	0.108	(0.175)	312.1	19.3
Q2	0.618	(0.135)	308.4	19.7	Q2	0.266	(0.291)	324.7	19.6
Q3	0.483	(0.136)	347.2	19.1	Q3	0.880	(0.151)	345.3	19.5
Q4	0.280	(0.125)	461.2	18.8	Q4	0.690	(0.138)	394.3	19.1
Sex					Direction				
Female	0.464	(0.153)	346.7	27.6	Down	0.022	(0.220)	342.4	39.4
Male	0.442	(0.074)	344.8	50.0	Up	0.843	(0.141)	348.3	38.2
Age									
<30	0.346	(0.147)	280.8	5.5					
30-50	0.475	(0.079)	362.3	48.2					
>50	0.298	(0.221)	372.8	23.9					

Notes: E[c] denotes mean consumption in three years prior to the switch and denominated in thousands of real 2018 NOK. N is the number of observations and also in thousands.

fects are significantly different from zero, including spending in categories, such as clothing, furnishing, and in particular hotels and restaurants, we also note some categories, such as food (or expenditure in groceries), transportation, and communication services (like broadband and mobile phone) are not statistically significantly different from zero. These findings show that pay increases due to job mobility is associated with changes in consumption, but suggest that responses are mainly driven by non-essential consumption goods.

Figure 5: MPC across categories proportional to consumption share for singles



Notes: The figure reports the MPCs across the sub-categories. Estimates are adjusted by the average share of each category of total consumption.

5 Inequality

Having established that employees respond strongly to pay premiums and that these responses are broad based, this section quantifies the extent to which firms matter for consumption inequality. To do so, we lay out a simple revealed preference framework that

allows us to make inference about inequality in welfare without taking a stance on the exact mechanisms.

5.1 Framework

To contribute to understanding the role of firms in shaping inequality in welfare, we develop a revealed preference approach that solely rely on workers' actual consumption behavior to assess the role of firms in welfare inequality. A benefit of this approach is that we can remain agnostic about the underlying mechanisms driving consumption decisions. As the measure of welfare, we apply the consumption equivalent variation (CEV) to quantify contemporaneous utility differences across firms: how much consumption must worker i give up (gain) in firm j in order to move to firm k such that the current period utility is equalized? Let this CEV measure, γ_{ijk}, solve

$$u(c_{ij}(1 + \gamma_{ijk})) = u(c_{ik}), \tag{5}$$

where u is the worker's utility function. We make the following additional assumptions. First, we require that $u(\cdot)$ is an increasing function. Note that we have not specified any particular functional form, and that the utility function can, in principle, differ between individuals. Second, we assume that log(c) is additively separable in a worker effect and a firm effect. This additive separability implies homogeneous treatment effects, meaning that $\gamma_{ijk} = \gamma_{jk}$ in (5). While this is a restrictive assumption, we consider it a natural benchmark and evaluate how mis-specification may alter our conclusions in the next section.[11] With these assumptions, we may write

$$u(c_{ij}(1 + \gamma_{jk})) = u(c_{ik}) \implies c_{ij}(1 + \gamma_{jk}) = c_{ik}. \tag{6}$$

The fact that utility is increasing in consumption implies that equalizing welfare across firms requires equalizing consumption across firms, as shown following the implication. Given a baseline firm k, simple manipulation allows us to write

$$var_j(\gamma_{jk}) \approx var_j(\log(c_{ij})) = var_j(\psi^c_j), \tag{7}$$

where the approximation comes from $\log(1 + \gamma_{jk}) \approx \gamma_{jk}$. Given a log-additive consumption equation, taking the variance of $\log(c_{ij})$ over j is equal to the variance of firm consumption effects. Equation (7) states that the variance of firm consumption effects is equivalent to the

[11] For example, the additive separability between individual and firm fixed effects in consumption is violated when workers' asset holdings and skill determine the consumption response from switching firms.

variance of welfare across firms.[12] Imposing the standard assumption that wages are also log-additively separable, it is straightforward to write down the following ratio

$$\frac{\text{var}_j(\gamma_{jk})}{\text{var}_j(\gamma_{jk}^w)} \approx \frac{\text{var}_j(\psi_j^c)}{\text{var}_j(\psi_j^w)}, \tag{8}$$

where γ_{jk}^w is a welfare measure that solves (15) but where we replace c with w. The ratio of firm effects from consumption and wages informs us directly how much one would overstate (or understate) static welfare inequality if we were to focus on wages as opposed to consumption. We will see next that the variance of consumption is lower empirically, and show in the final section that an off-the-shelf model of search with consumption savings decisions is fully consistent with this.

5.2 Variance Decomposition

Next, we perform a variance decomposition using the log of wages and consumption expenditure. To assess the role of firms and workers in shaping inequality, we follow the two-way fixed effects framework of Abowd et al. (1999) and apply the bias-corrected variance decomposition estimator of Kline et al. (2020) (henceforth, KSS). However, we must impose some additional data restrictions to facilitate this analysis.

First, while our data covers the vast majority of electronic payments, a key missing object is transactions made with cash. The extent to which a household relying on cash for expenditure is a fixed characteristic, the person- or household fixed effect would explain an overly large fraction of the total variation in spending.[13] To minimize the influence of these observations, our baseline variance decomposition is conducted on a trimmed sample, where we remove entirely from the sample workers whose household consumption ever drop below the 5th percentile of the consumption distribution. Appendix Figure 4(b) plots the consumption-to-wage ratios on this sample. Reassuringly the ratio is now much more stable across the consumption distribution.[14] Second, we focus on single-person household to filter out variation in expenditure due to household size.

[12] Importantly, note that the choice of the benchmark firm k is innocuous. This simply adds a constant to $\log(1 + \gamma_{jk})$ and does not alter the variance.

[13] To assess this concern, we examine the distribution of log consumption relative to log gross wages. Appendix Figure A3 shows that consumption exhibits a longer and fatter left tail than wages. Moreover, sorting workers by percentile of consumption in Appendix Figure 4(a), reveals a sizable portion of very low consumption-to-wage workers at the bottom of the consumption distribution, indicating that low consumption is not due to low wages, but likely cash as the mode of payment.

[14] We also show that the relative variance of consumption expenditure to wage income remains stable across the trimming vs winsorizing approaches over time.

Results Table 8 reports variance decompositions for log earnings and log expenditures among singles, contrasting standard AKM estimator to the bias-corrected KSS estimator (Kline et al., 2020). Consistent with limited mobility bias inflating the role of firm effects, we find that the firm effect accounts for a larger share in the standard AKM decomposition – particularly pronounced in the decomposition of log expenditures. The role of firms shrinks from 0.02 in the AKM model to 0.011 in KSS, where the correlation between firms and worker-fixed effects goes from negative to 0.22. Similarly, the variance of firm effects for consumption shrinks from 0.017 to 0.006 when applying the KSS approach, where the correlation between firms and workers is 0.068. By contrast, a much larger share of the explained variance in both outcomes is attributable to person effects. The KSS estimates indicating that 76.8% of the variance in log earnings and 93.2% in log expenditures is due to persistent individual heterogeneity.

Table 8: Variance Decompositions Single-Person Households

	log earnings		log expenditure	
	AKM	KSS	AKM	KSS
variance	0.116	0.108	0.158	0.138
variance of firm effects	0.020 (0.170)	0.011 (0.104)	0.017 (0.105)	0.006 (0.041)
variance of person effects	0.098 (0.840)	0.083 (0.769)	0.156 (0.990)	0.128 (0.932)
covariance of firm and person effects	-0.001 (-0.010)	0.007 (0.127)	-0.008 (-0.096)	0.002 (0.027)
correlation of worker and person effects	-0.014	0.225	-0.148	0.068
R squared	0.862	0.799	0.719	0.628

Notes: Variance decomposition of log earnings and log expenditures for single-person households on the leave-one-out connected set. Log earnings measured at the individual level, log expenditures at the household level (equivalized). Wages winsorized at the 1st and 99th percentile, consumption at the 99th. We then compute the 5th percentile of consumption and drop all observations for workers whose consumption ever drop below this threshold. 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 refer to fraction of explained variance.

The key take-away from this decomposition is that the variance of firm effects on consumption (0.006 under KSS) is roughly half the size of the variance of firm effects on earnings (0.011 under KSS). Through the lens of the consumption equivalent variation appraoch, the evidence suggests that firms play a nontrivial role in shaping welfare, although the role for household consumption is considerably attenuated relative to wages.

We contrast the variance decomposition results to the raw data and a version of the data where we include workers with missing consumption, but winsorize at the 5th percentile. We report summary statistics of these samples in Appendix Table A4, and the results from KSS variance decompositions in Panel A Appendix Table A5. Notably, the ratio of the variance of firm effects in consumption and wages is remarkably similar across specifications: the variance of firm effects on consumption is roughly half the size of the corresponding variance for wages, across all specifications. This suggests that while firms influence both earnings and consumption, the fraction of explained variance depends on how we deal with worker heterogeneity. Excluding low-consumption households primarily reduces the variance of the person effect, which drops sharply in the trimmed sample, suggesting that measurement error in spending is a persistent trait at the household level.

Lastly, we extend the analysis to all workers, including single- and multi-person households. We perform variance decompositions using individual earnings and consumption expenditure, and a household member equivalent, where we divide earnings and consumption by the square root of the number of household members. Both samples produce similar estimates for the role of firms, reported in Panel B and C of Appendix Table A5.

6 Model

This section demonstrates that an off-the-shelf model incorporating search frictions and consumption savings can rationalize our empirical results. We use the calibrated model to relax the log-linear assumption on consumption, and that our revealed preference approach delivers a close approximation to the variance of firms for welfare inequality.

6.1 Environment

A continuum of measure 1 workers and a continuum of measure 1 firms populate the economy. Workers discount the future at rate $\hat{\rho}$ and leave the model at rate d, which combine to produce a social discount rate of $\rho = \hat{\rho} + 1/d$.

Workers derive utility, u, from consumption, c. We assume that the workers utility func-

tion is of the CRRA variety written as

$$u(c) = \frac{c^{1-\alpha} - 1}{1-\alpha}. \tag{9}$$

Workers can save. Consumption is chosen by maximizing the workers' expected present discounted stream of consumption

$$\mathbb{E}_0 \int_{t=0}^{\infty} e^{-\rho t} u(c_t) dt, \tag{10}$$

where the evolution of assets follows

$$\dot{a} = ra + i - c. \tag{11}$$

There is a two-sided heterogeneity. Workers are indexed by a permanent ability level, $\varepsilon \in [1-\Delta_\varepsilon, 1+\Delta_\varepsilon]$, and firms by a permanent productivity level $z \in [0,1]$. Workers differ ex ante due to ability differences, but also differ ex post within each ability level due to different assets resulting from their varying employment histories.

Search frictions are present. Workers search on and off the job. During unemployment, workers receive offers at rate λ and at rate $s\lambda$ while employed. The offer is the firm ability level, ε, drawn from an offer distribution F, which describes the job ladder. We also allow for reallocation shocks whilst employed which occur at rate $s^r\lambda$, where workers are forced to "take it or leave it".[15]

We assume production is multiplicative in the worker and firm type. Wages are determined in the bargaining setup following Binmore et al. (1986), Hall and Milgrom (2008) and Elsby and Gottfries (2022) where threats to disruption are only credible. We provide a complete description of the two stage bargaining game in the Appendix, but the upshot is a wage equation of the following form

$$w(\varepsilon, z) = \beta \varepsilon z + \omega_0, \tag{12}$$

where β can be interpreted as worker bargaining power,[16] and ω_0 collects terms describing the cost of disruption to the worker and firm. There are three other common routes to take when determining wages. First, we could have adopted the standard Nash bargaining

[15] These shocks are included since downward wage moves are an important feature of the data, but also for connected set arguments when running two-way fixed effect regressions. Take it or leave it implies that the worker's outside option becomes unemployment following the shock at rate $s^r\lambda$.

[16] But is more precisely the probability that the worker makes the offer in the second stage.

setup where workers and firms in some way split the surplus of the match Pissarides (2000). Second, we could have considered the wage-posting game in Burdett and Mortensen (1998) where wage dispersion is endogenous due to the dynamics of the job ladder. Finally, in the sequential auctioning framework of Postel-Vinay and Robin (2002). We view the bargaining setup used as a simplification of the Pissarides (2000), but that does not come with it concerns over non-convexity of the bargaining set in models of on-the-job search (Shimer, 2006). Moreover, our setup complicates the wage-posting game due to heterogeneity in the reservation wage by workers in unemployment as a result of differences in assets.[17]

Consumption choice and search frictions

This concludes the setup of the model. In short, it resembles Lise (2013), but we include worker heterogeneity to enable us to consider worker effects in an AKM setup. In the next section we will target the MPC and AKM decomposition results, and use the model structure to carefully quantify the importance of firms for welfare inequality, and to what extent workers mitigate firm driven wage inequality. The consumption choice is critical to this, which we now describe in this class of consumption savings models.

Consumption solves the Euler equation written as

$$u'(c) = V_a \implies c(\varepsilon, z, a) = \varkappa\big(V_a(\varepsilon, z, a)\big), \tag{13}$$

where V_a is the partial derivative of the present discounted value of the employment to the worker with respect to assets - the shadow value of wealth. Consumption is a function of the shadow value of wealth where $\varkappa$ is the inverse function of the marginal utility of consumption. We make explicit that $c(\cdot)$ depends on worker characteristics, firm characteristics, and the workers' asset level. Consumption is not additively separable into an individual effect and firm effect. In particular, the consumption choice depends on a worker's asset level which is heterogeneous between workers of identical skill, and interacts with the firm the worker is in - treatment effects are heterogeneous in worker assets. While this is true, later in this section we will see that the model produces very similar welfare results to those described in the previous section where we take into account this heterogeneity in treatment effects.

The model also admits a simple representation of the MPC which can be written as

$$c_w = \frac{c_a + c_{wa}\dot{a}}{\tilde{\rho} - r + c_a\left(1 - \dot{a}\frac{u'''(c)}{u''(c)}\right)}, \tag{14}$$

[17] We note that one could circumvent this with a binding minimum wage for all asset levels.

where $\tilde{\rho} = \rho + s^e\lambda[1 - F(p)] + s^r\lambda + \delta$ is the implied social discount rate. There are two properties of the model-implied MPC that we wish to highlight. First, holding all else equal, the MPC is decreasing in the origin firms productivity. This is consistent with our heterogeneity analysis, which shows that movements away from firms lower down the job ladder are associated with larger MPCs (the analogous is also true). Secondly, the consumption response depends critically on the implied social discount rate $\tilde{\rho}$. A higher $\tilde{\rho}$ reflects greater impatience, but also a greater separation hazard, captured by the arrival rate of job offers, λ. When λ is low and job offers arrive infrequently, workers cannot anticipate future transitions and therefore struggle to smooth consumption across jobs. This leads to discrete jumps in consumption following job switches. As λ increases, the consumption response diminishes, since the worker's current job becomes less important for determining lifetime earnings. Thus, search frictions play a central role in shaping consumption behavior following job transitions and are key to understanding how firm-level pay heterogeneity contributes to welfare inequality.

6.2 Calibration

We calibrate the model to the Norwegian economy and choose certain parameters to replicate the variance of firm effects for consumption and wages, and our estimated MPC along the job ladder. In doing so, our model is primed for carefully considering how firm pay heterogeneity impacts welfare inequality, which we will consider in the final subsection.

To calibrate the model, we make two parametric assumptions on the firm and worker distributions. We parameterize the offer distribution to be a Beta distribution, and the worker distribution to be uniform. This results in 14 parameters. We set the $\hat{\rho}$ to equal a 5% annual social discount rate, d to represent a 40 year working life on average, and an interest rate, r, of 3% annually. We choose the worker bargaining power, β, and the cost of disruption, ω_0, to both be 0.5, close to Elsby and Gottfries (2022).

We calibrate the remaining 9 parameters to reproduce standard moments from Norway and our new non-standard moments from Norway. First, λ is chosen to reproduce the unemployment to employment rate. δ is chosen to reproduce the employment to unemployment rate. s and s^r are chosen to reproduce the job-to-job rate, and the proportion of job switching that results in wage cuts. Finally, we choose b to replicate the average replacement rate.

The remaining parameters are chosen to replicate statistics from the previous sections. Broadly, we chose the CRRA parameter, α, and the 2 parameters describing firm productivity heterogeneity, β_z^1 and β_z^2, to reproduce the average MPC, and the variance of firm consumption and wage effects which we will further discuss later in this section. Finally,

we choose the dispersion in worker ability, Δ_ε, to reproduce the variance of person wage effects. See Appendix Section C for further details on the calibration process.

Table 9 shows the model fit, we reproduce all moments almost perfectly. We also show the dynamics of the MPC in Figure A7. We observe no change in consumption prior to the switch and the large and persistent jump afterward consistent with our data. A key target is our estimated MPC and firm consumption effects from the previous sections. To generate these responses, the value of the CRRA coefficient is crucial. In Figure A6 we show that the MPC and variance of firm consumption effects relative to wages are both strongly decreasing in α. As we approach log utility ($\alpha \to 1$), the MPC and ratio of variances approach 1. On the other hand, the two measures are around 0.2 for $\alpha = 5$. This is intuitive: as workers become more risk averse they more aggressively save for periods of unemployment at the top of the job ladder, thus reducing the pass through of wage inequality to consumption inequality, and also reducing the MPC out of job moves. We find a value of around 2.3. This is close to the standard value of 2 often used in the literature, and well within the range of values reported in Landais and Spinnewijn (2021). Thus, a calibrated off the shelf search model with a sensible degree of risk aversion is fully consistent with our empirical findings from the previous sections.

Table 9: Calibrated parameters and model fit

Moments	Data	Model	Parameters	Value
UE rate	0.20	0.20	λ	0.064
EU rate	0.02	0.02	δ	0.02
JJ rate	0.02	0.02	s	1.01
Wage cuts	0.33	0.34	s^r	0.35
Replacement rate	0.50	0.49	b	0.52
MPC	0.50	0.51	α	2.31
Firm c effect variance	0.006	0.005	β_ε^1	0.40
Firm w effect variance	0.011	0.012	β_ε^2	3.44
Person w effect variance	0.084	0.084	Δ_α	0.39

Notes: The table reports data moments, model moments, parameters and parameter values. The parameter values are positioned next to the moments that are particular informative for them. Although, as discussed in the text, α informs both the MPC and the variance of firm consumption effects.

6.3 Welfare Inequality in the Model

In Section 5 we showed that the variance of firm effects in log consumption corresponds directly to the role of firms in explaining welfare inequality. However, this result relied on the assumption that when workers move from firm j to firm k, all workers increase consumption by the same proportional amount - a homogeneous treatment effect. In the typical consumption saving model, like that presented in the present paper, consumption is not log-additive. Instead the CEV gains from moving between firms is dependent on worker characteristics, and solves

$$u(c_{ij}(1+\gamma_{ijk})) = u(c_{ik}) \implies c_{ij}(1+\gamma_{ijk}) = c_{ik}. \tag{15}$$

While the static case is the better comparison to the variance of firm wages and firm consumption effects, we can also use the model structure to take into account the future path of consumption. The dynamic comparison solves

$$\mathbb{E}_0 \int_0^\infty e^{-\rho t} u\Big(c_{ijt}(1+\gamma_{ijk})\Big)\, dt = \mathbb{E}_0 \int_0^\infty e^{-\rho t} u(c_{ikt}) dt, \tag{16}$$

where now we equalize the *present discounted* utility between firms j and k.

Since the consumption-equivalent welfare gains from moving across firms - denoted by γ_{ijk} - are heterogeneous and depend on individual i, we must aggregate across individuals within a firm. This means that the underlying distribution of workers (i.e., heterogeneity in ability and assets) matters for the importance of firm pay heterogeneity in welfare inequality. Moreover, the choice of destination firm is no longer innocuous. This is because the consumption equation is not log-additive: the growth in consumption from switching firms depends on both the individual's wealth and their skill.

We therefore compute the firm-level aggregate CEV as

$$\gamma_k = \iint g_{ij}\, \gamma_{ijk}\, di\, dj, \tag{17}$$

where g_{ij} denotes the ergodic joint distribution of individuals i (ability and asset level) in firm j. This expression gives the average CEV gains or costs of moving a worker, drawn from the ergodic population distribution, and moving them to firm k. Firms at the bottom of the wage distribution will have negative γ_k, reflecting that movements to these firms entail welfare losses, while firms at the top will have positive γ_k, reflecting welfare gains.

Table 10 reports the results. In Panel A, for comparison, we report in row 1 the model-implied variance of firm wage and consumption effects, and in row 2 the ratio of each to the variance of firm wage effects. The remaining panels follow the same convention. In Panel

B, row 1, we report the variance of "static" welfare inequality, analogous to Section 5, when taking wages as the input to utility and when taking consumption as the input to utility. The variance of the wage based welfare measure is very close to the variance of firm wage effects and, importantly, the variance of the consumption based welfare measure if very close to the variance of firm consumption effects. Neither the wage nor the consumption equations are log-additive; nevertheless, the variance of welfare measures accounting for this non-additivity in the model (shown in Panel B) is very close to the variance of firm wage and consumption effects estimated under additive separability. According to the model, moving one standard deviation up the firm pay ladder, conditional on the distribution of workers in the firm, results in average CEV gains of around 7%. Moreover, firms appear less important for consumption inequality than for wages because of workers' optimal consumption decisions, which compress the "firm consumption distribution." This optimal behavior reduces the impact of firm pay heterogeneity on welfare inequality by roughly a half.

Finally, Panel C shows the variance of welfare driven by firms when taking into account the future path of consumption. We see that the variance of welfare as a result of firm heterogeneity shrinks considerably. This is very natural, workers at the bottom of the wage distribution at any point in time use their search behavior to climb the ladder reducing the consequences of being at the lowest paying firm. Moreover, workers at the top of the distribution do not expect to stay there forever. Thus, job mobility and job loss mean job ladder driven welfare heterogeneity are muted. Nevertheless, consistent with our empirical results, the ratio of the variance of consumption based welfare to the variance of wage based welfare is still significantly lower than one, again highlighting that optimal consumption decisions of workers reduces the welfare inequality driven by firms.

Table 10: Firm Effects and Welfare Inequality

	A. Firm Effects		B. Welfare (Static)		C. Welfare (Dynamic)	
Variance	Wages	Consumption	Wages	Consumption	Wages	Consumption
Level	1.18	0.46	1.06	0.43	0.10	0.07
Cons/Wages	1.00	0.39	1.00	0.41	1.00	0.71

Notes: This table reports the variance$\times$100 and variance of each column to the variance of wages in each section. Section A refers to the model calibrated variance of firm consumption effects. Section B refers to static welfare inequality analogous to Section 5. Section C refers to full dynamic welfare inequality. All figures are from the model simulations.

7 Conclusion

A large literature has established that firms matter for workers' pay. In this paper, we have examined whether firms affect consumption behavior of workers. This question has important positive implications, as it informs us about the importance of search frictions and workers' income expectations. It also has normative implications for household welfare. However, making progress on this quesiton has proven difficult due to several empirical challenges.

This paper circumvents the measurement and conceptual issues by combining a revealed preference approach with newly collected data on consumption expenditure and employment from Norway. Using a movers-based IV design, we show that firm pay differentials lead to important consumption differentials: each $1 in pay inequality leads to 45-cent differences in consumption. We found that the key mitigating factor is smoothing behavior, as we found no change in consumption as workers move down the pay ladder, while the spending response to upward moves was almost one-to-one. We identified important nuances in the data, for example, that pay differentials primarily cause non-essential spending differentials, such as those for hotels and restaurants. In contrast, pay differentials had little impact on expenditure on groceries, which remained unresponsive to pay differentials.

Next, we performed a variance decomposition to quantify the relative importance of firms in shaping consumption vs. wage inequality. We found that firm effects explained roughly half as much of consumption inequality as of wage inequality, indicating that household spending decisions compress part of the variation in pay. Lastly, we calibrated a simple search-and-savings model to show that it reproduces these patterns and supports a revealed-preference interpretation of the firm consumption effects. Overall, our results show that firms are an important source of welfare inequality. At the same time, the firm's impact on household well-being is mitigated through workers' consumption-smoothing behavior.

References

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

ACHDOU, Y., J. HAN, J.-M. LASRY, P.-L. LIONS, AND B. MOLL (2022): "Income and wealth distribution in macroeconomics: A continuous-time approach," *The review of economic studies*, 89, 45–86.

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.

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.

BINMORE, K., A. RUBINSTEIN, AND A. WOLINSKY (1986): "The Nash bargaining solution in economic modelling," *The RAND Journal of Economics*, 176–188.

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.

BLUNDELL, R. AND I. PRESTON (1998): "Consumption inequality and income uncertainty," *The Quarterly Journal of Economics*, 113, 603–640.

BOROVICKOVÁ, K. AND R. SHIMER (2024): "Assortative Matching and Wages: The Role of Selection," .

BURDETT, K. AND D. T. MORTENSEN (1998): "Wage Differentials, Employer Size, and Unemployment," *International Economic Review*, 39, 257–273.

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

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.

CHETTY, R. AND A. SZEIDL (2007): "Consumption commitments and risk preferences," *The Quarterly Journal of Economics*, 122, 831–877.

CONLON, J. J., L. PILOSSOPH, M. WISWALL, AND B. ZAFAR (2018): "Labor market search with imperfect information and learning," Tech. rep., National Bureau of Economic Research.

CUTLER, D. M. AND L. F. KATZ (1992): "Rising Inequality? Changes in the Distribution of Income and Consumption in the 1980's," *American Economic Review*, 82, 546–551.

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

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.

FAGERENG, A., M. B. HOLM, AND G. J. NATVIK (2021): "MPC heterogeneity and household balance sheets," *American Economic Journal: Macroeconomics*, 13, 1–54.

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

GANONG, P., D. JONES, P. J. NOEL, F. E. GREIG, D. FARRELL, AND C. WHEAT (2020): "Wealth, Race, and Consumption Smoothing of Typical Income Shocks," Tech. rep., 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.

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

GOURIEROUX, C., A. MONFORT, AND E. RENAULT (1993): "Indirect inference," *Journal of applied econometrics*, 8, S85–S118.

GRINDAKER, M. AND M. SIMMONS (2025): "Unemployment Insurance and Worker Reallocation," .

HALL, R. E. AND P. R. MILGROM (2008): "The limited influence of unemployment on the wage bargain," *American economic review*, 98, 1653–1674.

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

KAPLAN, G. AND G. L. VIOLANTE (2014): "A model of the consumption response to fiscal stimulus payments," *Econometrica*, 82, 1199–1239.

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*.

LACHOWSKA, M., I. SORKIN, AND S. A. WOODBURY (2025): "Employers and Unemployment Insurance Take-Up," *American Economic Review*, 115, 2529–2573.

LAMADON, T., J. LISE, C. MEGHIR, AND J.-M. ROBIN (2024): "Labor market matching, wages, and amenities," Tech. rep., National Bureau of Economic Research.

LANDAIS, C. AND J. SPINNEWIJN (2021): "The value of unemployment insurance," *The Review of Economic Studies*, 88, 3041–3085.

LENTZ, R., S. PIYAPROMDEE, AND J.-M. ROBIN (2023): "The anatomy of sortingâ€"evidence from danish data," *Econometrica*, 91, 2409–2455.

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

MORCHIO, I. AND C. MOSER (2024): "The gender pay gap: Micro sources and macro consequences," Tech. rep., National Bureau of Economic Research.

PINHEIRO, R. AND L. VISSCHERS (2015): "Unemployment risk and wage differentials," *Journal of Economic Theory*, 157, 397–424.

PISSARIDES, C. A. (2000): *Equilibrium unemployment theory*, MIT press.

POSTEL-VINAY, F. AND J.-M. ROBIN (2002): "Equilibrium wage dispersion with worker and employer heterogeneity," *Econometrica*, 70, 2295–2350.

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

SORKIN, I. (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.

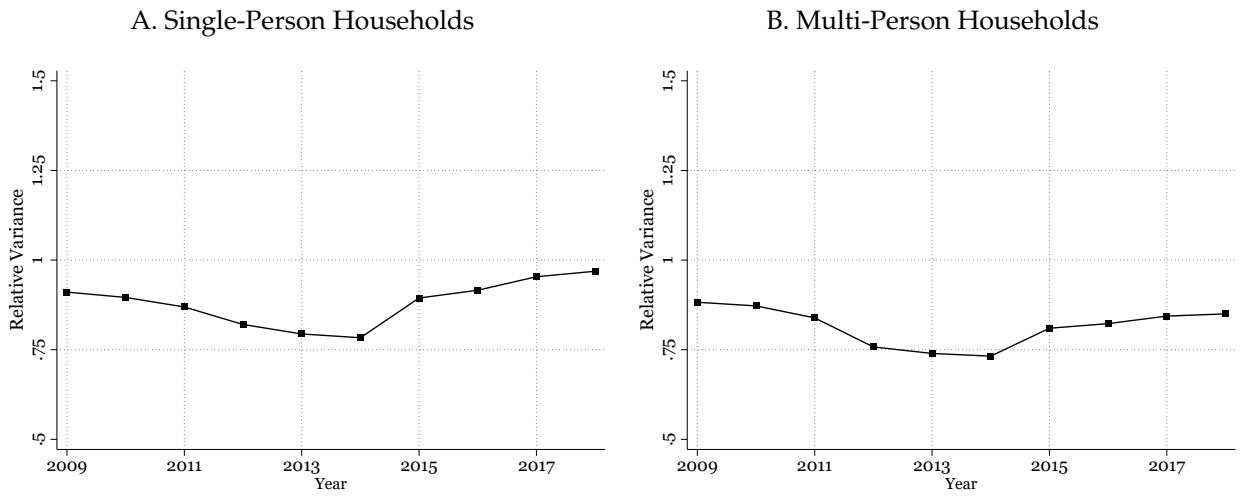
YAGAN, D. (2019): "Employment Hysteresis from the Great Recession," *Journal of Political Economy*, 127, 2505–2558, publisher: University of Chicago.

Online Appendix

A Additional Material

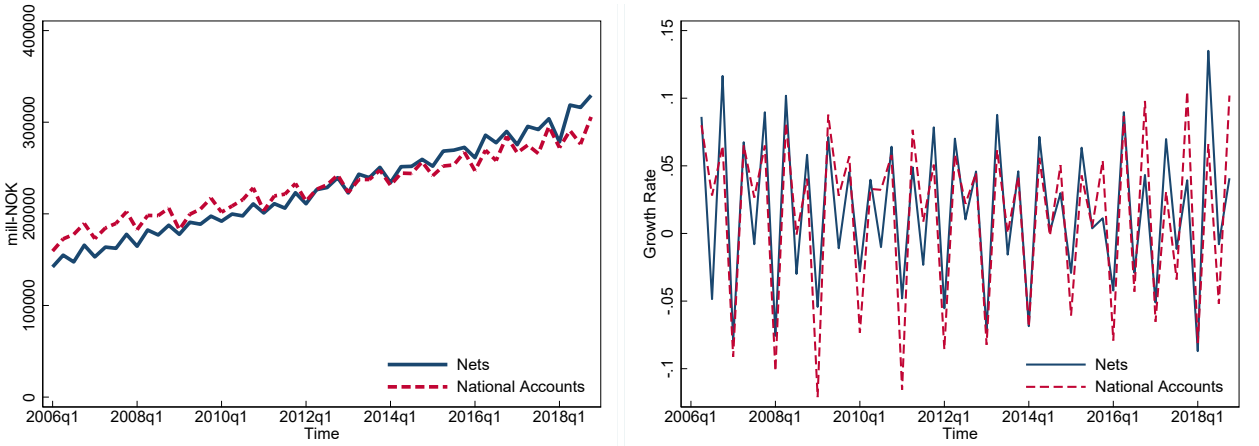
A.1 Figures

Figure A1: Relative Variance of Consumption



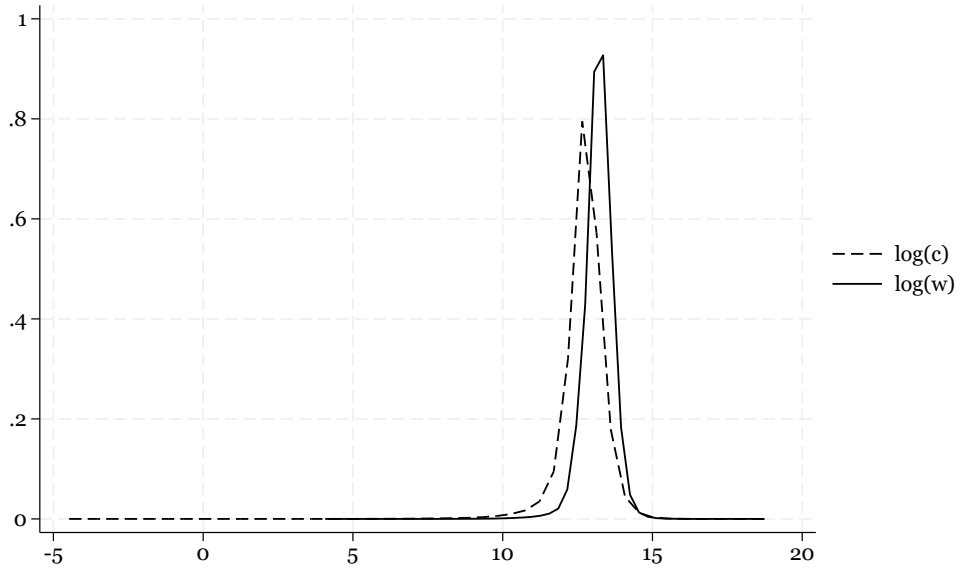
Notes: This figure plots the time series of the ratio of cross-sectional variance of consumption to the variance of wage income. Left-hand-side panel restricts the sample to single-person households, and corresponds to the first two rows of the single-person column in Table 1. Right-hand-side panel restricts to multi-person households, and corresponds to the first two rows of the multi-person column in Table 1.

Figure A2: Aggregated micro-level expenditure data compared with National Accounts data.



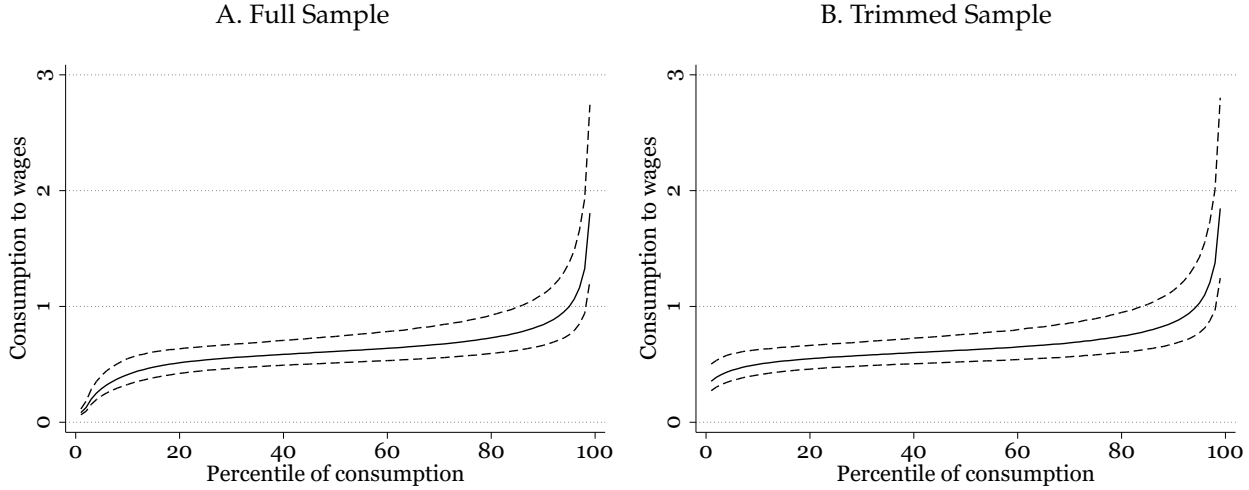
Notes: The figures compare the level of consumption from our data and the national accounts, and the quarterly growth rates from 2006 to 2018. Our data has attrition in bank account links before 2012.

Figure A3: Density of Log Consumption Expenditure and Gross Wages



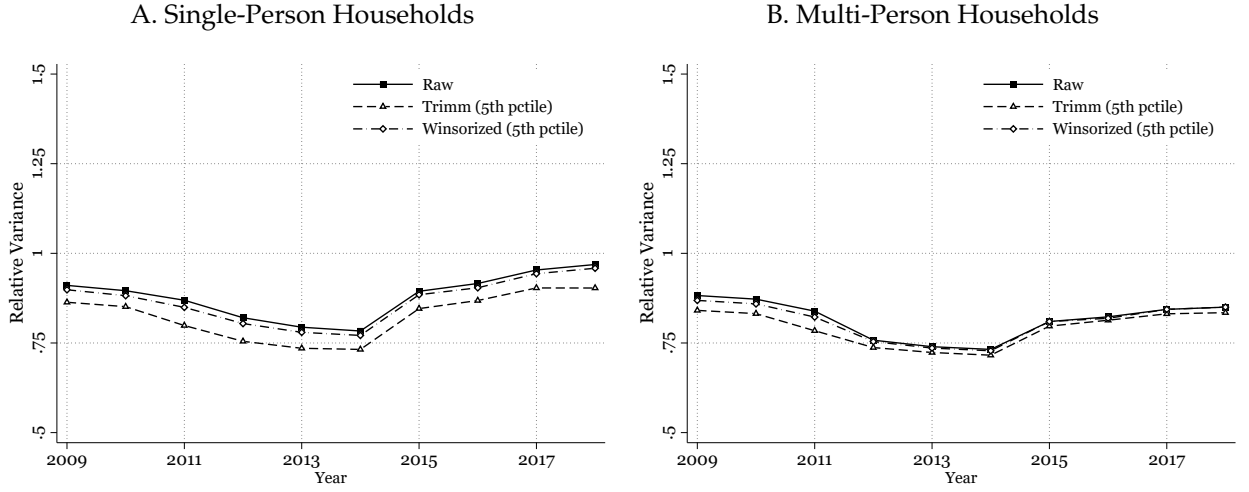
Notes: The figure plots the density of log consumption expenditure and gross wages for single-person households.

Figure A4: Consumption-to-Wages Across Consumption



Notes: This figure plots the median and interquartile range of consumption to wage income ratios, within percentiles of consumption income, for the full distribution of workers. The left panel is for the full sample (consumption and wages winsorized at the 1st and 99th percentile). The right panel is for the trimmed sample. Both wages and consumption are aggregated to the household level, and the unit of observation is the worker.

Figure A5: Relative Variance of Consumption (across samples)



Notes: This figure plots the time series of the ratio of cross-sectional variance of consumption to the variance of wage income, for various ways of cleaning the consumption data. Left-hand-side panel restricts the sample to single-person households. Right-hand-side panel restricts to multi-person households.

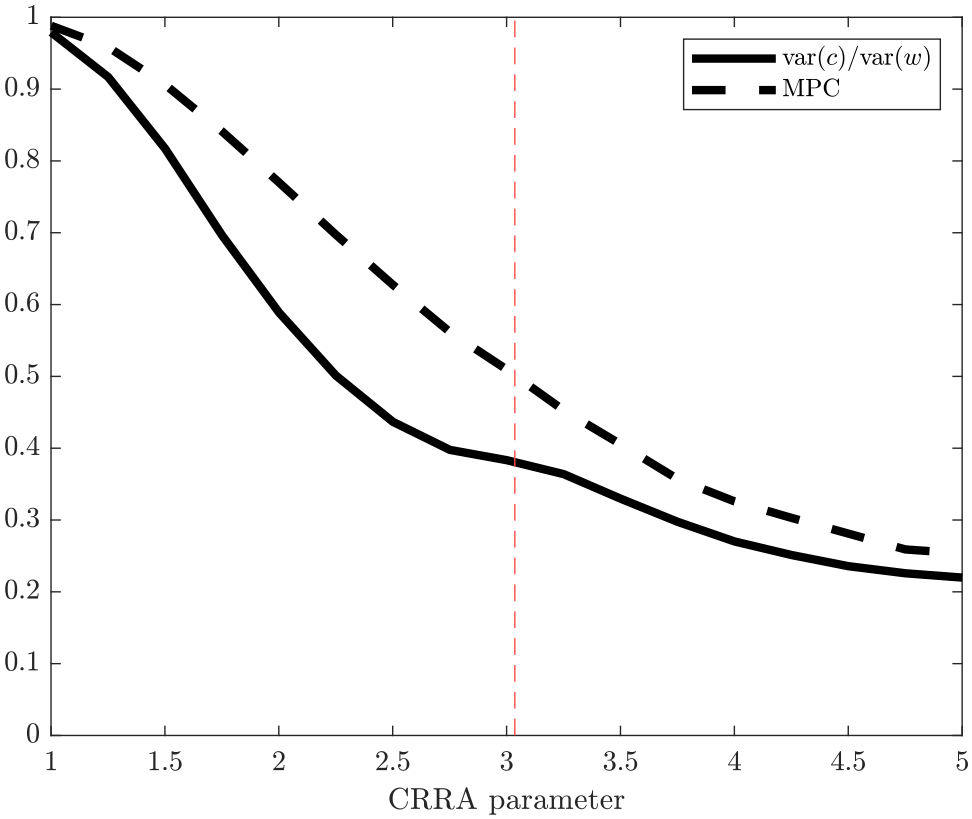


Figure A6: Varying the CRRA parameter

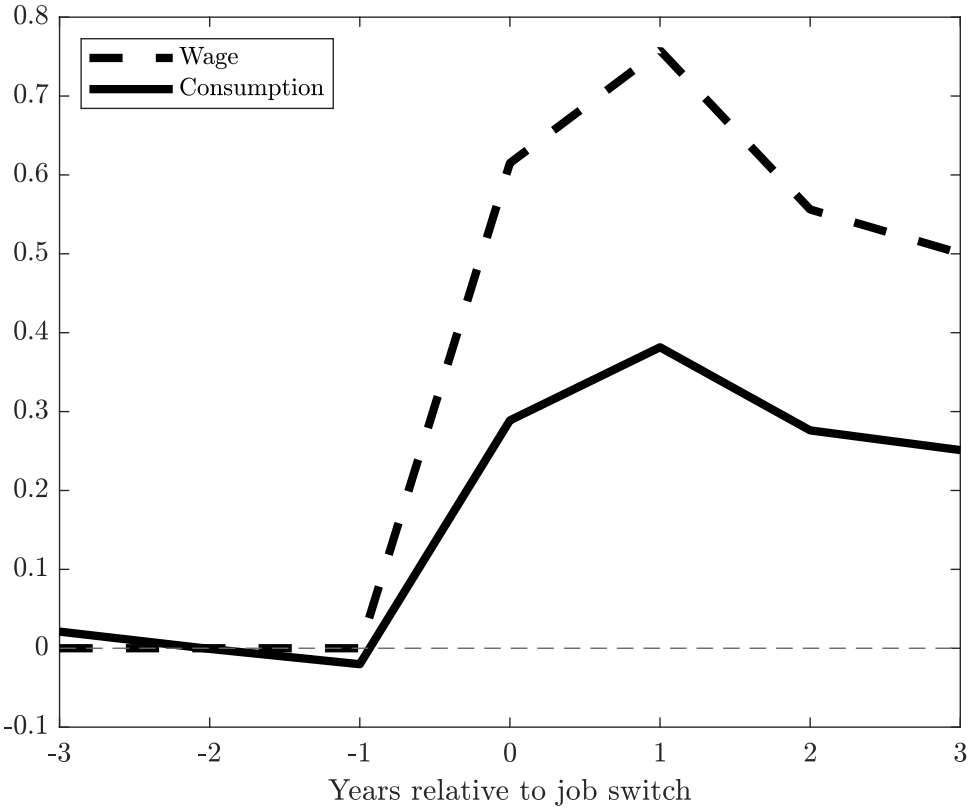


Figure A7: Model MPC

A.2 Tables

Table A1: MPC out of firm pay premia with alternative household definitions

	(1)	(2)	(1)	(2)	(3)
	Panel A: Single households		*Panel B*: Multi person households		
	Split sample	Co-worker	Household	Own	Spouse
Net-income	0.151	0.125	0.205		
(First stage)	(0.009)	(0.007)	(0.056)		
Consumption	0.068	0.066	0.083	0.084	-0.001
(Reduced form)	(0.011)	(0.008)	(0.012)	(0.007)	(0.009)
MPC (2sls)	0.448	0.526	0.403	0.409	-0.006
	(0.070)	(0.057)	(0.120)	(0.115)	(0.046)
N	54372	56356	252955	252955	252955

Notes: The difference between this table and Table 5 is that singles here are households with strictly one person in them. Multi-person households include singles with children. The table reports pooled (over years 1, 2, and 3) coefficients from the first stage (2), reduced form (3), and two-stage least squares regressions for singles in Panel A and couples in Panel B. In Panel A, column (1) uses firm effects as our preferred measure of firm pay, measured in log wages from an AKM regression shown in (1). Column (2) uses co-worker average wages, while column (3) uses firm effects from a split-sample AKM. Panel B presents the analysis for couples: column (4) reports total household results, while columns (5) and (6) report results split by own consumption change and spousal consumption change, respectively. Standard errors are clustered at the firm level.

Table A2: Changes in unemployment insurance risk

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
		Panel A			*Panel B*		
$\psi_k - \psi_j$	0.032		0.032	0.031		0.031	0.030
	(0.005)		(0.005)	(0.005)		(0.005)	(0.005)
$\Delta_k - \Delta_j$		-0.012	-0.012	-0.007			
		(0.005)	(0.005)	(0.005)			
$\Delta_k^{fe} - \Delta_j^{fe}$					-0.016	-0.015	-0.010
					(0.005)	(0.006)	(0.005)
Controls	N	N	N	Y	N	N	Y
Obs.	77605	79778	77605	77605	79674	77504	77504

Notes: The table reports standardized coefficient estimates from variations of equation (4). Panel A uses the average proportion of individuals experiencing unemployment for a firm as a measure of firm-level unemployment risk. Panel B uses estimates of firm effects from estimating two-way fixed effects regressions in the spirit of Lachowska et al. (2025), where the dependent variable is an indicator of whether a worker experiences unemployment the following year or not. In both panels, we compute standardized coefficients, which adjust for the variance in the underlying variable, to enable comparison between $\hat{\gamma}^c$ and $\hat{\gamma}^\Delta$. Control variables include calendar year, age, and gender dummies. Standard errors are clustered at the firm level.

Table A3: MPC across consumption categories for singles

	(1)	(2)	(3)	(4)
Category	Normalized MPC	SE	MPC	Proportion
A. Groceries	0.005	(0.082)	0.001	0.121
B. Clothes	0.480	(0.113)	0.014	0.029
C. Furnish	0.727	(0.171)	0.027	0.037
D. Cars	0.486	(0.468)	0.016	0.033
E. Hotel	1.014	(0.274)	0.004	0.004
F. Restaurant	0.730	(0.117)	0.025	0.034
G. Beverage	0.435	(0.133)	0.006	0.013
H. Utilities	0.599	(0.250)	0.057	0.095
I. Transport	0.276	(0.154)	0.012	0.044
J. Communication	-0.051	(0.104)	-0.001	0.028
K. Electro	-0.733	(0.310)	-0.004	0.006
L. Recreation	0.241	(0.125)	0.011	0.047
M. Culture	0.234	(0.161)	0.002	0.010
N. Personal care	0.662	(0.153)	0.006	0.009
O. Insurance	0.266	(0.103)	0.010	0.038
P. Finance	0.737	(0.193)	0.030	0.040
Q. Service	-0.235	(0.257)	-0.014	0.058
R. Other	0.042	(0.164)	0.001	0.020

Notes: The Normalized MPC in (1) is measured as the MPC (3) divided by the proportion of the consumption category in total consumption (4).

Table A4: Summary Statistics Across Samples

	Single-Person Households		**Multi-Person Households**	
	Mean	SD	Mean	SD
Panel A: Full Sample				
Consumption (household)	318.27	249.10	789.75	453.48
Wages (household)	574.50	259.35	1139.37	503.19
Consumption (own)			396.69	295.21
Wages (own)			628.85	285.00
Observations	2,570,293		10,280,952	
Panel B: Trimmed Sample				
Cons. (household)	372.26	251.61	827.99	446.27
Wages (household)	606.77	267.36	1164.13	502.13
Cons. (own)			416.09	296.70
Wages (own)			638.48	288.07
Observations	1,975,943		9,397,129	
Panel C: (5–99) Sample				
Cons. (household)	329.92	238.88	790.70	452.05
Wages (household)	574.50	259.35	1139.37	503.19
Cons. (own)			398.78	292.59
Wages (own)			628.85	285.00
Observations	2,570,293		10,280,952	

Notes: Single-person refers to household with only one member, while multi-person to those with more than one member. Unit of observation is the worker. First two lines report the workers' total household spending and wage income, while the last two report the workers own spending and wage income. Variables denominated in thousands real 2018 NOK.

Table A5: KSS Variance Decomposition: Robustness

	Raw		Trimmed		Winsor. (5-99)	
	W	**C**	**W**	**C**	**W**	**C**
(A) Single-person households						
Var expl.	0.122	0.403	0.108	0.138	0.122	0.219
Var(firm)	0.014	0.010	0.011	0.006	0.014	0.006
	(0.113)	*(0.024)*	*(0.104)*	*(0.041)*	*(0.113)*	*(0.027)*
Var(person)	0.088	0.375	0.083	0.128	0.088	0.206
	(0.723)	*(0.930)*	*(0.769)*	*(0.932)*	*(0.723)*	*(0.939)*
Cov(firm,person)	0.010	0.009	0.007	0.002	0.010	0.004
	(0.163)	*(0.046)*	*(0.127)*	*(0.027)*	*(0.163)*	*(0.034)*
Corr(firm,person)	0.286	0.154	0.225	0.068	0.286	0.107
R squared	0.792	0.778	0.799	0.628	0.792	0.729
(B) All workers: individual earnings and consumption						
Var expl.	0.135	0.407	0.131	0.329	0.135	0.301
Var(firm)	0.012	0.008	0.011	0.007	0.012	0.006
	(0.085)	*(0.019)*	*(0.083)*	*(0.020)*	*(0.085)*	*(0.019)*
Var(person)	0.104	0.390	0.101	0.317	0.104	0.287
	(0.767)	*(0.958)*	*(0.775)*	*(0.965)*	*(0.767)*	*(0.955)*
Cov(firm,person)	0.010	0.005	0.009	0.002	0.010	0.004
	(0.148)	*(0.023)*	*(0.143)*	*(0.015)*	*(0.148)*	*(0.026)*
Corr(firm,person)	0.290	0.084	0.282	0.053	0.290	0.095
R squared	0.817	0.732	0.822	0.693	0.817	0.734
(C) All workers: equivalized earnings and consumption						
Var expl.	0.142	0.253	0.133	0.130	0.142	0.187
Var(firm)	0.007	0.005	0.006	0.003	0.007	0.003
	(0.050)	*(0.019)*	*(0.048)*	*(0.024)*	*(0.050)*	*(0.019)*
Var(person)	0.126	0.244	0.119	0.126	0.126	0.181
	(0.888)	*(0.966)*	*(0.898)*	*(0.973)*	*(0.888)*	*(0.970)*
Cov(firm,person)	0.004	0.002	0.004	0.000	0.004	0.001
	(0.063)	*(0.015)*	*(0.053)*	*(0.002)*	*(0.063)*	*(0.011)*
Corr(firm,person)	0.149	0.057	0.128	0.008	0.149	0.041
R squared	0.768	0.695	0.767	0.608	0.768	0.677

Notes: KSS variance decomposition across different samples and measurements. The baseline KSS results reported in Table 8 corresponds to the Trimmed column in Panel (A).

B Consumption Expenditure Data Appendix

B.1 Electronic Payments: Data Construction

The electronic payment 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 C.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.

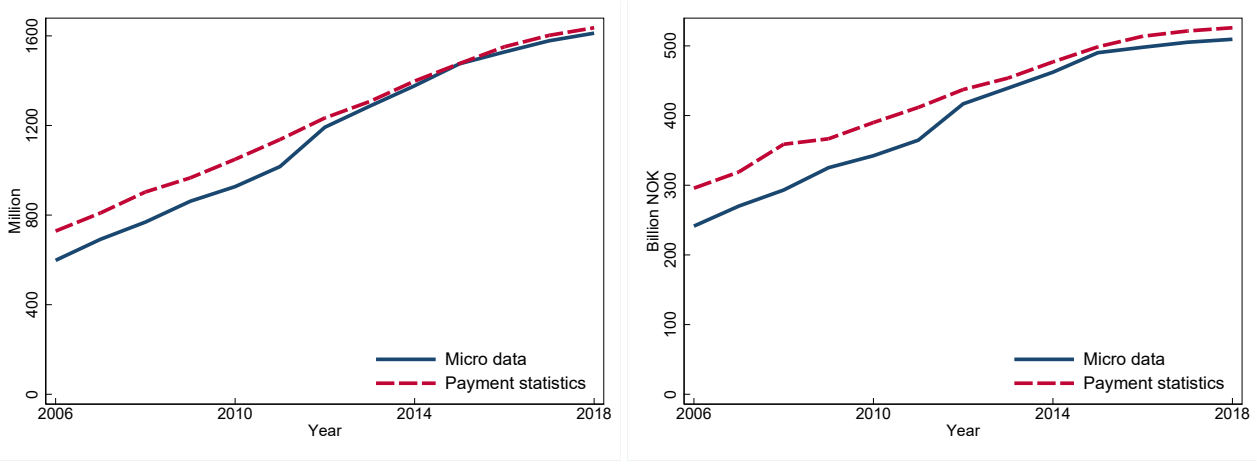
Table C.1: Consumption Categories

	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 C.1 displays the total number and value of transactions in our micro data versus aggregate values reported in official statistics.

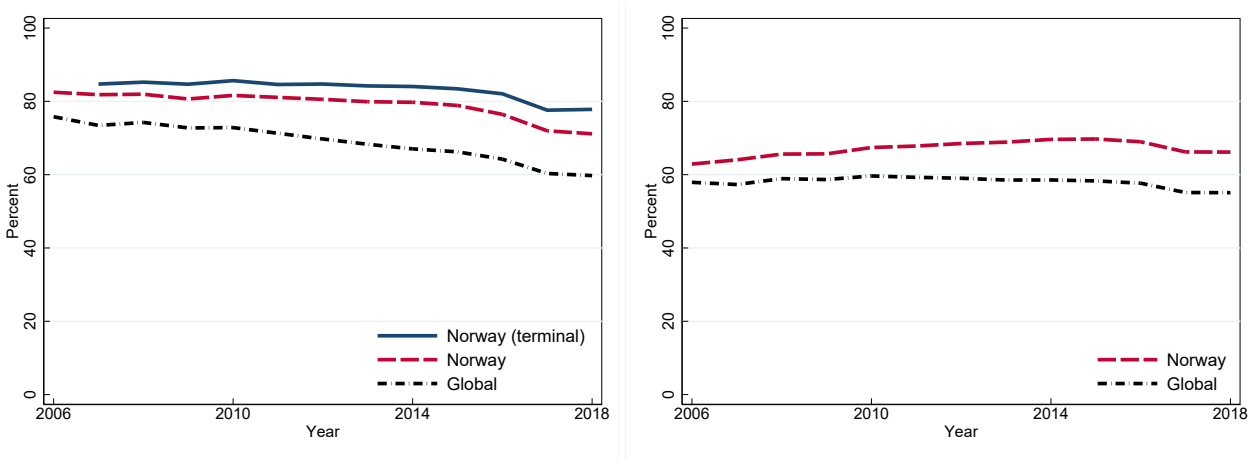
Figure C.1: BankAxept transactions and total amount.



Overall BankAxept is the dominant means of card payment by Norwegians as shown in Figure C.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 C.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.

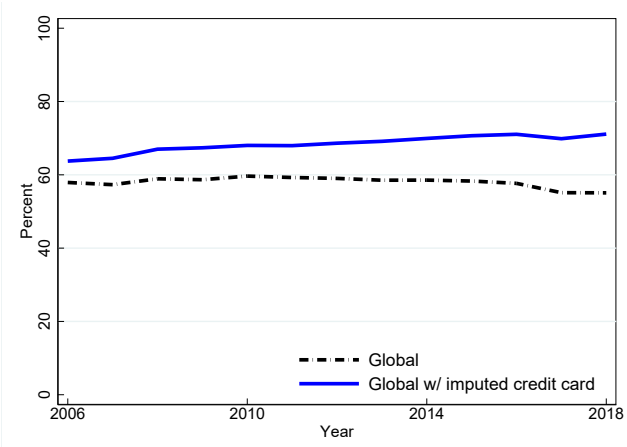
Of the missing payments in Figure C.2, panel (b), credit card accounts for roughly 40 percent the payments. Despite not having the actual credit card purchases in our data, the granularity allows us to impute payments of credit card bills. In our wire transfer measure, payments to banks is a separate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

Figure C.2: BankAxept share of payments.



credit card payments covers on average 63 percent of total credit card payments from the official statistics. In figure C.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 C.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.

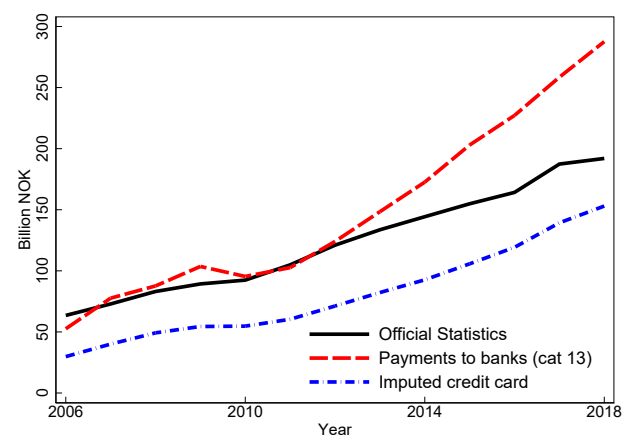
5. Temporary non-mortgage: Re-classify as imputed debt service if more than 75 percent of observations in a window of +/- 2 observations are classified as mortgage.

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 C.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 C.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 Further model and calibration details

C.1 Hamilton-Jacobi-Bellman equations

Let the present discounted value of unemployment and employment be written as $V^u(z,a)$ and $V^e(z,\varepsilon,a)$. z is the productivity of the firm, ε is the worker's skill, and a is a workers asset stock. Define the relative acceptance strategies, $z^u(\varepsilon,a)$ and $z^e(z,\varepsilon,a)$

$$V^e\left(z^u(\varepsilon,a),\varepsilon,a\right) \coloneqq V^u(\varepsilon,a) \text{ and } V^e\left(z^e(z,\varepsilon,a),\varepsilon,a\right) \coloneqq V^e(z,\varepsilon,a), \tag{18}$$

where the thresholds are unique since $V^e_z > 0$. Of course, we may immediately write that $z^e(z,\varepsilon,a) = z$. The associated Hamilton-Jacobi-Bellman equations for unemployment and employment may, respectively, be written as

$$\rho V^u(\varepsilon,a) = \max_{c^u}\left\{u(c^u) + \lambda\int_{z^u(\varepsilon,a)}\left(V^e(x,\varepsilon,a) - V^u(\varepsilon,a)\right)dF(x) + V^u_a(\varepsilon,a)\dot{a}\right\}, \tag{19}$$

and

$$\begin{aligned}\rho V^e(z,\varepsilon,a) = \max_{c^e}\Bigg\{&u(c^e) + s\lambda\int_z\left(V^e(x,\varepsilon,a) - V^e(z,\varepsilon,a)\right)dF(x)\\ &+ s^r\lambda\int_{z^u(z,\varepsilon,a)}\left(V^e(x,\varepsilon,a) - V^e(z,\varepsilon,a)\right)dF(x)\\ &+ \left(\delta + s^r\lambda F(z^u(z,\varepsilon,a))\right)\left[V^u(\varepsilon,a) - V^e(z,\varepsilon,a)\right] + V^e_a(z,\varepsilon,a)\dot{a}\Bigg\},\end{aligned} \tag{20}$$

Unemployed workers enjoy utility from consumption, find jobs at rate λ, and the annuity value changes through asset growth. Employed workers enjoy utility from consumption, find jobs at rate $s\lambda$, experience reallocation shocks at rate $s^r\lambda$, the worker becomes unemployed if the job is destroyed at rate δ or the offer following a reallocation shock is not accepted, and again the annuity value changes through asset growth.

C.2 Wages

We follow an ex-post bargaining setup with continuous renegotiation, where the threat to disruption of production is only credible. This follows that of Binmore et al. (1986), Hall and Milgrom (2008) and Elsby and Gottfries (2022).[5] The continual renegotiation means that

[5] See Grindaker and Simmons (2025) who adopt the same setup we use here.

current wages do not influence future turnover decisions, and the surplus of the match is independent of the current wage.

Workers and firms engage in a two-period bargaining game. In the first stage, workers make a take-it-or-leave-it wage offer. If an agreement is not reached, the worker receives a payoff $\omega - \beta K$, and the firm receives a payoff $-(1-\beta)K$. These represent the threat points in the bargaining game. Under disagreement, in the second stage workers make an offer with probability β, while firms make an offer with probability $1-\beta$. If the worker makes the offer in the second stage, the firm will accept any offer such that $p - w \geq -(1-\beta)K$, thus the worker demands $w = p + (1-\beta)K$. If the firm makes the offer, the worker will accept if $w \geq \omega - \beta K$, thus the firm offers $w = \omega - \beta K$.

In the first stage, the worker chooses a wage that makes the firm indifferent between accepting now and receiving $p - w$ or proceeding to the second stage. The expected payoff from the stage 2, the lottery stage, is $\beta(p-(p+(1-\beta)K)) + (1-\beta)(p-(\omega-\beta K))$. The equilibrium wage, w^*, is then equal to

$$w^* = \beta p + (1-\beta)\omega. \tag{21}$$

This is the subgame perfect equilibrium wage of the alternating-offer bargaining game that includes temporary disruption costs.

C.3 Computation

We follow Achdou et al. (2022) when solving the model. Let $\mathbf{V}$ be the column vector that stacks the value functions. Let $\mathbf{\Lambda}$ represent the matrix governing workers returns from job mobility. Let $\mathbf{\Delta}$ represent workers returns from job destruction or reallocation shocks. Finally, let $\mathbf{P}$ describe the returns from asset accumulation. We can then write down the transitions matrix $\mathbf{T}$ as

$$\mathbf{T} = \mathbf{\Lambda} + \mathbf{\Delta} + \mathbf{P} \tag{22}$$

Once the matrices are built the solution iterates until convergence using the implicit method. With an appropriate value for Γ, which can in principle be arbitrarily large, we iterate on the following equation

$$\mathbf{V}^{n+1} = \left(I(\frac{1}{\Gamma} + \rho) - \mathbf{T}\right)^{-1}\left(u(\mathbf{c}^n) + \frac{1}{\Gamma}\mathbf{V}^n\right), \tag{23}$$

until some threshold level, ε^{thresh} is reached which we take as occurring when $\varepsilon^{thresh} = 10^{-6} > \max |\mathbf{V}^n - \mathbf{V}^{n+1}|$. We are then able to calculate the equilibrium distribution, $\mathbf{g}$, since

$$0 = \mathbf{T}^T\mathbf{g}, \tag{24}$$

where superscript T is the transpose. Note that we must adjust $\mathbf{T}$ to also incorporate the death process which replaces a worker with identical skill but pushes them into unemployment with zero assets.

We calibrate the parameters using Indirect Inference Gourieroux et al. (1993). We compute the majority of the moments using the ergodic distribution from (24); compute the MPC using simulated data of a sample of 10000 workers observed for 15 years and calculate the MPC on the final 10 years; and compute the variance of firm effects and worker effects using the methodology in the following subsection. Stack these moments into the vector $\mathbf{m}^{\mathbf{S}}(\vartheta)$, where $\mathbf{j}$ is the vector of parameters, and denote the corresponding empirical moments from the data in the vector $\mathbf{m}(\vartheta_0)$, where ϑ_0 is the vector of "true" parameter values. We then estimate ϑ using the following criterion

$$\hat{\vartheta} = \arg\min_{\vartheta} \left\{ \left[\mathbf{m}(\vartheta_0) - \mathbf{m}^{\mathbf{S}}(\vartheta)\right]' \mathbf{W} \left[\mathbf{m}(\vartheta_0) - \mathbf{m}^{\mathbf{S}}(\vartheta)\right] \right\}. \tag{25}$$

We choose the weighting matrix, $\mathbf{W}$, as a diagonal matrix containing the reciprocal of the empirical moments.

C.4 Estimating firm effects in the model

Instead of running corrections for limited mobility bias in model simulated data when calibrating the model parameters, we proceed following Borovicková and Shimer (2024), to which we describe next and largely follow their notation. We focus on describing the AKM in the natural logarithm of consumption. Given a log-consumption equation, the two first-order conditions for minimizing the sum of squared errors (running OLS) are written as

$$\sum_{t \in \mathcal{T}_i} \left(\log c_{it} - \hat{\alpha}_i - \hat{\psi}_{J_{(it)}} \right) = 0, \tag{26}$$

where $\mathcal{T}_i$ are the periods the worker earns a wage (is employed). And the second is written as

$$\sum_{t}^{T} \sum_{t \in \mathcal{I}_{jt}} \left(\log c_{it} - \hat{\alpha}_i - \hat{\psi}_j \right) = 0, \tag{27}$$

where $\mathcal{I}_{jt}$ is the set of employees employer j employs at time t.

We assume that we have access to a very long panel of data, that is $T \to \infty$, and there are a very large number of workers and firms. While workers leave the labor market in our environment (this is important in the model to keep workers from increasing assets indefinitely if we calibrate the interest rate to be larger than the discount rate), they are replaced by an

individual of identical ability. Thus, we observe consecutive sets of individuals of identical ability for a very long time. The workers consumption choice is dependent on their income, but also their asset level. But since we observe individuals forever, they or their successor will with some probability be at each grid point on the asset spectrum with positive probability. Thus, a workers unobserved type, let us call it x_i, must be there ability level. This is also true for the firm unobserved type y_i, but less nuanced for the firm since firms do not "leave the labor market". Let $\bar{\alpha}_x$ be the limiting value of $\hat{\alpha}_i$, and $\bar{\psi}_y$ be the limiting value of $\hat{\psi}_i$. Again, $\bar{\psi}_y$ is the firms firm effect estimated on data we have described ($T \to \infty$, and there are a very large number of workers and firms).

Using the FOC conditions from above for α_i ((26)), and ψ_j ((27)), we can solve for $\bar{\alpha}_x$ and $\bar{\psi}_y$, up to an additive constant, on a computer using

$$\bar{\alpha_x} = \frac{\sum_y (\log(c_{xy}) - \bar{\psi}_y) g_{xy}}{\sum_y g_{xy}} \quad \text{and} \quad \bar{\psi_y} = \frac{\sum_x (\log(c_{xy}) - \bar{\alpha}_x) g_{xy}}{\sum_x g_{xy}}, \tag{28}$$

where g_{xy} is the mass of type x workers in type y firms. Where again these correspond to the fixed abilities for both. Note two things. First, the workers current asset level plays no role in determining type, and so there is a positive variance of the error term due to asset heterogeneity. Second, we reiterate, that the worker and firm effects are only identified up to an additive constant. Of course, this is not a concern when interested in variance components.

This procedure should be thought of as analogous to adjusting the variance components for limited mobility bias following Kline et al. (2020), as we have done in the empirical analysis in the main paper. Instead, we make use of the fact that we have access to the data generating process, and therefore know the ergodic distribution.