



ROCKWOOL Foundation Berlin

Institute for the Economy and the Future of Work (RFBerlin)

DISCUSSION PAPER SERIES

271/26

Firm Pay and Consumption Inequality

Sigurd Galaasen, Andreas Kostøl, Michael Simmons

Firm Pay and Consumption Inequality

Authors

Sigurd Galaasen, Andreas Kostøl, Michael Simmons

Reference

JEL Codes: D12, E21, J31, J62, J64

Keywords: Consumption inequality; firms; marginal propensity to consume; search frictions

Recommended Citation: Sigurd Galaasen, Andreas Kostøl, Michael Simmons (2026): Firm Pay and Consumption Inequality. RFBerlin Discussion Paper No. 271/26

Access

Papers can be downloaded free of charge from the RFBerlin website: <https://www.rfberlin.com/discussion-papers>

Discussion Papers of RFBerlin are indexed on RePEc: <https://ideas.repec.org/s/crm/wpaper.html>

Disclaimer

Opinions and views expressed in this paper are those of the author(s) and not those of RFBerlin. Research disseminated in this discussion paper series may include views on policy, but RFBerlin takes no institutional policy positions. RFBerlin is an independent research institute.

RFBerlin Discussion Papers often represent preliminary or incomplete work and have not been peer-reviewed. Citation and use of research disseminated in this series should take into account the provisional nature of the work. Discussion papers are shared to encourage feedback and foster academic discussion.

All materials were provided by the authors, who are responsible for proper attribution and rights clearance. While every effort has been made to ensure proper attribution and accuracy, should any issues arise regarding authorship, citation, or rights, please contact RFBerlin to request a correction.

These materials may not be used for the development or training of artificial intelligence systems.

Imprint

RFBerlin
ROCKWOOL Foundation Berlin –
Institute for the Economy
and the Future of Work

Gormannstrasse 22, 10119 Berlin
Tel: +49 (0) 151 143 444 67
E-mail: info@rfberlin.com
Web: www.rfberlin.com



Firm Pay and Consumption Inequality*

Sigurd Galaasen[†] Andreas R. Kostøl[‡] Michael Simmons[§]

September 9, 2026

Abstract

An extensive body of research has documented that firms play an important role in wage inequality. However, much less is known about how firm pay differentials translate into household spending, and under what conditions those spending differences map into welfare differences. In this paper, we combine newly collected expenditure data with employer-employee records from Norway to examine how firm pay inequality is transmitted to household spending. We show that when workers move up the firm pay ladder, the resulting increase in spending is roughly half of the increase in disposable income: the implied marginal propensity to consume (MPC out of income) is approximately 0.5. These spending responses are highly persistent. The response is concentrated in discretionary and housing-related spending: hotel stays display the largest relative response, while housing-related expenditures account for the largest share of the aggregate MPC. Next, we decompose wages and spending into worker- and firm-level components, and find that firms account for roughly half the variance in consumption inequality relative to wages. Finally, a search model with consumption-savings decisions reproduces these results and shows that firm-level consumption differences closely track firms' contribution to welfare inequality. Our findings thus indicate that firms play an important role not only in wage inequality but also in welfare inequality.

JEL codes: D12, E21, J31, J62, J64.

*This project received financial support from the Research Council of Norway 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 David Berger, Axel Gottfries, Daniel Haanwinckel, Rasmus Lentz, Luigi Pistaferri, and seminar participants at Aarhus University, ASU, Bocconi Workshop on Labor Markets and Public Policy, Collegio Carlo Alberto, EEA 2025, SOFI Stockholm, SED 2025, UCD, UCLA, UCSB, University of Edinburgh, University of Manchester, University of Oslo, Uppsala University, University of Strathclyde, University of York, IFN Workshop on Firms and Human Capital, and the 2nd Oslo Labor Economics Workshop, for helpful comments and suggestions.

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

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

[§]Uppsala University. E-mail: michael.simmons@nek.uu.se

1 Introduction

An extensive body of research has shown that firms play an increasingly important role in wage inequality (e.g., [Abowd et al., 1999](#), [Kline, 2024](#)). These findings have fueled growing interest in understanding their underlying sources (e.g., [Kline, 2025](#)), with rising power to mark down wages being a leading explanation for the widening dispersion in earnings across firms.¹ However, the welfare implications of these changes remain poorly understood. This lack of evidence reflects in part that the mapping from firm pay differentials to welfare depends on features that are difficult to observe directly, including consumption smoothing, households’ expectations about future income changes ([Friedman, 1957](#)), access to liquidity ([Jappelli, 1990](#)), and the extent of labor market frictions ([Lise, 2013](#)).

This paper addresses these challenges by studying how firm pay differentials are transmitted to household spending. What makes our contribution possible is the combination of newly collected spending data with administrative records on employment transitions and annual gross earnings in Norway. Our starting point is that consumption brings us one step closer to welfare than wages alone. Because consumption is forward-looking, it reflects several forces that drive a wedge between pay and welfare: consumption smoothing, expectations about future income, and access to liquidity. The marginal propensity to consume (MPC) out of firm pay differentials therefore carries information about their welfare consequences that wages alone cannot provide.

We estimate spending responses to firm pay differentials with a movers design. To this end, we estimate firm pay premia following [Abowd et al. \(1999\)](#), and relate changes in household spending to differences in permanent pay premia. Our design identifies the MPC out of firm pay premia using a split-sample procedure, which mitigates the small-sample bias arising when movers’ spending changes are correlated with estimated firm effects. Scaling the 3-year spending response by the corresponding response of disposable income then yields the MPC out of income induced by firm pay premia. This scaling serves two purposes: it absorbs the attenuation from measurement error in estimated firm premia, and it expresses the spending response per dollar of net income rather than gross labor income. The object we estimate is thus the spending response to changes in current and future income from moving along the firm pay ladder. A move to a higher-paying firm raises current and expected future income relative to a move to a lower-paying firm.

¹For 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](#)).

Our analysis yields three main findings on the transmission of firm pay to spending. The first is that, after adjusting for measurement error in the estimated premia, a 10% differential in firm pay raises disposable income by 5.3% of its pre-move level over the three years after the move. The corresponding increase in household spending is 2.8%, also measured relative to pre-move disposable income: a \$1 increase in disposable income induced by an upward move in the pay ladder leads to approximately \$0.5 in additional spending. Second, the spending and income responses are highly persistent and remain economically and statistically significant even seven years after the move. Third, we find important heterogeneity across spending categories. Housing-related expenditures, including utilities and refurbishment expenses, account for the largest share of the aggregate spending response. By contrast, hotel stays are the most sensitive category: they exhibit the largest implied elasticity with respect to income. We find no statistically significant effect on more essential spending, such as groceries and communication.

To interpret our results as causal responses to firm pay differentials, the move must be unrelated to unobserved shocks to earnings capacity or spending needs, and to the quality of the worker-firm match. Earnings trends before the move are similar for workers moving up and down the pay ladder, suggesting that downward mobility is not driven by unobserved negative shocks to earnings capacity. Consistent with prior studies, we also find that earnings changes are symmetric across upward and downward movements, as expected when moves are unrelated to match quality. Our MPC out-of-income interpretation additionally requires that firm pay differentials do not simultaneously shift non-pay attributes of the job—such as amenities, household composition, housing, or local consumption opportunities—that affect household spending through channels other than income. We show that the estimates hold after controlling for characteristics designed to capture non-income channels through which firms can affect consumption.

We next investigate how firms contribute to consumption inequality and under what conditions those differences are informative about welfare. To this end, we formalize a simple consumption-equivalent variation (CEV) framework that motivates a variance decomposition in the spirit of [Abowd et al. \(1999\)](#) and [Kline et al. \(2020\)](#). We build on the extensive literature showing that the bias-corrected variance of firm fixed effects explains a substantial share of overall wage inequality. While a log-linear specification of consumption expenditure is theoretically restrictive, it provides a useful benchmark for organizing the empirical analysis. Empirically, we find that the variance of firm effects in consumption is about half the size of that for pay, a ratio that is stable across data subsamples and robust to alternative treatments of outliers.

The CEV framework provides a benchmark way to relate the variance of firm consump-

tion components to welfare. This interpretation relies on a static framework in which the person and firm components of consumption are log-additive, a restriction that we relax in Section 6. While utility from consumption may take a general form, an important nuance arises when individuals value non-pay amenities differently, as in the classic monopsony model of [Card et al. \(2018\)](#). In that case, wage dispersion is partly attributable to compensating differentials rather than to welfare loss. Under this benchmark interpretation, the decomposition of firm consumption effects is informative about how firms contribute to welfare inequality via consumption differences. While it is an upper bound if compensating differentials account for part of firm wage inequality, it is a lower bound if high-paying firms also offer better amenities, as several recent studies reviewed in [Kline \(2024\)](#) find.²

The magnitudes of the spending response and the variance of firm spending effects are informative about how workers perceive firm pay premiums. On the one hand, a frictionless permanent-income benchmark predicts an MPC close to one if workers view the firm pay premium as a permanent shift in their income path. Our estimate of a persistent effect below one instead suggests that firm pay premiums are considered risky, and that workers save a significant portion of the income increase. On the other hand, the evidence is also consistent with important search frictions, in which moving to a lower-paying firm entails the expectation of moving to a higher-paying one in the future.

We explore whether an off-the-shelf search model with savings and two-sided heterogeneity can account for these new facts, and use it to interpret the empirical results and clarify their welfare implications within the model. We start by showing that, under a reasonable degree of risk aversion, the model can quantitatively match both the estimated MPC and the variance decomposition.³ Through the lens of the model, the large responses to firm pay differentials are explained by the fact that workers do not know when job-switching opportunities will arise. As a consequence, workers experience a substantial revision to the expected net present value of their careers when they move to a new firm. Within the model, these changes generate discrete jumps in consumption and welfare at job changes and limit workers' ability to smooth consumption while they remain on lower rungs of the job ladder.

Lastly, we use the quantitative model to answer two questions. First, does the mapping from firm consumption effects to welfare survive when consumption is not log-additive? In the model, the spending response to a move varies with a worker's assets and skill. We compute, for each worker, the consumption-equivalent variation that leaves the worker indifferent between two firms, average within firms, and take the variance across firms. This

²For Norway, [Audoly et al. \(2024\)](#) document a negative correlation between wages and amenities, which points to the upper-bound case.

³The coefficient of relative risk aversion is tightly linked to the MPC and the variance of firm effects on consumption; we calibrate a value of 3.2.

variance is half the variance of firm wage effects, close to the ratio observed for consumption in the data. Second, how large is the revision to lifetime welfare when climbing the pay ladder? We find that a one-standard-deviation move up the pay ladder is worth a permanent 2.3% increase in consumption.

Our paper contributes to the long and influential literature on the role of firms in wage inequality, building on the seminal framework in [Abowd et al. \(1999\)](#). A robust finding, refined with modern econometric techniques, is that the variance of firm pay premiums is quantitatively important (see, e.g., [Kline, 2024](#) for an extensive review). Moreover, firm pay effects have been shown to be strongly related to observable measures of firm quality, including value added and firm size (e.g., [Bloom et al., 2018](#) and [Card et al., 2016](#)), as well as to other desirable aspects of the job (e.g., [Sorkin, 2018](#) and [Bagger and Lentz, 2019](#)). We contribute to this literature in two ways. First, we show that firm pay premiums are reflected not only in wages, but also in household spending, and we provide quasi-experimental evidence of this link using job movers. Second, we show that the variance of firm effects in consumption is substantially smaller than in wages. We interpret this attenuation as evidence that firm pay differences transmit only imperfectly to household spending, and then study under what conditions that attenuation is informative about welfare.

Our study also extends research on measuring consumption inequality. [Deaton \(1992\)](#), [Krueger and Perri \(2006\)](#) 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.⁴ [Krueger and Perri \(2006\)](#) document that while income inequality widened during the eighties and nineties, consumption inequality widened much less. We contribute to research on consumption inequality 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 that the consumption response to firm pay differentials is consistent with the notion that households perceive these premiums as persistent.

⁴More recent studies have used de-identified bank account data to study consumption responses to income and unemployment shocks. [Baker \(2018\)](#) and [Ganong and Noel \(2019\)](#) provide evidence from the US, [Gerard and Naritomi \(2021\)](#) for Brazil, and [Andersen et al. \(2023\)](#) for evidence from Denmark. Relative to existing research, our measure covers the entire population and allows us to link household members. Overall, we show our measure of consumption expenditure is consistent with theories of life-cycle and consumption-smoothing behavior, where average consumption expenditure is lower than disposable income and that the standard deviation of consumption is lower than that of gross wages.

Finally, we offer new empirical content to 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 [Borovičková 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 the expected future path of income changes, since workers randomly and infrequently receive opportunities to climb 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 relative to wages along the job ladder. Using an extended version of [Lise \(2013\)](#), we show that the model quantitatively reproduces the estimated MPC along the job ladder and the contribution of firms to 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 in Section 4, we implement it. In Section 5, we describe the revealed preference approach and estimate the variance contributions for wages and consumption. In Section 6, we consider the firm’s role in shaping welfare inequality through the lens of the model, and Section 7 concludes.

2 Data

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

2.1 Linked Expenditure and Employment Data

The main data source is the Norwegian retail clearing institution, Nets Branch Norway (henceforth Nets), collected and maintained by Norges Bank. The data set contains information on all debit card payments and all online bank-wire payments cleared via the

Norwegian Interbank Clearing System (NICS).⁵ Norges Bank received individual-level data aggregated at the weekly frequency, with information on total expenditure, number of transactions, and cashback at the zip-code level and for 24 consumption categories. Moreover, the electronic payments are classified by purpose using the United Nations' 1999 COICOP system.⁶ Using the industry-code information, transactions to and from banks and the public sector are placed into two separate categories. We further refine the bank-payments category by removing payments likely related to mortgage servicing or other investment-related debt, and show that the remaining bank payments align well with aggregate statistics on credit-card payments. Appendix B provides additional details on the data-cleaning procedures.

What makes our analysis unique is that we can link the expenditure data to administrative records maintained by Statistics Norway using unique anonymized personal identifiers. The first key linkage is to the universe of employment relationships from 2006 to 2018.⁷ 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.

⁵BankAxept 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 made through VISA or Mastercard. NICS is the interbank clearing system for the Norwegian Krone (NOK). It is used by all banks operating in Norway that take part in the Norwegian banking community's payments infrastructure. Transactions via NICS include 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 information in this paper.

⁶Card 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, and the level at which our data are aggregated, remained mostly unchanged. The main change was the introduction of more granular categories at a fourth level. Our categorization contains 24 COICOP groups and includes all 12 top-level codes, covering categories such as "Food and beverages," "Restaurants," "Hotels," and "Clothing and footwear." Some of the top-level codes are further divided into second-level COICOP groups, resulting in a total of 24 categories. Table B.1 in Appendix B.1 lists all our categories and provides details on how transactions are linked to a particular COICOP code.

⁷Before 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.

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 of its pooled distribution across workers, which is used throughout the analysis. To illustrate what these variables measure in 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.⁸ Appendix Figure A1 shows that the relative dispersion of consumption and wages is 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.

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.


⁸We caution against reading this comparison as evidence of consumption smoothing: because consumption is below wages on average, a proportional spending rule with no smoothing at all would deliver a lower standard deviation of consumption in levels, and in logs the dispersion of consumption in fact exceeds that of wages, partly reflecting measurement error in spending. We return to measurement error in the estimated firm premia when adjusting the MPC estimates for noise in Section 4—measurement error in spending, if classical, leaves those estimates unbiased and only widens confidence intervals—and to measurement error in spending itself when trimming the sample for the variance decompositions in Section 5; the evidence on consumption smoothing comes from the marginal propensities to consume estimated in Section 4, not from these summary statistics.

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

	A. Single-Person Household		B. Multi-Person Household	
	Mean	SD	Mean	SD
Consumption. (household)	318.41	249.90	790.10	454.16
Wages (household)	573.96	259.42	1138.61	503.10
Consumption. (own)			397.07	295.91
Wages (own)			628.46	284.99
Observations	2,586,778		10,347,506	

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-year. 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).

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. The left panel shows single-person households and the right panel shows multi-person 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 A4 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

estimate firm pay differentials and clarify the restrictions needed for a causal interpretation.

3.1 Estimating Firm Pay Differentials

We apply a two-way fixed-effects model to the log of annual gross earnings to estimate firm pay differentials. Following [Abowd et al. \(1999\)](#), we assume that log annual earnings are additively separable in worker and firm components, which allows us to separately identify firm pay differentials. We implement this idea by specifying a standard Abowd, Kramarz, and Margolis (AKM) model, in which $y_{i,t}$ denotes log annual gross earnings for individual i in year t ,

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

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.

To estimate the pay premiums, we use annual gross earnings from the tax data. We estimate the model for full-time workers aged 20-66 over the 2006-2018 period.⁹ In addition to person and firm dummies, we include age and calendar-year indicators. We estimate ψ on the largest connected set using three variants: the full-sample AKM measure, a half-sample split measure, and a leave-1%-out measure. Our baseline MPC analysis uses the leave-1%-out split-sample measure, while the half-sample and full-sample measures are used for diagnosing the role of measurement error in the firm pay premiums. For the variance decompositions, we likewise implement the leave-one-out connected-set approach. For our empirical analysis, we first restrict the AKM estimation sample to the largest connected set of firms and then normalize firm effects to have a mean of zero in that estimation sample.¹⁰

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 similar to the split sample estimate, which is true for both samples. Stayers tend to work for firms with slightly higher average pay premiums than workers who move between employers (i.e., movers).

⁹While the premia are estimated on this broader age window, all subsequent analyses restrict attention to full-time workers aged 25-60.

¹⁰As emphasized by [Kline \(2024\)](#), firm-effect levels are identified only up to a constant within a connected set. Our implementation, therefore, keeps the largest connected set of firms in the mobility network and centers the estimated firm effects in that sample before using them in downstream analysis.

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

Sample:	A. Movers		B. Stayers	
	Mean	St. Dev	Mean	St. Dev
Age	38.844	9.304	42.394	9.774
Male	0.644		0.595	
Single	0.306		0.281	
Number of Children	1.171	1.118	1.192	1.117
Annual Earnings (NOK1,000)	558.070	279.419	566.014	272.729
Overall ψ	0.005	0.167	0.008	0.156
Split-sample $\tilde{\psi}$	0.005	0.176	0.009	0.160
Number of Observations	116,019		1,122,103	

Notes: The table reports summary statistics for the sample of full-time employed workers aged 25–60 in 2013. The split-sample $\tilde{\psi}$ row reports the leave-1%-out split-sample estimate of the firm pay premium; the overall ψ row reports the full-sample estimate.

Next, we document a statistically significant association between the firm pay differentials, disposable income, and consumption differentials among stayers. To do so, we regress net income and household spending on the estimated firm-pay premium $\tilde{\psi}_{j(i)}$. In Table 3, we report the association between firm pay premia and disposable income in Panel A and the association between firm pay premia and consumption in Panel B. We see that workers in a firm with $\tilde{\psi} = 1$ have NOK 401,000 (approximately USD 40,000, converting at 10 NOK per USD) higher disposable income, and NOK 258,000 higher consumption, relative to workers in a firm with $\tilde{\psi} = 0$. The ratio of these two comparisons implies an MPC of 0.64. This implied marginal propensity to consume reflects not only the effect of higher pay, but also systematic differences between workers who sort into different firms. We add indicators for each age group as controls in column (2), a dummy for gender in column (3), and an indicator for each number of children below age 18 in column (4). Somewhat surprisingly, we find that adding controls does not materially change the estimated differences in disposable income and consumption.

3.2 Identifying Assumptions for Movers Design

We maintain the assumptions under which the estimated firm fixed effect can be interpreted as a firm pay premium, as discussed above and in [Card et al. \(2013\)](#) and [Kline \(2024\)](#). In particular, the causal interpretation of mover-induced changes in firm pay premia requires four restrictions. First, mobility across firms must be unrelated to unobserved shocks to workers' earnings capacity. Second, current earnings must depend on the pay policy of the current employer rather than on the identity of past or future employers. Third, the relevant firm-specific pay component must be sufficiently stable over the sample period.

Table 3: Income, Consumption, and Firm Pay Premia

	(1)	(2)	(3)	(4)
A. Dep. variable: Net Income (NOK1,000)	401.783 (36.166)	402.672 (35.120)	395.922 (35.061)	408.797 (34.674)
B. Dep. variable: Consumption (NOK1,000)	258.313 (20.969)	259.872 (19.864)	264.410 (19.660)	277.389 (19.255)
Age control		Yes	Yes	Yes
Female control			Yes	Yes
N Children control				Yes
Observations	156145	156145	156145	156140

Notes: The table reports coefficient estimates from regressions of (A) total net income and (B) total expenditure on the leave-1%-out split-sample firm fixed effect. The sample is full-time employed singles with and without children aged 25–60 in 2013. Robust standard errors (in parentheses) are clustered at the firm level.

Lastly, mobility must be unrelated to the match-specific component of pay. [Kline \(2024\)](#) shows that the first three restrictions identify the average wage effect for workers who move between a given pair of firms. The fourth is what allows us to compare firms that never exchange workers, and to treat the estimated premia as a common pay ladder along which workers move.


Figure 2 provides support for the first and the last of these restrictions.¹¹ There is no evidence of systematically different pre-move earnings trajectories for workers moving to higher- versus lower-premium firms. As in [Card et al. \(2013\)](#), wage changes associated with upward and downward mobility across the firm-pay distribution are also close to symmetric. Symmetry is what we expect when moves are unrelated to match quality. If workers moved to firms where they had drawn a good match, losses from downward moves would be smaller than gains from upward moves.

These assumptions identify the effect of moving to a higher- or lower-premium firm on individual pay. A causal interpretation of spending responses to pay differentials additionally requires that the move itself be unrelated to unobserved shocks to households' spending needs. Reverse causality, for example, in which a worker moves to a higher-paying firm because spending needs have risen, would invalidate this interpretation. The checks we present in Section 4, which restrict attention to movers who remain in the same commuting zone, the same house, or the same household composition, and show no differential pre-

¹¹[Di Addario et al. \(2023\)](#) provide supportive evidence for independence between current pay and past employers, and [Lachowska et al. \(2023\)](#) find evidence of stable firm effects over time.

move spending trends, address this threat for both the reduced-form and scaled estimates. Lastly, interpreting the *ratio* of the spending and income responses as an MPC out of income requires one additional assumption, which we state and probe in Section 4: the change in pay premia must not simultaneously shift non-pay attributes of the job that affect spending through channels other than current and future income.

Figure 2: Symmetry of Wage Effects of Up- and Downwards Mobility



Notes: The figure reports mean log wages by quartiles of the leave-1%-out split-sample estimate of ψ for the estimation sample of movers in 2011-2015 (excluding movers with a change in full-time employment and multiple movers in a household in the same year), aged 25-60, with at least 36 months of tenure.

3.3 Scaling Spending Responses by Income Responses

Our design estimates two reduced-form responses to the *change in estimated firm-fixed effects*: the response of disposable income and the response of household spending. Scaling the spending response by the income response (via indirect least squares) is numerically equivalent to using the change in firm effects as an instrument for disposable income in two-stage least squares (2SLS). This strategy identifies the MPC out of income induced by firm pay premia. However, the movers-design creates a mechanical correlation between individual (i.e., mover) pay and the firm component (i.e., the “firm premium”).

We address this mechanical bias by using the leave-1%-out split-sample procedure that assigns each worker firm premiums estimated from all workers except those in the worker’s own 1% subsample.¹² In the first estimating equation, we regress the change in net after-tax income induced on the difference in estimated pay policies between the destination k and the origin firm j ,

¹²Relative to the 50% split, this measure minimizes the measurement error in the instrumental variable.

$$\underbrace{y_{i,t} - y_{i,s-1}}_{\text{Change in net income}} = \alpha_t^y + \underbrace{\gamma_t^y [\tilde{\psi}_{k(i),s} - \tilde{\psi}_{j(i),s-1}]}_{\text{Change in Firm Fixed-Effect}} + W'_{i,t} \vartheta^y + \eta_{i,t} \quad (2)$$

where a worker moves to the new firm k in event year s . Following [Yagan \(2019\)](#), we difference out permanent characteristics of households using net income in the year before the move, $y_{i,s-1}$, so that outcomes are determined by the average change in income α_t^y , the change in firm fixed effects (i.e., the firm pay premium), the baseline controls $W_{i,t}$ (defined in the next paragraph, and distinct from the wage-equation controls $X_{i,t}$ in (1)) 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 may induce a change in household spending from year s to t ,

$$\underbrace{c_{i,t} - c_{i,s-1}}_{\text{Change in Spending}} = \alpha_t^c + \underbrace{\gamma_t^c [\tilde{\psi}_{k(i),s} - \tilde{\psi}_{j(i),s-1}]}_{\text{Change in Firm Fixed-Effect}} + W'_{i,t} \vartheta^c + \varepsilon_{i,t} \quad (3)$$

where γ_t^c captures the spending response to the pay differential between firms k and j . Again, we difference out fixed characteristics of the household by subtracting the level of spending the year before the move. In the baseline regressions, $W_{i,t}$ includes a parsimonious set of controls of the worker's side of the match: indicators for age, which absorb life-cycle profiles in income and spending; an indicator for gender, which absorbs level differences in earnings and spending between men and women; indicators for the number of children below age 18, which absorb differences in family spending needs; calendar-year indicators, which absorb aggregate shocks and common trends; and origin-region indicators, which absorb differences in local prices and labor-market conditions across the regions from which workers move. We deliberately exclude characteristics of the destination—such as the destination region or attributes of the destination firm—from the baseline: these are chosen jointly with the new employer and are therefore themselves potential outcomes of the move. Appendix Table [A5](#) shows that the estimates are insensitive to the choice of control set, including specifications that add destination-side firm characteristics, and Appendix Table [A4](#) shows that the estimates are unchanged in subsamples where specific post-move outcomes (location, residence, household size) are held fixed.

Finally, the estimates of γ_t^y and γ_t^c will, in general, be attenuated by noise from having estimated $\tilde{\psi}_{k(i)} - \tilde{\psi}_{j(i)}$. Under classical measurement error, the attenuation bias depends on the ratio of the variance of the true difference to the variance of the noise. While the ratio of these two estimates cancels out noise from the numerator and denominator, yielding a consistent estimate of the MPC out of income induced by firm pay premia, we report the

reduced-form parameter estimates, the implied MPC, and an estimate of the signal-to-noise ratio.

3.4 Key Variables and Sample Construction

Our measure of disposable income is annual and includes labor and capital income, as well as government transfers and taxes paid. The measure of consumption expenditure is the sum of spending across all categories for all household members in a year, but excludes individual transactions exceeding NOK 500K. Next, we calculate the difference of income and consumption between time t and $s - 1$, and normalize it by disposable income at $s - 1$.¹³ Before constructing these normalized outcomes, we winsorize the underlying income and total expenditure at the top and bottom 1% of their distributions.

We provide summary statistics for the estimation sample of movers in Table 4. As expected, singles are younger, have fewer children, and are more likely to work and live in Oslo, Norway's capital, than married and cohabiting individuals. The level of pay and disposable income is substantially lower among singles, who also exhibit a lower consumption-to-net-income ratio. The spending share on food and clothing is lower than that of couples, possibly reflecting fewer children, whereas spending shares in the selected set of other categories are broadly similar.

The estimation sample consists of full-time workers aged 25-60 who change employers between 2011 and 2015. These restrictions ensure that we observe workers and households in the data for at least three years prior to and three years after a move; within this window, each regression uses the observations with non-missing consumption and income, i.e., $t - s \in [-3, 3]$. We furthermore require workers to have least 36 months of tenure at their origin 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 full-time work indicator. We also exclude a small number of individuals whose spouses move to a new firm in the year s . These additional restrictions alleviate concerns that unobserved factors drive both changes in firm pay premiums and changes in household consumption.

Throughout the empirical analysis, we estimate the average MPC out of income separately for single individuals and those in cohabiting relationships. This sample split is performed by defining singles as those in the year before the move and in the outcome year. Hence, a worker contributes to the singles household spending only in the outcome years in which they are still single, and does not require workers to remain single in every post-move year. The couples sample analogously requires that they be married or cohabiting in

¹³While this normalization helps with precision, the difference yields similar results. These results are available upon request.

Table 4: Key Descriptives for Single and Married/Cohabiting Movers

Sample:	A. Single		B. Cohabiting & Married	
	Mean	St. Dev	Mean	St. Dev
Age	41.313	9.905	43.133	9.273
Male	0.655		0.737	
Lives in Oslo Region	0.379		0.326	
Number of kids	0.410	0.726	1.489	1.077
Household Gross Wealth	2,756.827	2,737.354	5,036.476	3,879.087
Household Debt	1,374.707	1,284.677	2,411.134	1,765.878
Household Annual Earnings	658.276	288.941	1,202.320	491.678
Household Annual Net Income	511.133	213.115	970.489	332.911
Household Annual Expenditure	397.438	298.262	826.950	455.141
Food	47.111	34.108	109.410	53.253
Clothing	12.081	14.086	32.147	22.201
Furnishing	15.088	22.354	32.641	33.256
Motor Vehicles	12.838	48.516	28.338	74.564
Restaurants	13.164	14.674	20.652	17.699
Hotels	1.545	3.443	3.154	5.178
Number of Observations	201,760		559,617	

Notes: The table reports summary statistics for movers in 2009-2015 (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.

the year before the move and in the outcome year. 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. Appendix Table A1 summarizes the impacts of each of these sample restrictions on the total number of individual workers in the estimation sample.

4 The Marginal Propensity to Consume out of Pay Premia

This section presents evidence on the marginal propensity to consume (MPC) out of firm pay differentials 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 advance and earn higher wages in the future, the current pay differential may have little effect on consumption. Second, if workers face liquidity constraints and cannot smooth income, they may consume a large share of a pay increase. More generally, any mechanism that prevents workers from fully

accessing lifetime wealth implies that firm pay policies can affect household consumption.

4.1 Reduced Form Evidence

We begin by examining the impact of leave-1%-out firm pay premia on normalized net income among singles. We do so year by year, from three years before the job change (from firm j to k) to three years after, and plot the coefficient estimates from equation 2 by the solid line in Figure 3. The vertical lines around each point estimate represent 95 percent confidence intervals, with standard errors clustered at the firm level. The horizontal axis denotes event time, $t - s$, and the vertical axis shows the change in income and consumption expenditure, both normalized by disposable income in the year before the move (i.e., event time $t - s = -1$).¹⁴ We find an immediate and highly statistically significant positive effect of firm pay differentials. On average, the move occurs in the summer of year 0, so the full effect of the pay gain is not realized until year 1. The effect then persists over the following two years. Turning to the consumption expenditure to firm pay differentials, we plot the coefficient estimates from equation 3 with the dashed line in Figure 3. Again, the full spending effect of pay changes is abrupt and is fully realized by the first year after. Notably, for both outcomes, we find no statistically significant changes before the move, further supporting the exogenous mobility assumption.


We summarize the reduced-form coefficient estimates over the three-year horizon in Table 5: the estimated effect of the change in firm pay premia on disposable income in Panel A (i.e., the first stage) and on household spending in Panel B (i.e., the reduced form – the MPC out of pay differentials). Within each panel, we report the estimates for two measures of the premium change—the half-sample split, whose two independent replicates allow a direct estimate of measurement noise, and the leave-1%-out split, our baseline—estimated on the common sample on which both measures are available.¹⁵ While the first stage is strong by conventional standards, neither coefficient should be read at face value. The income coefficient is well below one, and the spending coefficient is well below the income coefficient, for three distinct reasons that we take up in turn: measurement error in the estimated firm premia, taxes and transfers, and household saving.

First, all of the coefficients in Table 5 are attenuated by noise in the estimated firm premia, with the attenuation factor given by each measure’s signal-to-noise ratio. For the half-

¹⁴We normalize by disposable income at event time -1 because this is closely related to the log specification of wages used to estimate the firm effects. Indeed, to a first-order approximation, if the left-hand side of the first stage were gross labor income, the specification would be exact under AKM.

¹⁵We also perform the same regressions on a sample of single-person households, which removes singles with children from the sample. These estimates are reported in Panel A of Appendix Table A2, and are virtually identical.

Figure 3: Net-income and consumption responses to changes in firm pay: Singles



Notes: The figure reports coefficient estimates from the first-stage (2) and reduced-form (3) equations for the baseline using the leave-1%-out split-sample firm effects for singles. The sample consists of full-time workers aged 25-60 who moved between firms between 2011 and 2015. 95% confidence intervals are based on robust standard errors, clustered at the firm level.

sample split, this ratio can be estimated directly. The half-sample columns assign each mover the premium change estimated from the random half of the workforce that excludes them, so the estimation error in this cross-assigned measure is independent of the error in the measure constructed from the mover's own half. Hence, regressing the own-half measure on the cross-assigned measure identifies the signal-to-noise ratio. Columns (1) and (3) also report noise-adjusted coefficients—the estimates divided by this ratio or, equivalently, split-sample instrumental-variables estimates that instrument the own-half measure with the cross-assigned one: 0.529 for income and 0.281 for spending. Second, even absent noise, the income coefficient would remain below one because the outcome is disposable income whereas the firm effects are estimated on gross labor income: the noise-adjusted coefficient of 0.529 absorbs the offsetting forces of the tax and transfer system as well as any behavioral responses of the worker, such as adjustments in hours or capital income.¹⁶ Third, the noise-adjusted spending response is 0.281 per unit change in the pay premium. The ratio of the two adjusted coefficients—roughly one half—is the fraction of the induced income change that households spend; the residual is the savings response. Note that this statement uses only the two reduced-form parameters, without reference to any particular transmission

¹⁶Transfers work in the same direction as progressive taxation: for a worker receiving a transfer amount b , a doubling of gross pay from 1 to 2 raises $\log(b + \text{pay})$ by less than $\log 2$.

channel (cf. Autor et al., 2019).¹⁷

Table 5: Reduced-form effects of firm pay premia on income and spending

	Panel A: Net income (First stage)		Panel B: Spending (Reduced form)	
	Half-sample (1)	Leave-1%-out (2)	Half-sample (3)	Leave-1%-out (4)
Coefficient	0.177 (0.013)	0.273 (0.014)	0.094 (0.017)	0.135 (0.020)
Signal-to-noise ratio	0.335	–	0.335	–
Noise-adjusted coefficient	0.529	–	0.281	–
Implied MPC out of income			0.530	0.493
N	51781	51781	51781	51781

Notes: The table reports pooled (over years 1, 2, and 3) coefficients from the first stage (2) (Panel A) and reduced form (3) (Panel B) for singles, using two measures of the change in firm pay premia from an AKM regression as in (1): the half-sample split in columns (1) and (3), in which the workforce is divided into two random halves and each mover is assigned the firm effects estimated on the half that excludes them, and the leave-1%-out split in columns (2) and (4), our baseline, which removes the mover’s own wage from the construction of the pay measure. The signal-to-noise ratio of the half-sample measure is estimated by regressing the mover’s own-half measure of the firm-pay change on the cross-assigned measure; because the estimation errors in the two halves are independent, this regression identifies the share of signal in the variance of the cross-assigned measure. The noise-adjusted coefficients divide the coefficients in columns (1) and (3) by this ratio; equivalently, they are split-sample instrumental-variables estimates in which the own-half measure is instrumented with the cross-assigned measure. No analogous adjustment exists for the leave-1%-out measure, which has no independent replicate. The “Implied MPC out of income” row, reported under Panel B, gives the ratio of the spending response to the income response within each pay measure (column (3) relative to column (1), and column (4) relative to column (2)); because the same signal-to-noise ratio divides both half-sample coefficients, it cancels in the ratio, which is therefore identical whether computed from the estimated or the noise-adjusted coefficients. All specifications include the baseline controls $W_{i,t}$ (indicators for age, gender, number of children, calendar year, and origin region) and are estimated on the common sample on which all pay measures are available; N is reported separately for each column, and the common-sample construction implies the four counts coincide. The sample consists of single full-time workers aged 25-60 who moved between firms between 2011 and 2015. Robust standard errors in parentheses are clustered at the firm level.

To attach a standard error to the ratio, we estimate it by two-stage least squares, using the change in split-sample firm effects as an instrument for disposable income. Table 6 reports the estimates. The MPC out of income induced by firm pay premia for singles is 0.487 (standard error 0.071): households spend roughly half of the induced change in disposable

¹⁷Note that the implied MPC out of income from the half-sample split is 0.530, while the leave-1%-out columns imply 0.493. The small difference across the two measures reflects sampling variation in the estimated premia rather than a failure of the adjustment. Likewise, 0.493 differs slightly from the two-stage least squares estimate of 0.487 in Table 6 below only because Table 5 is estimated on the common sample on which both pay measures are available, whereas Table 6 uses the full baseline sample.

income and save the rest. The estimate is precise enough to reject both complete smoothing (a value of zero) and hand-to-mouth behavior (a value of one).¹⁸ While our preferred specification uses a leave-1%-out estimate of firm effects to eliminate the mover’s own wage from the construction of the baseline firm-pay measure, we can also use differences in the average wage of coworkers (i.e., stayers) as the exogenous regressor (and instrument). Column (2) of Table 6 shows that the estimate using the co-worker wage as an instrument is statistically indistinguishable from the estimate in column (1).

Thus far, our focus has been on single workers’ consumption responses to changes in disposable income following a job change. We next ask to what extent the pay premium is shared with other adults in the household, and whether the response among couples differs from that of singles. To do so, we redefine the mover sample to comprise married or cohabiting couples and aggregate consumption at the household level. The object of interest is the total effect on household income and spending of one partner moving up or down the firm-pay ladder. For consistency, we keep the normalization fixed across samples: household and spousal changes are scaled by the mover’s own pre-move disposable income, exactly as for singles. The coefficients therefore answer the same question as for singles—the spending response per krone of the induced change in household disposable income—and the own and spousal responses sum to the household response. We report the point estimates, averaged over years 1–3, in columns (3)–(5) of Table 6; Appendix Figure A5 shows the corresponding year-by-year responses. The household-level estimate is similar in magnitude to that for singles, and the increase in household consumption is driven primarily by the mover’s own spending: the spouse’s response in column (5) is close to zero and statistically insignificant. Together with the similarity of the household-level and singles estimates, this indicates that our focus on singles does not miss economically important responses that would only appear in the sample of married and cohabiting couples.¹⁹

4.2 Understanding the Nature of the Pay Shock

This section examines the nature of the pay “shock”. To do so, we assess two questions. First, to what extent is the firm’s pay premia shock persistent vs transitory? Second, what conditions must hold to interpret the findings as causal spending responses to changes in current and future disposable income?

¹⁸We perform diagnostics of the noise in Appendix Table A3, where we compare a half-sample split, the leave-1%-out split-sample baseline, and firm effects estimated on the whole sample, and show that changing the amount of noise substantially changes the first stage and reduced form but has no material impact on the estimated ratio.

¹⁹Again, we perform the same regressions on a sample of multi-person households, which adds singles with children to the sample. These estimates are reported in Panel B of Appendix Table A2, and are statistically indistinguishable from the baseline sample of married and cohabiting individuals.

Table 6: The MPC out of income induced by firm pay premia: 2SLS estimates

	(1)	(2)	(3)	(4)	(5)
	<i>Panel A: Singles</i>		<i>Panel B: Couples</i>		
	Split sample	Co-worker	Household	Own	Spouse
MPC (2SLS)	0.487 (0.071)	0.522 (0.103)	0.589 (0.084)	0.659 (0.067)	-0.069 (0.071)
N	52853	50849	174900	174900	174900

Notes: The table reports pooled (over years 1, 2, and 3) two-stage least squares estimates of the MPC out of income for singles in Panel A and couples in Panel B. In Panel A, column (1) uses leave-1%-out split-sample firm effects as the instrument for disposable income, and column (2) uses co-worker average wages. Panel B presents the analysis for couples: column (3) reports total household results, while columns (4) and (5) split the household response into the mover’s own consumption change and the spousal consumption change, respectively. All outcomes in Panel B capture the total household response to one partner’s move and are normalized by the mover’s own pre-move disposable income, the same scaling as in Panel A, so the coefficients are comparable across panels and the own and spousal responses in columns (4) and (5) sum to the household response in column (3). In all three columns the endogenous regressor is the change in household disposable income. All specifications include the baseline controls $W_{i,t}$: indicators for age, gender, number of children, calendar year, and origin region. The sample consists of full-time workers aged 25-60 who moved between firms between 2011 and 2015. Robust standard errors in parentheses are clustered at the firm level.

In the previous section, we considered mobility between firms over the years 2011–2015 to avoid contamination from the Great Recession. Here, we extend the mobility sample back to 2009 and examine a longer horizon, up to 7 years after the mobility event; each mover cohort is followed as far as the data permit, so the later event years are estimated on the earlier cohorts. Figure 4 shows the coefficient estimates from the first-stage and reduced-form equations. We find clear persistence in the earnings effect and consumption response. The spending response is stable over time, while the income effect has been declining in the last three years. The implied MPC out of income is below one throughout, consistent with workers perceiving firm pay premiums as durable but not fully certain. The year-by-year point estimates drift upward, from below 0.5 in the first year after the move to 0.7 in the last year of our observation window, but the yearly estimates are imprecise: the confidence intervals of the early and late years overlap widely in Figure 4, and we therefore do not attach a structural interpretation to the drift. Appendix Table A6 makes the horizon comparison in pooled form on a balanced sample of movers followed for seven years: pooling all seven post-move years yields an MPC out of income close to the one-to-three-year pooling estimated on the same sample.


Two features of our design clarify what the estimates capture. First, we identify a “reduced form” MPC, in the sense that we track the response to moving to firm k regardless of whether a worker later moves to other firms. In practice, later moves are rare: in the baseline sample, 87 percent of workers are still in the destination firm three years after the move, and

in the longer-horizon sample just analyzed, 78 percent remain after seven years. Second, anticipation is only a problem for our estimates if workers adjust spending *before* the move. If workers expect the pay gain from an upcoming job switch, they may start spending part of it in advance. This would mute the measured consumption response at time $t \geq s$ and bias the estimated MPC out of income downward. We find no such pre-move spending response in the data. Anticipation of income *after* the move is different. It is part of what we want to measure, as the firm-pay shock captures the change in net income, including the level and risk associated with the job's net present value.

The key threat to identification is that other changes affecting spending, such as household reorganization or housing adjustments, occur simultaneously with the job move. These confounding events raise concerns about reverse causality, potentially leading workers to move to a higher-paying firm because they have higher spending demand, rather than the other way around. In Table A4, we show that the estimated MPC is essentially unchanged when we restrict attention to individuals who switch employers but remain in the same commuting zone, remain in the same house, or have the same number of household members. We assess sensitivity to controlling for other demographic factors by including a rich set of worker-level controls in our estimating equations. Table A5 shows that including individual-level controls barely moves the estimates. We also address the potential confounding effects of market-level shocks by including indicators for calendar year and origin region. Again, the estimated MPC barely moves.

The MPC out of income should be read as the spending response to the change in *current and future* income induced by the move: our baseline pools the three post-move years, and workers may, in addition, consume out of income changes they expect beyond that window. This anticipation is not a threat to our design: consumption today encoding expectations about the income path attached to the new employer is precisely what makes spending informative about welfare. Figure 4 shows that spending remains flat even as the income effect tapers in the final years of the window, and the year-1 point estimates lie within the year-7 confidence intervals. When we instead pool the post-move years, the estimates are similar across the two horizons (Appendix Table A6). Higher-paying firms may also differ in the job-loss risk they expose workers to. In our framework, such risk is part of the object of interest rather than a confounder: the firm-pay shock captures changes in both the level and the risk of the income path associated with the new employer, and realized job losses are reflected in the disposable-income response. It is nevertheless informative to ask how much of the estimated response reflects systematic differences in unemployment risk across firms. We follow Lachowska et al. (2025) in measuring a systematic firm component of unemployment risk. To do so, we regress an indicator of whether an employee of firm j

Figure 4: Net-income and consumption responses to changes in firm pay: Singles, seven-year horizon



Notes: The figure reports coefficient estimates from the first stage (2) and reduced form (3) equations for the baseline using the leave-1%-out split-sample firm effects for singles. The sample consists of full-time workers aged 25-60 who moved between firms between 2009 and 2015, with each mover cohort followed for up to seven post-move years as far as the data permit. 95% confidence intervals are based on robust standard errors, clustered at the firm level.

receives any unemployment benefits in year $t+1$, controlling for permanent worker-level differences in unemployment risk while identifying the firm-level component among workers who experience job loss across multiple employers. In Table A5, we show that the estimated MPC changes only slightly when we control for this component together with average firm employment growth and the size of the origin and destination firms, indicating that the response is not driven by differential exposure to job loss.

A potential concern is that higher-paying firms offer fewer amenities (e.g., Sorkin 2018) or differ in their workforce composition, leading to differential peer consumption behavior that confounds the estimated MPC (e.g., De Giorgi et al., 2020). Unfortunately, we do not have direct measures of amenities in the data, but we can proxy workplace flexibility using the measured variability in hours. We can also control for compositional differences by including other measured differences between the origin and destination firms, such as the share of employees with young children, and the squared difference in spending levels among stayers to control for firm peer effects on changes in household consumption. In Table A5, we include all of these firm-level variables as controls, and find that the estimated MPC barely moves.

4.3 Consumption categories

This section examines the extent to which the spending response is broad-based or concentrated in certain categories. To implement this test, we estimate the following version of the second-stage equation,

$$c_{i,t}^g - c_{i,s-1}^g = \alpha_t^g + \beta_t^g[y_{i,t} - y_{i,s-1}] + W_{i,t}'\vartheta^g + \varepsilon_{i,t}^g \quad (4)$$

where g denotes the spending category, and β_t^g provides a decomposition of the overall MPC effect at time t . We perform 20 such regressions and report the estimates and standard errors in Panel A of Table 7. The third column shows the share of the total effect attributable to each consumption category, and shows that the overall spending response is driven by a few categories. The largest contributions come from the utilities category, which includes utilities and other housing-related expenditures, such as refurbishment expenses.²⁰ We also see that other discretionary categories, including furnishing and recreation, account for a larger share of the overall MPC.²¹ By contrast, we find little role of changes in spending on caregiving or groceries. Moreover, there is no evidence that spending on financial services drives the overall MPC either. Because financial services include contributions to private retirement accounts, the evidence shows that, in the Norwegian context, firm pay differentials are not primarily a substitute for generous retirement plans.²²

To assess the extent to which some spending categories are more sensitive to income changes due to higher firm pay premiums, we calculate the implied spending elasticities. This facilitates comparisons across categories with different spending shares, which may have different welfare implications (see, e.g., Chetty and Szeidl, 2007 and Attanasio and Pistaferri, 2016). For example, changes in consumption in categories such as necessities like clothing and food, or in spending on goods with high adjustment costs, are more directly tied to economic well-being than other categories. By contrast, increases in spending on luxury items and other nonessential goods, such as hotels and restaurants, may be less important for workers' overall well-being. We report the spending share in column (4) and the adjusted MPC (implied spending elasticity) in column (5) of Table 7, Panel B. We define nor-

²⁰Several issues are worth noting. Utilities include spending on gasoline, electricity in primary and secondary homes, as well as rent paid to professional landlords. As we documented in the previous section, the overall MPC effect barely moves if we instead focus on households that switch firms without changing address.

²¹Table 7 excludes uncategorized spending (credit card payments, cash withdrawals, payments with missing category, and payments to public administration). These expenditures account for roughly 30 percent of total spending in our sample. The implied elasticity for this residual spending category is 0.59.

²²Of course, such expenditure would be misclassified as saving, implying that the MPC is biased upward. Nevertheless, the overall importance of this category is small.

Table 7: MPC across consumption categories for singles

	(1)	(2)	(3)	(4)	(5)
	<i>Panel A: Decomposing spending effect</i>			<i>Panel B: Differential spending elasticities</i>	
Category	MPC	St. Err.	Decomp	Spending share	Implied elasticities
A. Groceries	0.012	(0.010)	0.025	0.117	0.101
B. Clothes	0.017	(0.003)	0.037	0.029	0.607
C. Furnish	0.033	(0.007)	0.072	0.037	0.897
D. Cars	0.016	(0.020)	0.034	0.033	0.480
E. Hotel	0.005	(0.001)	0.010	0.004	1.173
F. Restaurant	0.023	(0.004)	0.049	0.031	0.727
G. Beverage	0.007	(0.002)	0.015	0.013	0.533
H. Utilities	0.095	(0.025)	0.204	0.096	0.991
I. Transport	0.018	(0.007)	0.040	0.042	0.443
J. Communication	0.002	(0.003)	0.005	0.029	0.076
K. Electro	-0.001	(0.002)	-0.002	0.006	-0.132
L. Recreation	0.026	(0.008)	0.057	0.047	0.560
M. Culture	0.003	(0.002)	0.006	0.010	0.283
N. Personal care	0.002	(0.001)	0.005	0.008	0.299
O. Personal effects	0.002	(0.001)	0.003	0.004	0.359
P. Insurance	0.005	(0.003)	0.012	0.038	0.144
Q. Finance	0.018	(0.008)	0.039	0.043	0.422
R. Service	-0.022	(0.017)	-0.048	0.057	-0.387
S. Caregiving	0.000	(0.000)	0.001	0.002	0.168
T. Education & health	0.011	(0.004)	0.023	0.019	0.563


Notes: Panel A reports estimates from the estimating equation 4 in Column (1), and the share of the overall MPC each category can account for in Column (3). Panel B reports implied elasticities in (5), measured as the MPC (1) divided by the share of the consumption category in total expenditure (4). The table excludes uncategorized spending (imputed credit card payments, cash withdrawals, payments to the government sector, and payments with missing category codes), which accounts for roughly a third of total spending in our sample; consequently, the category-level MPCs and budget shares do not sum to the total MPC and one, respectively. The spending response of uncategorized consumption is similar to the average of the categorized spending reported here. The sample consists of single full-time workers aged 25-60 who moved between firms between 2011 and 2015. Robust standard errors in parentheses are clustered at the firm level.

malized spending growth across all 20 spending categories as the movers-2SLS coefficient, adjusted for the initial share of each category in total spending. We report these normalized 2SLS estimates, along with 95% confidence intervals, in Figure 5. These adjustments enable straightforward comparisons across goods and are similar to the estimation of Engel curves. Because households also save, the relevant benchmark is the overall MPC out of income rather than one: categories with an implied elasticity above the overall MPC out of

income increase their share of total consumption as income rises, consistent with luxuries, while categories below it see their consumption share decline. To obtain income elasticities comparable to traditional Engel-curve estimates, one would rescale the implied elasticities by the overall MPC out of income.

Overall, we document a broad-based increase in consumption driven by higher incomes. We see that spending effects are statistically and economically significant and different from zero in categories such as clothing and furnishings, and especially in hotels and restaurants. However, the more essential categories, such as food and communication services, are not statistically significantly different from zero. We also observe a small but positive rise in transport spending and an increase in personal care expenditures.²³ Overall, these findings underscore that the consumption response to mobility-induced pay increases is multi-faceted.

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



Notes: The figure reports implied elasticities by spending category measured as the estimated MPC divided by the share of the consumption category in total expenditure. The sample consists of single full-time workers aged 25-60 who moved between firms between 2011 and 2015.


²³It is possible that some increases in spending may reflect adverse job characteristics, such as higher transport costs, reflecting compensating wage differentials for longer commutes, or greater health-related expenses from more demanding work conditions.

4.4 Heterogeneity

So far, we have estimated the transmission of firm pay differentials to consumption by pooling positive and negative changes. We begin by asking whether the effects are symmetric with respect to the direction of the move. To do so, we split the sample into positive- and negative-change groups and examine symmetry in MPCs. Figure 6 reports the estimates from three years before to three years after the move, where we multiply the estimates by negative one for downward moves.

First, we do not see clear pre-trends for either upward or downward moves, suggesting no obvious anticipation effects by move direction. Second, for both upward and downward moves, we see clear changes in consumption, although the downward response is not statistically different from zero. The point estimates for upward and downward moves are of broadly similar magnitude, although the downward-move estimates are less precise.

Figure 6: Up vs down moves



Notes: This figure plots the coefficient estimates separately for up (left) vs down (right) moves. For expositional clarity, the latter figure is adjusted to display negative coefficients, even though the regression coefficient is positive.


We report estimates across additional firm sample splits in Figure 7, with point estimates, standard errors, the level of consumption, and the number of observations reported in Appendix Table A7. Notably, the estimated MPCs differ across different rungs of the job ladder, as measured by firm effect, although we often lack precision. We observe that movements from lower-paying firms are associated with higher MPCs than movements from higher-paying firms.

Turning to heterogeneity by worker characteristics, we split the sample by household asset quartiles. We note that a standard consumption-savings model implies a higher marginal utility of consumption for the less affluent and a greater likelihood of binding credit or liq-

uidity constraints (see, e.g., Jappelli, 1990). We find evidence that low-asset households exhibit higher MPCs than those in the highest-asset quartile. Moreover, we see that even among workers with assets likely to place them beyond borrowing constraints, they exhibit sizable consumption responses. If firm pay policies solely matter through borrowing constraints, we would expect near-zero MPCs among wealthy workers.

Finally, we find some evidence that workers in the highest-earning types exhibit a somewhat smaller consumption response, although we cannot reject equality with the estimates for low-wage types. Similarly, women appear to have somewhat larger responses than men, but the difference is not statistically significant, and we find no evidence of differential effects across age groups. Taken together, the stability of the spending responses across these broad groups provides some support for the generalizability of our results beyond the specific group of movers, who may be overrepresented by younger, more mobile workers.

Figure 7: MPC heterogeneity across different sample splits for singles



Notes: The figure reports 2SLS estimates of the MPC out of income across different sample splits, corresponding to Appendix Table A7. All specifications include the baseline controls of Table 6. Worker effects are taken from the estimated wage equation. Direction refers to movements up or down in firm pay effects.

5 Consumption Inequality

Having established that employees' spending responds strongly to changes in pay premia, this section quantifies the extent to which firms matter for spending inequality relative to wage inequality, and then discusses what additional assumptions are needed for a welfare

interpretation. To do so, we lay out a simple revealed-preference framework that provides a benchmark mapping from the data to welfare, rather than a fully model-free welfare measure.

5.1 Framework

To assess the role of firms in welfare inequality, we develop a revealed-preference benchmark built on workers' observed consumption behavior. The benefit of this approach is that it clarifies the connection between the reduced-form evidence and welfare while making the required assumptions transparent. As our welfare measure, we use the consumption-equivalent variation (CEV) to quantify contemporaneous utility differences across firms: how much consumption must worker i give up or gain in firm j in order to move to firm k while keeping current-period utility constant? Let this CEV measure, γ_{ijk} , solve

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

where u is the worker's utility function. We make two additional assumptions. First, we require that $u(\cdot)$ is increasing. Note that we do not impose any particular functional form, and the utility function can, in principle, differ across individuals. Second, we assume that $\log(c)$ is additively separable into a worker effect and a firm effect. This additive separability implies homogeneous treatment effects, meaning that $\gamma_{ijk} = \gamma_{jk}$ in (5). While restrictive, we treat this as a natural benchmark and evaluate in the next section how deviations from it may affect our conclusions.²⁴ With these assumptions, we may write

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

Because utility is increasing in consumption, equalizing welfare across firms requires equalizing consumption across firms, as shown by the implication above. Given a baseline firm k , simple manipulation yields

$$\text{var}_j(\gamma_{jk}) \approx \text{var}_j(\log(c_{ij})) = \text{var}_j(\psi_j^c), \quad (7)$$

where the approximation comes from $\log(1 + \gamma_{jk}) \approx \gamma_{jk}$. Under a log-additive consumption equation, taking the variance of $\log(c_{ij})$ over j is equivalent to taking the variance of firm consumption effects. Equation (7) therefore provides a benchmark mapping from firm consumption effects to welfare across firms under these assumptions.²⁵ Imposing the standard assumption that wages are also log-additively separable, it is straightforward to write the

²⁴For example, additive separability between individual and firm fixed effects in consumption is violated when workers' asset holdings and skill determine the consumption response to switching firms.

²⁵Importantly, the choice of benchmark firm k is innocuous. It simply adds a constant to $\log(1 + \gamma_{jk})$ and therefore does not alter the variance.

following ratio

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

where γ_{jk}^w is the analogous welfare measure that solves (6) after replacing c with w . The ratio of firm effects from consumption and wages, therefore, tells us how much one would overstate, or understate, static welfare inequality by focusing on wages rather than consumption.

5.2 Variance Decomposition

Next, we perform a variance decomposition using log wages and log 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). Applying this framework to consumption requires that log consumption, like log wages, be additively separable into person and firm components. Appendix Figure A7 probes this assumption directly, replicating for consumption the diagnostics developed for wages by [Card et al. \(2013\)](#). Mean log consumption of movers across firm-premium quartiles displays the same patterns as wages: gains for workers moving up the firm-premium distribution mirror the losses of workers moving down, and pre-move consumption levels line up with the origin quartile rather than the destination quartile, showing no evidence of differential trends. Symmetry of gains and losses and the absence of sorting on trends are precisely the patterns additive separability predicts; their presence in consumption supports carrying the wage decomposition over to spending. For the KSS analysis, we estimate the decomposition on the leave-one-out connected set. 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 cash transactions. To the extent that a household's reliance on cash is a fixed characteristic, the person or household fixed effect would explain an overly large share of the total variation in spending.²⁶ To minimize the influence of these observations, our baseline variance decomposition is conducted on a trimmed sample, where we remove workers whose household consumption ever drops below the 5th percentile of the consumption distribution. Appendix Figure A3(b) plots the consumption-to-wage ratios in this sample. Reassur-

²⁶To assess this concern, we examine the distribution of log consumption relative to log gross wages. Appendix Figure A2 shows that consumption exhibits a longer and fatter left tail than wages. Moreover, sorting workers by percentile of consumption in Appendix Figure A3(a) reveals a sizable group 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 reflects cash as the mode of payment.

ingly, the ratio is much more stable across the consumption distribution.²⁷ Second, we focus on single-person households to filter out variation in expenditure due to household size.

Results Table 8 reports variance decompositions for log earnings and log expenditures among singles, contrasting the standard AKM estimator with 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, a pattern that is particularly pronounced for log expenditures. The role of firms in earnings shrinks from 0.024 in AKM to 0.014 in KSS, while the correlation between firm and worker fixed effects increases from 0.019 to 0.290. Similarly, the variance of firm effects for consumption shrinks from 0.017 to 0.006 under KSS, 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 indicate that 72.3% of the variance in log earnings and 93.2% of the variance in log expenditures is due to persistent individual heterogeneity.

The key takeaway 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 (0.014 under KSS). Through the lens of the consumption-equivalent-variation approach, the evidence suggests that firms play a nontrivial role in shaping welfare, although their role through household consumption is considerably attenuated relative to wages.

We compare these consumption variance-decomposition results with those from the raw data and from a version of the data in which we instead winsorize at the 5th percentile. Panel A of Appendix Table A8 reports the corresponding KSS decompositions. Notably, the ratio of the variance of firm effects in consumption to wages is remarkably similar across specifications: the variance of firm effects on consumption is roughly half the size of the corresponding variance for wages in the baseline and winsorized samples, and below the wage variance in the raw data as well, though by a smaller margin there. This suggests that while firms influence both earnings and consumption, the relative contribution of firms is robust to how we handle these measurement issues. 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 household trait.

Lastly, we extend the analysis to all workers, including single- and multi-person households. We perform variance decompositions using individual earnings and consumption expenditure, as well as household-member equivalents, where we divide earnings and consumption by the square root of the number of household members. Both samples yield

²⁷We also show that the relative variance of consumption expenditure to wage income remains stable over time across the trimming and winsorizing approaches (Appendix Figure A6).

Table 8: Variance Decompositions Single-Person Households

	log earnings		log expenditure	
	AKM	KSS	AKM	KSS
explained variance	0.133	0.122	0.158	0.138
variance of firm effects	0.024 (0.183)	0.014 (0.112)	0.017 (0.105)	0.006 (0.041)
variance of person effects	0.107 (0.803)	0.089 (0.723)	0.156 (0.990)	0.128 (0.932)
covariance of firm and person effects	0.001 (0.014)	0.010 (0.165)	-0.008 (-0.095)	0.002 (0.027)
correlation of firm and person effects	0.019	0.290	-0.148	0.068
R squared	0.858	0.791	0.719	0.628

Notes: Variance decomposition of log earnings and log expenditures for single-person households on the leave-one-out connected set. In the implementation, we first restrict the AKM sample to the largest connected set of firms and center firm effects to have mean zero within that sample. Earnings winsorized at the 1st and 99th percentiles. Consumption winsorized at the 99th percentile. We then compute the 5th percentile of consumption and drop all consumption 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 parentheses refer to the fraction of explained variance.

similar estimates for the role of firms, as reported in Panels B and C of Appendix Table [A8](#).

5.3 Extensions

Our simple revealed-preference approach abstracts from aspects of firms that may affect worker welfare beyond their contribution to expected lifetime resources. Examples include workplace safety, flexible schedules, commuting convenience, workplace culture, and other job amenities that workers value directly. In such cases, consumption may not fully summarize welfare differences across firms. For example, workers may be willing to accept lower consumption in exchange for more desirable working conditions. If these amenities are normal goods, wealthier or higher-consumption individuals may sort disproportionately into firms offering such amenities. In the model of [Luo and Mongey \(2019\)](#), the testable implication is that high-consumption or high-wealth individuals should sort into lower-consumption firms that provide desirable non-pay amenities. However, the covari-

ance reported in Table 8 has the opposite sign and is weakly positive, accounting for only about two-and-a-half percent of the total explained variance. We also find that MPC estimates are relatively stable across households with high and low asset levels. Taken together, these findings provide limited evidence that sorting on non-pay amenities is an important driver of the gap between firm pay and firm consumption differentials in our data.²⁸

6 Model

This section demonstrates that an off-the-shelf model with search frictions and consumption-savings decisions can rationalize our empirical results, and then uses the model to reassess the role of firms in welfare inequality. The model serves three purposes. First, it formalizes what the estimated MPC out of income measures: without a model, it is not always clear which structural object the ratio of spending and income responses corresponds to. Second, it shows that a standard framework, without additional bells and whistles, quantitatively reproduces both the MPC and the variance-decomposition evidence. Third, it relaxes the log-additivity assumption underlying the sufficient-statistics approach of Section 5, allowing us to make the welfare statements more precise.

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 $\rho = \hat{\rho} + d$.

Workers derive utility, u , from consumption, c . We assume that the worker's utility function takes the CRRA form

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

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

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

where the evolution of assets follows

$$\dot{a} = ra + i - c. \quad (11)$$

where i is income in the given period.

²⁸We note that the MPC is somewhat lower in the top wealth and worker-ability quartiles, but we lack the statistical precision to draw firm conclusions from this evidence.

There is 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 because of ability differences, but they also differ ex post within each ability level because assets vary with their employment histories.

Search frictions are present, and workers search both on and off the job. During unemployment, workers receive offers at rate λ , and while employed they receive offers at rate $s\lambda$. An offer consists of a firm productivity level, z , drawn from an offer distribution F that describes the job ladder. We also allow for reallocation shocks while employed, arriving at rate $s^r\lambda$, under which workers are forced to “take it or leave it.”²⁹

We assume production is multiplicative in worker and firm type. Wages are determined in a bargaining setup following [Binmore et al. \(1986\)](#), [Hall and Milgrom \(2008\)](#), and [Elsby and Gottfries \(2022\)](#), in which threats to disrupt are credible. We provide a complete description of the two-stage bargaining game in the Appendix, but the key implication is a wage equation of the form

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

where β can be interpreted as worker bargaining power,³⁰ and ω_0 collects terms describing the cost of disruption to the worker and the firm. There are three other common approaches to wage determination. First, we could adopt the standard Nash bargaining setup, in which workers and firms split the surplus of the match in some way ([Pissarides, 2000](#)). Second, we could consider the wage-posting game in [Burdett and Mortensen \(1998\)](#), where wage dispersion is endogenous because of the dynamics of the job ladder. Third, we could use the sequential auction framework of [Postel-Vinay and Robin \(2002\)](#). We view our bargaining setup as a simplification of [Pissarides \(2000\)](#), but one that avoids concerns about non-convexity of the bargaining set in models with on-the-job search ([Shimer, 2006](#)). Moreover, allowing for asset heterogeneity complicates the wage-posting approach because reservation wages differ across unemployed workers.³¹ Workers receive a share b of the previous wage as in the Norwegian case.

Finally, as is common in the literature, in Sections 4 and 5 we run the AKM regression on gross labor income. We therefore need to take a stand on labor income taxation. This matters because progressive income taxation, as in Norway, compresses the pass-through from gross wages to consumption. We use a simple but accurate nonlinear tax rule, first

²⁹These shocks are included because downward wage moves are an important feature of the data, and because they help with connected-set arguments when running two-way fixed-effects regressions. “Take it or leave it” implies that the worker’s outside option becomes unemployment following the shock at rate $s^r\lambda$.

³⁰More precisely, it is the probability that the worker makes the offer in the second stage.

³¹We note that one could circumvent this with a binding minimum wage for all asset levels.

used in [Feldstein \(1969\)](#) and more recently popularized and validated in [Heathcote et al. \(2017\)](#). After-tax income is written as

$$w(\varepsilon, z) = \Lambda \tilde{w}(\varepsilon, z)^{1-\tau}, \quad (13)$$

where τ determines the degree of progressivity. If $\tau = 0$, the government operates a linear tax rule. If $\tau = 1$, there is full redistribution and no heterogeneity in net pay between firms. In our calibration, we set $\Lambda = 1$, and estimate a value of $\tau = 0.15$ which we discuss further in [Section C](#). We note that taxation slightly complicates the wage bargain. Maintaining consistency, our implicit assumption here is that firms and workers bargain over gross labor income.

Consumption choice

This concludes the model setup. In short, it resembles [Lise \(2013\)](#), but we include worker heterogeneity so that we can study worker effects in an AKM setting. In the next section, we target the MPC and AKM decomposition results and use the model structure to quantify the importance of firms for welfare inequality, as well as the extent to which workers mitigate firm-driven wage inequality. The consumption choice is central to this exercise, so we now describe it in this class of consumption-savings models.

Consumption solves the first-order-condition (FOC) written as

$$u'(c) = V_a \implies c(\varepsilon, z, a) = \chi(V_a(\varepsilon, z, a)), \quad (14)$$

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

Structural interpretation of the MPC along the ladder

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

$$c_w = \frac{V_{wa}}{u''(c)} = \frac{c_a + c_{wa}\dot{a}}{\tilde{\rho} - r + c_a \left(1 - \dot{a} \frac{u'''(c)}{u''(c)} \right)}, \quad (15)$$

where $\tilde{\rho} = \rho + s\lambda[1 - F(p)] + s^r\lambda + \delta$ is the implied social discount rate. Since $u''(c) < 0$, a larger consumption response c_w corresponds to a larger V_{wa} in absolute value.³² We derive these results in Appendix Section C. The consumption response and the variance of consumption across firms depend on how the marginal value of wealth changes with wages, V_{wa} . A higher V_{wa} implies that wage changes have a larger effect on the marginal value of wealth, reflecting weaker insurance against wage shocks along the job ladder.

The traditional notion of an MPC is the consumption response out of a transitory income shock, represented here by c_a . The MPC along the ladder, c_w , is a function of this object, scaled by the expected duration of the wage change. That duration depends on the effective discount rate $\tilde{\rho}$, which embeds not only impatience (ρ) but the frictions and risks of the environment: the on-the-job search hazard $s\lambda[1 - F(p)]$, the reallocation term $s^r\lambda$, and the separation hazard δ . A higher $\tilde{\rho}$ shortens the effective horizon over which a wage change is enjoyed, producing a smaller jump in consumption at a job transition.

The offer arrival rate λ shapes this in two ways. When λ is low, job changes are rare and largely unanticipated, so workers cannot pre-smooth and consumption jumps discretely at transitions. As λ rises, the current job becomes less important for lifetime earnings: V_{wa} falls, and the consumption response to any given change in firm pay, falls. Search frictions are thus central to consumption behavior following job transitions, and to how firm-level pay heterogeneity translates into welfare inequality.

Finally, note that the above results imply misspecification in the AKM consumption equation, since V_{wa} is heterogeneous across worker types. In contrast, the AKM in consumption imposes homogenous treatment out of pay differentials. In the remaining sections, we target these moments and consider how informative they are for welfare inequality in the calibrated model.

³²For simplicity, this derivation ignores the dependence of acceptance out of reallocation on productivity, and the dependence of current benefits on the previous wage. Under a first order approximation, we can also write the variance of consumption across firms for a particular worker type as

$$\text{var}_z(\ln(c(\varepsilon, z, a))) \approx \left((1 - \tau) c_w \frac{\tilde{w}}{c} \right)^2 \bigg|_{z=\bar{z}} \text{var}(\ln(z)), \quad (16)$$

where $c_w = V_{wa}/u''(c)$ is the consumption response defined above.

6.2 Calibration

We calibrate the model to the Norwegian economy at a quarterly frequency and choose parameters to replicate the variance of firm effects in consumption and wages, as well as our estimated MPC along the job ladder. This allows us to use the model to assess how firm pay heterogeneity affects welfare inequality 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 16 parameters.

Externally set parameters. We set the $\hat{\rho}$ to equal a 5% annual social discount rate, d to represent a 40 year working life on average, implying a social discount rate of $\rho = (-\log(0.95) + 1/40)/4$. We choose the worker bargaining power, β , to be the typical value in the literature of 0.5, and set the cost of disruption, ω_0 , to be 0. The cost of disruption being zero has two implications. First, firms are willing to accept all matches, and second, this delivers an AKM structure in wages. This allows us to assess the content of firm consumption effects in a model where wages are additively separable into a worker effect (worker ability) and firm effect (firm productivity). As discussed earlier, we set $\Lambda = 1$, and estimate a value of $\tau = 0.15$ in labor income tax function which we describe further in Section C.

Internally calibrated parameters. We calibrate the remaining 10 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 report the dynamics of the MPC in Figure A9. We observe no change in consumption before the switch and a large, persistent jump afterward, consistent with the data. A key target is our estimated MPC and the firm consumption effects from the previous sections. To generate these responses, the value of the CRRA coefficient is crucial. Figure A8 shows that the model MPC is strongly decreasing in α , especially at lower values of α . As we approach log

utility ($\alpha \rightarrow 1$), the MPC approaches 0.9. By contrast, the MPC is around 0.45 for $\alpha = 5$. This is intuitive: as workers become more risk-averse, they save more aggressively for periods of unemployment at the top of the job ladder, which reduces the pass-through from wage inequality to consumption inequality and lowers the MPC out of income following job moves. Note, however, that even high levels of risk aversion still yield a large MPC due to frictions and borrowing constraints. Our calibrated value is around 3.2, well within the range 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.

Table 9: Internally calibrated parameters and model fit

Parameters	Value	Moments	Data	Model
λ	2.1	UE rate from Bhuller et al. (2023)	0.6	0.6
δ	0.01	EU rate from Bhuller et al. (2023)	0.041	0.04
s	0.015	JJ rate from Bhuller et al. (2023)	0.019	0.02
s^r	0.027	Wage cuts	0.3	0.3
b	0.49	Replacement rate	0.5	0.5
r	0.0072	Wealth-to-income ratio	0.5	0.5
α	3.2	MPC	0.5	0.49
β_z^1	0.15	Firm c effect variance	0.0057	0.0056
β_z^2	1.5	Firm w effect variance	0.014	0.013
Δ_ε	0.4	Person w effect variance	0.088	0.088

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.

An important feature of wage inequality is the substantial sorting present in the data: more productive workers tend to work in more productive firms. This is often summarized by the positive covariance between worker and firm effects, which we also document in [Section 5](#). We observe the same pattern for consumption, though to a lesser extent. The model reproduces a positive covariance for both wages and consumption. The mechanism is wealth sorting out of unemployment. High-ability workers accumulate more wealth, and because they have more wealth, they accept higher-paying jobs. This “wealth sorting” generates a positive covariance term for wages (0.001) and consumption (0.0005). While the model reproduces the direction qualitatively, the covariance term for wages in particular is substantially smaller than in the data.

To conclude this section, an off-the-shelf search model with on-the-job search and savings provides a comprehensive account of the empirical patterns we documented in the preceding sections. In the following subsection, we now use the calibrated model for welfare analysis.

6.3 Welfare Inequality in the Model

In Section 5, we showed that the variance of firm effects in log consumption maps directly onto 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 consumption-savings model studied here, consumption is not log-additive. Instead, the CEV gain from moving between firms depends on worker characteristics and solves

$$u(c_{ij}(1 + \gamma_{ijk})) = u(c_{ik}). \quad (17)$$

What distinguishes this welfare measure from Section 5 is that γ_{ijk} depends on the worker i : low-ability workers may adjust consumption by more than high-ability workers. We first aggregate consumption to the level of a worker's permanent type within each firm, averaging over the asset distribution, and then compute welfare, so that i indexes a worker's permanent ability rather than their assets. This mirrors how firm consumption effects are measured in Section 5, where worker type is a permanent characteristic. Accordingly, γ_{ijk} measures the gain or loss to a type- i worker of being at firm k rather than firm j in the long run, once the worker holds the asset levels typical of each firm. This also means that the asset heterogeneity across firms — itself a consequence of the wage differences workers experience along the job ladder — is retained in the measure rather than differenced away. We refer to this as the *static* welfare measure, because it evaluates current-period utility, $u(c_t)$; the *dynamic* measure below instead evaluates the discounted stream $\sum_t u(c_t)$. The distinction concerns only the information the welfare measure aggregates: workers are forward-looking in both cases, so the static measure already encodes expectations about future income through today's consumption choices.


Since the consumption-equivalent welfare gains γ_{ijk} are heterogeneous and depend on the individual i , we aggregate across individuals within a firm. We compute the firm-level average CEV as

$$\gamma_k = \int \int g_{ij} \gamma_{ijk} di dj, \quad (18)$$

where g_{ij} denotes the ergodic joint distribution of individuals i in firm j . Firms at the bottom of the wage distribution have negative γ_k , reflecting welfare losses from moves to those firms, while firms at the top have positive γ_k , reflecting welfare gains. The variance of γ_k across firms therefore captures the extent to which firm wage inequality contributes to welfare inequality.

Figure 8 shows the results in levels, while Table 10 reports them relative to inequality measured using gross labor income. In Panel A of Table 10, for comparison, we report the

Figure 8: Variance of firm effects and welfare in the model



Notes: The figure shows that the variance of firm consumption effects (Panel A, column (ii)) is close to, though slightly smaller than, the variance of welfare attributable to firms (Panel B, column (iii)). Panel A shows the variance of firm wage and consumption effects in the model which are close to identical to the empirical counterparts. Panel B shows relevant welfare measures. Wage in panel B, refers to taking a worker's input to utility as the gross wage they receive. Net wage in panel B, refers to taking a worker's input to utility as the net wage they receive. Finally, Consumption in panel B, refers to taking a worker's input to utility as their consumption. The dotted horizontal lines are at the variance of firm wage effects and firm consumption effects (Panel A) for comparison.

ratio of the variance of firm consumption effects to the variance of firm wage effects, which under log-additivity of consumption is equivalent to the welfare measure, as shown in Section 5. Since firm wage effects are defined on gross labor income, we report the welfare measures relative to the benchmark in which workers consume their gross labor income. In Panel B, we report the variance of welfare inequality. A value of 0.74 implies that the variance of welfare when disposable income is the input to utility is 74% of the variance obtained when gross pay is the input. The corresponding value of 0.49 for consumption suggests that taxation and consumption smoothing substantially, and equally, mitigate the welfare consequences of gross wage inequality.

Finally, we turn to the dynamic measure and compute the *present discounted value* welfare gains from moving up the pay ladder.³³ The model is well suited to this calculation because it jointly identifies the two forces that determine it: the extent of search frictions, and the response of consumption to moving along the ladder. According to the model, a one-standard-deviation move up the firm pay ladder is associated with dynamic welfare gains of 2.3%. By contrast, the corresponding change in lifetime gross pay is 12.2% and the change in disposable income 10.3%.

³³Measured as $\mathbb{E}_0 \int_0^\infty e^{-\rho t} u(c_{ijt}(1 + \gamma_{ijk})) dt = \mathbb{E}_0 \int_0^\infty e^{-\rho t} u(c_{ikt}) dt$.

Table 10: Firm Effects and Welfare Inequality

	A. Firm Effects	B. Welfare Inequality	
	Consumption	Disposable income	Consumption
Proportional to wages	0.42	0.74	0.49

Notes: This table reports the ratio of the variance of consumption firm effects to wage firm effects in panel A. Panel B shows static welfare inequality (solving (17), with consumption aggregated to the worker’s permanent type within each firm) when using disposable income as the input to utility, or consumption as the input to utility, both relative to welfare inequality when using gross income as the input to utility. All figures are computed from the ergodic distribution. For panel A, the variance of wage firm effects is 0.0132. For panel B, welfare inequality when using gross income as the input to utility is 0.0132.

7 Conclusion

A large literature has established that firms matter for workers’ earnings. In this paper, we extend this research by studying whether firms also affect household consumption. This question is nontrivial, because the transmission from firm pay to household spending depends on several unobserved factors. It also has important positive and normative implications, because it informs our understanding of search frictions, workers’ income expectations, and the welfare consequences of firm pay differentials. However, making progress on this question is difficult because of substantial conceptual and empirical challenges.

We address these measurement and conceptual challenges by combining newly collected data on consumption expenditure and employment from Norway with a revealed-preference approach. Using a movers design, we show that firm pay differentials translate into substantial consumption differences: each \$1 in additional income due to firm pay premia leads to approximately \$0.5 in additional consumption. This estimate falls well short of the permanent income benchmark of 1, because a firm’s fixed pay premium during the job need not be perceived as a permanent income shift by workers who anticipate further movement along the job ladder. We also document important heterogeneity across spending categories: pay differentials primarily translate into non-essential spending differences, such as those for hotels and restaurants. In contrast, we find little effect on grocery expenditure, which remains largely unresponsive to firm pay differentials.

We then perform a variance decomposition to quantify the relative importance of firms in shaping consumption and wage inequality. We find that firm effects explain roughly half as much of consumption inequality as of wage inequality, indicating that household spending decisions compress part of the dispersion in firm pay. Lastly, we use a simple search-and-savings model to interpret these patterns and to study when they are informa-

tive about welfare. Overall, our results show that firms are an important source of consumption inequality, and that the transmission of firm pay differences to household spending is substantially attenuated relative to wages. Through the lens of the model, these findings suggest that self-insurance and consumption smoothing attenuate the extent to which firm pay inequality maps into differences in household well-being.

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. JOHANNESSEN, 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.
- AUDOLY, R., M. BHULLER, AND T. A. REIREMO (2024): "The pay and non-pay content of job ads," *arXiv preprint arXiv:2407.13204*.
- AUTOR, D., A. KOSTOL, M. MOGSTAD, AND B. SETZLER (2019): "Disability benefits, consumption insurance, and household labor supply," *American Economic Review*, 109, 2613–2654.
- 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.
- BHULLER, M., D. FERRARO, A. R. KOSTØL, AND T. C. VIGTEL (2023): "The internet, search frictions and aggregate unemployment," Working paper, National Bureau of Economic Research.
- BINMORE, K., A. RUBINSTEIN, AND A. WOLINSKY (1986): "The Nash bargaining solution in economic modelling," *The RAND Journal of Economics*, 17, 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, vol. 108, 317–322.
- BLUNDELL, R., L. PISTAFERRI, AND I. PRESTON (2008): "Consumption Inequality and Partial Insurance," *American Economic Review*, 98, 1887–1921.
- BLUNDELL, R. AND I. PRESTON (1998): "Consumption inequality and income uncertainty," *The Quarterly Journal of Economics*, 113, 603–640.
- BOROVÍČKOVÁ, K. AND R. SHIMER (2024): "Assortative Matching and Wages: The Role of Selection,"

- Working Paper 33184, National Bureau of Economic Research.
- 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, 1075–1090.
- CARD, D., A. R. CARDOSO, J. HEINING, AND P. KLINE (2018): "Firms and labor market inequality: Evidence and some theory," *Journal of Labor Economics*, 36, S13–S70.
- CARD, D., A. R. CARDOSO, AND P. KLINE (2016): "Bargaining, sorting, and the gender wage gap: Quantifying the impact of firms on the relative pay of women," *The Quarterly Journal of Economics*, 131, 633–686.
- CARD, D., J. HEINING, AND P. KLINE (2013): "Workplace heterogeneity and the rise of West German wage inequality," *The Quarterly Journal of Economics*, 128, 967–1015.
- 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," Working paper, 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.
- DE GIORGI, G., A. FREDERIKSEN, AND L. PISTAFERRI (2020): "Consumption network effects," *The Review of Economic Studies*, 87, 130–163.
- DEATON, A. (1992): *Understanding consumption*, Oxford University Press.
- DI ADDARIO, S., P. KLINE, R. SAGGIO, AND M. SØLVSTEN (2023): "It ain't where you're from, it's where you're at: hiring origins, firm heterogeneity, and wages," *Journal of Econometrics*, 233, 340–374.
- 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.
- FELDSTEIN, M. S. (1969): "The effects of taxation on risk taking," *Journal of Political Economy*, 77, 755–764.
- FRIEDMAN, M. (1957): *A Theory of the Consumption Function*, Princeton, NJ: Princeton University Press.
- 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 (2021): "Job Displacement Insurance and (the Lack of) Consumption-Smoothing," *American Economic Review*, 111, 899–942.
- 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," *Umeå Economic Studies* 1039, Umeå University.
- HALL, R. E. AND P. R. MILGROM (2008): "The limited influence of unemployment on the wage


- bargain," *American Economic Review*, 98, 1653–1674.
- HEATHCOTE, J., K. STORESLETTEN, AND G. L. VIOLANTE (2017): "Optimal tax progressivity: An analytical framework," *The Quarterly Journal of Economics*, 132, 1693–1754.
- (2020): "How Should Tax Progressivity Respond to Rising Income Inequality?" *Journal of the European Economic Association*, 18, 2715–2754.
- JAPPELLI, T. (1990): "Who is credit constrained in the US economy?" *The Quarterly Journal of Economics*, 105, 219–234.
- 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," Working Paper 33084, National Bureau of Economic Research.
- (2025): "Labor Market Monopsony: Fundamentals and Frontiers," Working Paper 33467, National Bureau of Economic Research.
- KRUEGER, D. AND F. PERRI (2006): "Does income inequality lead to consumption inequality? Evidence and theory," *The Review of Economic Studies*, 73, 163–193.
- LACHOWSKA, M., A. MAS, R. SAGGIO, AND S. A. WOODBURY (2023): "Do firm effects drift? Evidence from Washington administrative data," *Journal of Econometrics*, 233, 375–395.
- 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," Working paper, 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 Sorting—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.
- LUO, M. AND S. MONGEY (2019): "Assets and job choice: Student debt, wages and amenities," Working paper, National Bureau of Economic Research.
- MORCHIO, I. AND C. MOSER (2024): "The gender pay gap: Micro sources and macro consequences," Working paper, National Bureau of Economic Research.
- NORGES OFFENTLIGE UTREDNINGER (2009): "Skatte- og avgiftssystemet – lov om skatt av formue og inntekt (skatteloven)," Tech. Rep. NOU 2009:10, Norwegian Ministry of Finance.
- 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*, 133, 1331–1393.
- 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.

Online Appendix


A Additional Material

A.1 Figures



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: Density of Log Consumption Expenditure and Gross Wages




Notes: The figure plots the density of log household consumption expenditure and log household gross wages, both equivalized by the square root of household size, for the full sample of workers (single- and multi-person households).

Figure A3: Consumption-to-Wages Across Consumption


A. Full Sample

B. Trimmed Sample




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 A4: 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 A5: Consumption responses to changes in firm pay: Married and cohabiting couples




Notes: The figure reports coefficient estimates from the reduced form equation (3) for the main person (i.e., the mover) and the spouse using split-sample firm effects.

Figure A6: Relative Variance of Consumption (across samples)

A. Single-Person Households




B. Multi-Person Households




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 A7: Mean log wages and log consumption around job moves




Notes: This figure plots mean log wages (left column) and mean log consumption (right column) around job moves. We bin movers' origin and destination firms into quartiles of the leave-1%-out split-sample firm effects, estimated from wages for the left column and from consumption for the right column. Q1 → Q4, for example, denotes a move from a firm in origin quartile 1 to a destination firm in quartile 4. The top row plots mean outcomes by year relative to the move for each origin–destination quartile pair. The bottom row assesses the symmetry of gains and losses: each point corresponds to a pair of quartiles and plots the mean change for upward moves (horizontal axis) against the mean change for downward moves between the same two quartiles (vertical axis), where changes are averaged over the three post-move years relative to the year before the move; the dashed line indicates gains and losses of equal magnitude. The sample consists of full-time workers aged 25–60 who moved between firms between 2011 and 2015.

Figure A8: Varying the CRRA parameter



Notes: The figure reports model outcomes under alternative values of the coefficient of relative risk aversion, holding the remaining parameters at their baseline calibrated values from Table 9.

Figure A9: Model MPC



Notes: The figure reports the marginal propensity to consume implied by the calibrated model, computed from the model's policy functions under the baseline calibration in Table 9.

A.2 Tables

Table A1: Sample Restrictions

Sample:	Number of observations
Movers in 2011-2015, aged 25-60, working 30 hours per week	157320
Only one mover per household	145217
At least 36 months of tenure	122567
No change in hours of work per week	87940
Single living alone or with children	23698

Notes: The table reports the number of individual workers (not worker-years) remaining after each successive sample restriction, for movers between 2011 and 2015. The estimation samples in Tables 5 and 6 pool observations over event years 1, 2, and 3 for these workers, and additionally require that marital status is observed in the year before the move and that the split-sample firm-pay measure and both outcomes are non-missing.

Table A2: First stage, reduced form, and 2SLS estimates: alternative household definitions

	<i>Panel A: Single person</i>		<i>Panel B: Multi-person households</i>		
	Split sample	Co-worker	Household	Own	Spouse
	(1)	(2)	(3)	(4)	(5)
Net-income (First stage)	0.272 (0.015)	0.086 (0.008)	0.220 (0.013)		
Consumption (Reduced form)	0.136 (0.023)	0.044 (0.010)	0.124 (0.019)	0.141 (0.013)	-0.017 (0.015)
MPC (2sls)	0.499 (0.079)	0.519 (0.112)	0.565 (0.080)	0.643 (0.064)	-0.078 (0.070)
N	38631	37142	196788	196788	196788

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 single-person households (singles without children) in Panel A and multi-person households (married and cohabiting couples together with singles who have children) in Panel B. In Panel A, column (1) uses leave-1%-out split-sample 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. In Panel B, column (3) reports total household results, while columns (4) and (5) split the household response into the mover's own and the spousal consumption change, respectively; for households without a partner, the spousal change is zero by construction. Outcomes are normalized as in Table 6. All specifications include the baseline controls of Table 6. The sample consists of full-time workers aged 25-60 who moved between firms between 2011 and 2015. Robust standard errors in parentheses are clustered at the firm level.

Table A3: First stage, reduced form, and 2SLS estimates under alternative firm-pay measures

	(1)	(2)	(3)
	Leave 50% out	Leave 1% out	Full sample
Net-income (First stage)	0.177 (0.012)	0.273 (0.014)	0.439 (0.017)
Consumption (Reduced form)	0.094 (0.017)	0.134 (0.020)	0.210 (0.023)
MPC (2sls)	0.528 (0.093)	0.491 (0.072)	0.479 (0.051)
N	51860	51860	51860

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. Column (1) uses firm effects measured using the half-sample split specification (AKM regression run twice). Column (2) reports the leave-1%-out specification on the common comparison sample. Column (3) uses firm effects measured using the whole sample (AKM regression run once). All three columns are estimated on the common sample of observations for which all three firm-pay measures are non-missing; this sample is slightly smaller than the baseline sample in Table 6, which explains the small difference between column (2) and the baseline estimate. The sample consists of single full-time workers aged 25-60 who moved between firms between 2011 and 2015. Robust standard errors in parentheses are clustered at the firm level.

Table A4: Sensitivity to other potential confounders

	Baseline (1)	Within CZ (2)	Within Residence (3)	Within HH size (4)
MPC 2SLS Estimate (standard error)	0.487 (0.071)	0.488 (0.071)	0.505 (0.080)	0.492 (0.072)
Observations	52853	52322	46591	48872

Notes: Column (1) reports the baseline 2SLS specification of Table 6. We then condition on a worker remaining in the commuting zone in Column (2), remaining at the same residence in Column (3), and maintaining the same number of individuals in the household in Column (4). The sample consists of single full-time workers aged 25-60 who moved between firms between 2011 and 2015. Robust standard errors are clustered at the firm level.

Table A5: Sensitivity to inclusion of controls

	No controls (1)	Worker (2)	+ Market (3)	+ Firm (4)	Amenity controls (5)
MPC 2SLS Estimate (standard error)	0.488 (0.068)	0.479 (0.069)	0.487 (0.071)	0.490 (0.071)	0.493 (0.071)
Observations	52853	52853	52853	52853	52853

Notes: Column (1) reports the 2SLS estimate with no control variables. Column (2) adds age, gender, and number of children indicators as control variables, and Column (3) additionally includes calendar year and origin-region indicators; Column (3) is therefore the baseline specification $W_{i,t}$ used throughout the paper. Column (4) adds to the Column (3) specification the change in the firm-specific unemployment risk component following [Lachowska et al. \(2025\)](#), average firm employment growth, and the size of the origin and destination firms. Column (5) instead adds to the Column (3) specification proxies for non-pay amenities and peer effects: the variability of contracted hours, the share of employees with young children, and the squared difference between the average consumption of stayers in the destination firm and the mover's own pre-move household consumption. The sample consists of single full-time workers aged 25-60 who moved between firms between 2011 and 2015. Robust standard errors are clustered at the firm level.

Table A6: Sensitivity to the pooling horizon: three vs seven post-move years

	Years 1–3	Years 1–7
Net-income (First stage)	0.187 (0.032)	0.186 (0.036)
Consumption (Reduced form)	0.103 (0.046)	0.099 (0.041)
MPC (2SLS)	0.548 (0.239)	0.535 (0.214)
N	10647	24843

Notes: The table reports coefficients from the first stage (2), reduced form (3), and two-stage least squares regressions on a balanced seven-year sample: movers in 2009–2011 who remain single and have income and spending observed in all seven post-move years. Column (1) pools post-move years 1–3 and column (2) pools years 1–7 on this same sample, so the two columns differ only in the pooling horizon and the number of observations per worker (three and seven, respectively). The specification, including the control set, is identical to the baseline in Table 6. Robust standard errors in parentheses are clustered at the firm level.

Table A7: MPC Heterogeneity

	(1)	(2)	(3)	(4)		(1)	(2)	(3)	(4)
<i>Panel A: worker types</i>					<i>Panel B: firm types</i>				
Group	MPC	SE	E[c]	N	Group	MPC	SE	E[c]	N
Assets					Firm origin				
Q1	0.557	(0.147)	284,766	13,911	Q1	0.548	(0.157)	306,304	13,342
Q2	0.503	(0.134)	332,763	12,966	Q2	0.626	(0.199)	326,943	13,079
Q3	0.534	(0.142)	356,379	13,002	Q3	0.470	(0.132)	351,689	13,328
Q4	0.372	(0.140)	412,031	12,974	Q4	0.306	(0.108)	392,831	13,104
Worker effect					Firm destination				
Q1	0.521	(0.174)	255,567	13,897	Q1	0.278	(0.224)	312,511	13,181
Q2	0.712	(0.135)	304,800	13,291	Q2	0.471	(0.189)	329,927	13,252
Q3	0.460	(0.138)	342,671	12,478	Q3	0.605	(0.145)	343,061	13,341
Q4	0.301	(0.130)	469,481	11,970	Q4	0.562	(0.110)	391,723	13,079
Sex					Direction				
Female	0.629	(0.158)	343,674	19,262	Down	0.305	(0.245)	340,311	25,988
Male	0.440	(0.077)	345,026	33,591	Up	0.611	(0.115)	348,622	26,865
Age									
<30	0.466	(0.160)	280,196	3,662					
30-50	0.460	(0.075)	362,380	31,495					
>50	0.561	(0.194)	370,945	17,696					

Notes: The table reports 2SLS estimates of the MPC out of income across sample splits. All specifications include the baseline controls of Table 6. E[c] denotes mean consumption in the three years prior to the switch, in real 2018 NOK. N is the number of observations. The sample consists of single full-time workers aged 25-60 who moved between firms between 2011 and 2015. Robust standard errors in parentheses are clustered at the firm level.

Table A8: KSS Variance Decomposition: Robustness

	Baseline		Winsorized	Raw data
	Annual wage income	Spending	Spending	Spending
(A)	Single-person households			
Var expl.	0.122	0.138	0.220	0.404
Var(firm)	0.014 (0.112)	0.006 (0.041)	0.006 (0.027)	0.010 (0.024)
Var(person)	0.089 (0.723)	0.128 (0.932)	0.206 (0.940)	0.376 (0.930)
Cov(firm,person)	0.010 (0.165)	0.002 (0.027)	0.004 (0.033)	0.009 (0.046)
Corr(firm,person)	0.290	0.068	0.104	0.154
R squared	0.791	0.628	0.729	0.778
(B)	All workers: individual earnings and consumption			
Var expl.	0.136	0.329	0.301	0.408
Var(firm)	0.012 (0.086)	0.007 (0.021)	0.006 (0.019)	0.008 (0.020)
Var(person)	0.104 (0.766)	0.317 (0.965)	0.288 (0.955)	0.390 (0.958)
Cov(firm,person)	0.010 (0.148)	0.002 (0.014)	0.004 (0.026)	0.005 (0.023)
Corr(firm,person)	0.288	0.050	0.094	0.083
R squared	0.817	0.693	0.734	0.732
(C)	All workers: equivalized earnings and consumption			
Var expl.	0.142	0.130	0.187	0.253
Var(firm)	0.007 (0.050)	0.003 (0.024)	0.003 (0.019)	0.005 (0.019)
Var(person)	0.126 (0.887)	0.126 (0.973)	0.181 (0.970)	0.244 (0.966)
Cov(firm,person)	0.004 (0.062)	0.000 (0.002)	0.001 (0.011)	0.002 (0.015)
Corr(firm,person)	0.148	0.008	0.041	0.057
R squared	0.768	0.608	0.677	0.695

Notes: KSS variance decomposition across different samples and measurements. The baseline KSS results reported in Table 8 correspond to the Baseline columns in Panel (A), where consumption is trimmed at the 5th percentile. The Winsorized column reports results when we instead winsorize at the 5th percentile, and the Raw data column when we winsorize at the 1st percentile.

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). We also have information on credit transactions, which are inbound transactions, typically including pay from employers, typically occurring once per month for full-time employees.²

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

Since COICOP classifications are not part of the meta-data stored by Nets Branch Norway, Norges Bank provided a crosswalk between each MCC and COICOP, and each 5 digit NACE and COICOP. The consumer categorization is based on these crosswalks. Among existing physical card terminals, 42,600 out of about 162,000 were classified with MCC "Miscellaneous and Specialty Retail Stores" (MCC = 5999) or had a missing MCC. These were classified using the registered information on the terminal, typically the name and NACE-code of the store.³

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.

¹MCC 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.

²We calculate the total of all inbound transactions in 2013 full-time workers, and find that the average monthly number of inbound transactions is 1.3, which includes pay and reimbursements.

³All cross-walks are available upon request.

Table B.1: Consumption Categories


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

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

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

Overall BankAxept is the dominant means of card payment by Norwegians as shown in Figure B.2, panel (a). Over the time period 2006-2018 Bankaxept accounts for above 80

Figure B.1: BankAxept transactions and total amount.




Notes: The figure compares the total number of BankAxept transactions (left) and the total transacted amount (right) in our micro data with the aggregate values reported in official payment statistics.

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


Of the missing payments in Figure B.2, panel (b), credit card accounts for roughly 40 percent the payments. Despite not having the actual credit card purchases in our data, the granularity allows us to impute payments of credit card bills. In our wire transfer measure, payments to banks is a separate consumer category. Using an imputation algorithm (described below), we can further separate out payments of credit card bills. Our imputed credit card payments covers on average 63 percent of total credit card payments from the official statistics. In figure B.3, when we add this imputed measure to total bankaxept payments, we cover around 70 percent of total card payments and cash withdrawals (ATM and terminal) by Norwegians over the sample period.

Figure B.2: BankAxept share of payments.



Notes: Panel (a) shows BankAxept payments as a share of domestic card payments in official statistics. Panel (b) adds cash withdrawals from ATMs and terminals to the denominator.

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



Notes: The figure shows BankAxept payments plus imputed credit card payments as a share of total card payments and cash withdrawals (ATM and terminal) in official statistics.

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 con-

sumption (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 if the 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.⁴

⁴We 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

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 payments⁵. Moreover, when we compare to aggregate statistics on credit card payments, Figure B.4 shows that our cleaned debit category 13 tracks the official statistics on credit card payments well. We do not expect the algorithm to pick up all credit card bill payments, as internal bill payments (made from an account in the same bank that issued the credit card) does not show up in our transaction data.

Figure B.4: Imputed credit card spending.



Notes: The figure compares our imputed credit card payments (cleaned category-13 wire transfers) with official statistics on credit card payments.

⁵For 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 worker's asset stock. Define the relative acceptance strategies, $z^u(\varepsilon, a)$ and $z^e(z, \varepsilon, a)$

$$V^e(z^u(z, \varepsilon, a), \varepsilon, a) := V^u(z, \varepsilon, a) \text{ and } V^e(z^e(z, \varepsilon, a), \varepsilon, a) := V^e(z, \varepsilon, a), \quad (19)$$

where the thresholds are unique since $V_z^e > 0$. 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

$$\rho V^u(z, \varepsilon, a) = \max_{c^u} \left\{ u(c^u) + \lambda \int_{z^u(\varepsilon, a)} (V^e(x, \varepsilon, a) - V^u(z, \varepsilon, a)) dF(x) + V_a^u(z, \varepsilon, a) \dot{a} \right\}, \quad (20)$$

and

$$\begin{aligned} \rho V^e(z, \varepsilon, a) = \max_{c^e} \left\{ u(c^e) + s\lambda \int_z (V^e(x, \varepsilon, a) - V^e(z, \varepsilon, a)) dF(x) \right. \\ \left. + s^r\lambda \int_{z^u(z, \varepsilon, a)} (V^e(x, \varepsilon, a) - V^e(z, \varepsilon, a)) dF(x) \right. \\ \left. + (\delta + s^r\lambda F(z^u(z, \varepsilon, a))) [V^u(z, \varepsilon, a) - V^e(z, \varepsilon, a)] + V_a^e(z, \varepsilon, a) \dot{a} \right\}, \quad (21) \end{aligned}$$

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 \(2022\)](#).⁶ The continual renegotiation means that 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

⁶See [Grindaker and Simmons \(2025\)](#) who adopt the same setup we use here.

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 $p - z \geq -(1 - \beta)K$, thus the worker demands $w = z + (1 - \beta)K$. 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 $\beta(z - (z + (1 - \beta)K)) + (1 - \beta)(z - (\omega - \beta K))$. The equilibrium wage, w^* , is then equal to

$$w^* = \beta z + (1 - \beta)\omega. \quad (22)$$

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

C.3 Consumption change derivations

Throughout, we ignore the dependence of z^u , the worker's benefits in unemployment, on p . Taking the derivative of the employed HJB equation with respect to a and then z , we arrive at

$$(\tilde{\rho} - r)V_{az}^e(z, \varepsilon, a) = V_{aa}^e(z, \varepsilon, a)\dot{a} + V_{aa}^e(z, \varepsilon, a)w'(z, \varepsilon). \quad (23)$$

Using the FOC for consumption, we can write

$$V_{az}^e(z, \varepsilon, a) = u''(c)c_z, \quad V_{aa}^e(z, \varepsilon, a) = u''(c)c_a, \quad \text{and} \quad V_{aaz}^e(z, \varepsilon, a) = u'''(c)c_z c_a + u''(c)c_{az}. \quad (24)$$

Substituting these into (23), we arrive at the MPC result in the main text.

For the variance, take a first order approximation of the natural logarithm of consumption around the mean of the natural logarithm of firm productivity, $\ln \bar{z}$,

$$\ln c(\varepsilon, z, a) \approx \ln c(\varepsilon, \bar{z}, a) + \left. \frac{\partial \ln c(\varepsilon, z, a)}{\partial \ln z} \right|_{z=\bar{z}} (\ln z - \ln \bar{z}) \quad (25)$$

$$\equiv \ln c(\varepsilon, \bar{z}, a) + \frac{z}{c} c_w(\varepsilon, z, a) \Big|_{z=\bar{z}} (\ln z - \ln \bar{z}). \quad (26)$$

Given the MPC along the ladder, we arrive at the expression in the text.

C.4 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

$$\mathbf{T} = \mathbf{\Lambda} + \mathbf{\Delta} + \mathbf{P} \quad (27)$$

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

$$\mathbf{V}^{n+1} = \left(I\left(\frac{1}{\Gamma} + \rho\right) - \mathbf{T}\right)^{-1} \left(u(\mathbf{c}^n) + \frac{1}{\Gamma}\mathbf{V}^n\right), \quad (28)$$

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

$$0 = \mathbf{T}^T \mathbf{g}, \quad (29)$$

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 all the moments using the ergodic distribution from (29). Stack these moments into the vector $\mathbf{m}^s(\vartheta)$, where ϑ 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

$$\hat{\vartheta} = \arg \min_{\vartheta} \left\{ [\mathbf{m}(\vartheta_0) - \mathbf{m}^s(\vartheta)]' \mathbf{W} [\mathbf{m}(\vartheta_0) - \mathbf{m}^s(\vartheta)] \right\}. \quad (30)$$

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

C.5 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 [Borovičková 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

$$\sum_{t \in \mathcal{T}_i} \left(\log c_{it} - \hat{\alpha}_i - \hat{\psi}_{J(it)} \right) = 0, \quad (31)$$

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

$$\sum_t \sum_{i \in \mathcal{I}_{jt}} \left(\log c_{it} - \hat{\alpha}_i - \hat{\psi}_j \right) = 0, \quad (32)$$

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 $T \rightarrow \infty$, 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 worker's 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 worker's unobserved type, let us call it x_i , must be their 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 $\hat{\psi}_j$. Again, $\bar{\psi}_y$ is the firm's firm effect estimated on data we have described ($T \rightarrow \infty$, and there are a very large number of workers and firms).

Using the FOC conditions from above for α_i ((31)), and ψ_j ((32)), we can solve for $\bar{\alpha}_x$ and $\bar{\psi}_y$, up to an additive constant, on a computer using

$$\bar{\alpha}_x = \frac{\sum_y (\log(c_{xy}) - \bar{\psi}_y) g_{xy}}{\sum_y g_{xy}} \quad \text{and} \quad \bar{\psi}_y = \frac{\sum_x (\log(c_{xy}) - \bar{\alpha}_x) g_{xy}}{\sum_x g_{xy}}, \quad (33)$$

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 can 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.

C.6 Estimation of Tax Progressivity

We estimate the tax progressivity parameter τ in equation (13) following the cross-country approach in [Heathcote et al. \(2020\)](#) (Table 2). Under the log-linear tax function, τ can be expressed as a function of the average marginal tax rate (AMTR) and the average tax rate (ATR):

$$\tau = \frac{\text{AMTR} - \text{ATR}}{1 - \text{ATR}}. \quad (34)$$

We obtain estimates of AMTR and ATR for Norway from the World Tax Indicators database and compute τ directly using this expression. Using 2005 data (last available year) we obtain $\tau = 0.154$.

As a simple validation, the statutory tax schedule for Norway in 2009 implies that average labour income tax rates increase approximately linearly from about 20 percent to 35 percent over the income range 200,000–800,000 NOK (see [Norges offentlige utredninger \(2009\)](#), Figure 12.2). This corresponds to a value of $\tau \approx 0.15$.