Job Mobility and Consumption Inequality

Itzik Fadlon, University of California, San Diego and NBER
Ezra Karger, Federal Reserve Bank of Chicago
Shanthi P. Ramnath, Federal Reserve Bank of Chicago

March 2025

Extended Abstract:

1. Motivation and Overview

Income inequality is a central topic of policy debate in the US and a key theme of a large and active body of work that aims to understand its root causes to help inform policy efforts. A recent series of influential studies (summarized in Card et al. 2018) emphasizes the pivotal role that employers could play in amplifying income inequality. Specifically, Song et al. (2019) find that two-thirds of the rise in earnings inequality in the US starting in the 1980s are attributable to the rise in earnings dispersion across firms. Despite their broad economic importance, however, measures of income—an input into the household's budget—may less accurately identify economic well-being than measures of consumption—a direct input into the household's utility (Attanasio and Pistaferri 2014, Meyer and Sullivan 2023). In particular, consumption decisions reflect the long-run effects of transitory earnings changes via the household's ability to self-insure and smooth income shocks, and they also account for the insurance protection of the tax and transfer system. As such, analyzing consumption inequality and its drivers is necessary for establishing the welfare consequences of variation in income or earnings.

This paper provides, to the best of our knowledge, the first quasi-experimental estimates for the role of firms in determining American households' consumption patterns and the resulting consumption inequality. We use linked financial transaction-level data spanning millions of US households' bank accounts and credit cards. The richness of the data allows us to construct a high-frequency panel of employer-employee matches with detailed information on income flows and expenditures along with geographic location. For identification, we exploit worker mobility in an event-study design around job changes to estimate employer effects on consumption dispersion. We then decompose these consumption effects to trace their underlying sources, focusing on a direct income component, a location component, and an unexplained group-level component, which incorporates all other workplace-level interactions such as peer effects. We will later map our findings to modeling fundamentals in the dynamic household consumption decision. Overall, our analysis deepens our understanding of the sources and dynamics of consumption inequality and has implications for policies aimed at reducing existing gaps.

2. Data and Analysis Sample

Data. We rely on data from a financial data aggregator that collects information from bank and credit card accounts of more than 5 million households. The data span the years 2011-2023 and the financial aggregator obtains the data by providing financial services software to banks and collecting data from bank customers. The data include transaction-level information for all bank and credit card accounts at partner financial institutions. This includes transactions from bank accounts (e.g., checking and savings deposits, ACH transfers, bill payments, checks, ATM withdrawals) and payment cards (credit and debit card purchases). Each transaction record comes with rich metadata, notably, date, amount, and the merchant or payee. Every entry is also tagged with a detailed category based on the transaction description. These categories cover both income and spending types, such as "Salary" (for paycheck deposits), "Groceries," "ATM Withdrawal," "Mortgage Payment," and "Medical Spending." In essence, the data provide a near-comprehensive view of cash inflows and outflows in household bank and card accounts.

For the majority of purchases, the dataset records the merchant's location down to the 5-digit ZIP code level. This allows us to see where spending occurs and while the data do not directly include the user's home address or state, we impute an account owner's 5-digit ZIP code based on where each account holder most frequently conducts in-person purchases.

We use the detailed transaction-level data to construct granular measures of income and consumption. We then define (overall) "income" as all inflows into accounts that are labeled as income by the data aggregator's categorization. The most common example is direct payroll deposits marked in the "Salary" or "Paycheck" categories. We also include as income other deposits that occur on a recurrent basis with exactly the same dollar amount. We define (overall) "consumption" similarly by aggregating all outflows from the accounts to relevant categories of payments. We exclude transfers between accounts to avoid over-counting spending.

Analysis Sample. With these data, we construct a sample of job movers who earn stable income from a single employer for six months leading up to a job transition, and then, within two months, begin to earn stable income from a second employer for a new six-month period. In Table 1 we provide details about our monthly panel of job movers. In our analysis below, we will aggregate the panel to a quarterly level to improve precision. Overall, our analysis sample of movers includes 267,527 unique individuals, with 2,837,497 person-month observations, and 20,097 unique workplaces.

***Table 1:* Summary Statistics of Analysis Sample**

	Mean	S.D.
Log Income	8.4504	0.8683
Income	7,261	11,730
Log Expenditure	8.3989	0.9392
Expenditure	7,123	12,225
Year of Observation	2016.05	2.31
Number of Unique Individuals	267,527	
Number of Person-Month Observations	2,837,497	
Number of Month-Level Observations Per Person	10.6	
Number of Employers	20,097	

Notes: This table provides baseline statistics for our analysis sample of job movers.

3. Analytical Framework

We employ the benchmark two-way fixed effects job ladder model, which includes household fixed effects and firm premia in the determination of workers' economic outcomes. There are now several well-developed approaches that one could take in the analysis of job mobility. Our model choice is for practical reasons so that we can achieve comparison of earnings premia and consumption premia using the workhorse model as a relatable benchmark. Our underlying statistical model would accordingly follow the structure set forth by Abowd, Kramarz, and Margolis (1999), where identification and estimation would follow the movers design implementation developed by Finkelstein, Gentzkow, and Williams (2016). We will employ bias corrections of Kline, Saggio, and Sølvsten (2020) when studying second moments, and we will employ Empirical Bayes shrinkage (following Walters 2024) when studying distributional characteristics of firm effects.

3.1. Setup

Consider worker i who works in firm j in year t. We let $J(i,t)$ index the employer of worker i in year t, and let $y_{i,t}$ denote the outcome of interest. We let the generating process of economic outcomes follow the structure:

$$y_{i,t} = \alpha_i + x_{i,t}\beta + \theta_j + \varepsilon_{i,t}, \tag{1}$$

where α_i is a person fixed effect, $x_{i,t}$ is a set of calendar month-year indicators, and $\varepsilon_{i,t}$ is a time-varying error component. θ_j is the contribution of employer $j = J(i,t)$ to an employee's outcome, that is, the "firm premium."

We can define cohorts of job switchers, indexed by c, according to their origin firm $o(c)$, destination firm $d(c)$, and move period $m(c)$. We let Ω_c denote the set of movers in cohort c. We focus on

movers who are in their origin firm for periods $t < m(c)$, in their destination firm for periods $t > m(c)$, and may be in either their origin firm or their destination firm in period $t = m(c)$. To focus on event time when studying movers, we let $r(c,t) \equiv t - m(c)$ index periods relative to a cohort's move period.

With this parameterization, we can rewrite equation (1) for workers from mover cohort c as:

$$y_{i,r} = \sigma_i + \gamma_r + x_{i,r}\beta + [\theta_{d(c)} - \theta_{o(c)}]\, \mathbb{I}\{r(i,t) > 0\} + \varepsilon_{i,r},$$

where $\sigma_i \equiv \alpha_i + \theta_{o(c)}$, and $\mathbb{I}\{\cdot\}$ is the indicator function. We explicitly allow for differential time patterns for movers, as captured by γ_r, above and beyond the general time trends, as captured by calendar month-year in $x_{i,t}$. It relaxes the "exogenous mobility" requirement for identification by allowing movers' dynamics in economic outcomes to systematically differ from general time patterns in the labor market. But it also highlights a restriction, which will be our identifying assumption as we describe below: it implies common relative time patterns across mover cohorts (as γ_r is not indexed by c). We can further generalize the model to allow for dynamics in firm effects, $\theta_{j,r}$, as we do in our empirical implementation.

3.2. Treatment Effect and Identification

We define the cohort-level treatment effect as:

$$\tau_c \equiv \theta_{d(c)} - \theta_{o(c)}.$$

We denote the difference in average outcomes of workers across a pair of employers by

$$\Delta_c \equiv E_i\big[y_{i,t} | J\big(i, m(c)\big) = d(c)\,\big] - E_i\big[y_{i,t} | J\big(i, m(c)\big) = o(c)\,\big],$$

and we scale to define the share of differentials causally attributable to firms as:

$$\lambda_c = \tfrac{\tau_c}{\Delta_c}.$$

This implies that the following holds for workers from mover cohort c:

$$y_{i,r} = \sigma_i + \gamma_r + x_{i,r}\beta + \lambda_{c,r}\mathbb{I}\{r(i,t) > 0\}\Delta_c + \varepsilon_{i,r},$$

where we allow for dynamic treatment effects as captured by the index r in $\lambda_{c,r}$. We define our estimand to be:

$$T_{c',r} \equiv E_{i\in\Omega_{c'}}\big[y_{i,r} - y_{i,-1}\big] - E_{i\in\Omega_{c''}}\big[y_{i,r} - y_{i,-1}\big],$$

where c' is a "treatment" cohort and c'' is a "control" cohort with zero treatment intensity, i.e., $\Delta_c = 0$. In practice, we take the following estimating equation to the data:

$$y_{i,r} = \sigma_i + \gamma_c + x_{i,r}\beta + \lambda_r\mathbb{I}\{r(i,t) > 0\}\Delta_c + \varepsilon_{i,r}, \tag{2}$$

which is a log-linear aggregation of all the pairwise difference-in-differences comparisons (in the form of $T_{c',r}$). Our parameter of interest is λ_r that characterizes the dynamics in average passthrough rates.

We can map the estimand to the model parameters as follows:

$$T_{c',r} = \underbrace{\lambda_{c',r}\Delta_{c'}}_{\substack{\textbf{Treatment}\\ \textbf{Effect}}} + \Big\{E_{i\in\Omega_{c',g}}\big[\varepsilon_{i,r}\big] - E_{i\in\Omega_{c'',g}}\big[\varepsilon_{i,r}\big]\Big\}.$$

This immediately leads to our identifying assumption.

Assumption 1: "Conditional Parallel Trends." *For mover cohort pairs c' and c'', the following holds*

$$E_{i\in\Omega_{c',g}}[\varepsilon_{i,r}] - E_{i\in\Omega_{c'',g}}[\varepsilon_{i,r}] = 0.$$

Assumption 1 requires that conditional on our set of fixed effects—an individual fixed effect, a firm fixed effect, calendar month-year indicators, and indicators for time relative to the job move in high frequency—outcomes across mover cohorts (c) would run in parallel in the absence of their differential treatment intensity (Δ_c). That is, taking into account a common relative time factor, γ_r, there remain no systematic differences across mover cohorts. Importantly, it allows for systematic selection in terms of patterns among those who become movers, so that moves need not be "exogenous." Comparability is required across counterfactuals of different mover cohorts instead.

As is natural in these settings, our identifying assumption delivers the testable implication of parallel trends in the periods prior to the move. In our setting, the test takes the following form:

Test 1: "Parallel Pre-Trends." *Under Assumption 1, $\lambda_{c',r} = 0$ for all $r < 0$, which in turn implies that $T_{c',r} = 0$ for all $r < 0$.*

3.3. Decomposition

There are multiple channels underlying the firm consumption premia, θ_j. The effect of firms on individual consumption could be driven by variation in earnings (the most obvious channel), variation across locations (that would take into account local markets, prices, purchasing power, and location-level norms), and a group-level unexplained residual that incorporates all types of interactions at the workplace levels, including firm policies and peer effects.

To dissect the degree to which the different channels contribute to the overall treatment effect that we identify, we will estimate equation (2) where we incrementally add appropriate sets of controls in $x_{i,r}$. These include flexible income controls (a quadratic in household income interacted with all relative time indicators) and geographic location fixed effects.

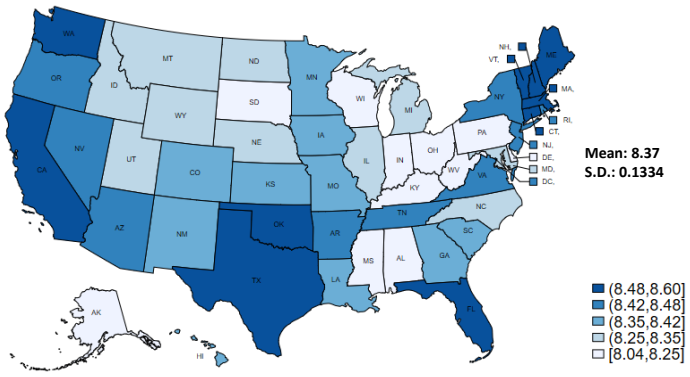
4. Preliminary Evidence

We provide preliminary findings from two sets of investigations. First, we study the role of firms in consumption variation. Specifically, we identify the share of the nationwide variation in consumption across jobs that is attributable to firms, directly speaking to the origins of consumption inequality in the US. Second, we study the roots of these firm effects in terms of the degree to which they capture income effects, location effects, and a residual group-level effect. This will allow us to draw informative conclusions about household models of consumption and the relative forces that govern American households' consumption decisions.

4.1. Baseline Statistics

Firm-Level Spatial Dispersion of Consumption. Figure 1 plots the spatial distribution of individuals' consumption at the workplace level. We calculate the average log expenditure of individuals in a given firm, and we then collapse the average firm-level means across all firms within a given state. This allows us to focus on comparisons at the firm level, whose causal effects on consumption we are after.

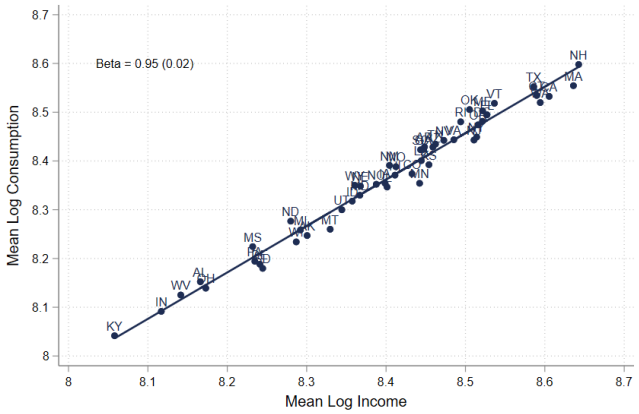
Figure 1: ***Firm-Level Dispersion of Consumption by State***



Notes: This figure plots averages for log consumption within firm by state. We calculate the average log expenditure of individuals in a given firm, and we then collapse the average firm-level means across all firms within a given state.

The figure shows meaningful spatial variation across states. The standard deviation is 13.34 log points, and with a nationwide mean of 8.37, state averages of log consumption range from 8.04 at the very bottom (KY) to 8.60 at the very top (MA). Figure 2 further explores the state-level correlation between log consumption and log income. The relationship between the two outcomes is very close to linear, i.e., a constant elasticity. The slope of 0.95 implies that, in equilibrium at the state level, 95% of additional income is consumed and 5% is saved at the aggregate.

***Figure 2:* State-Level Relationship Between Consumption and Income**



Notes: This figure plots the state-level relationship between log consumption and log income.

Job Transition Intensity. As an initial step in the analysis of moves, we characterize job moves in our sample in terms of the distribution of treatment intensity, Δ_c. Throughout the analysis, we compare our main findings for consumption with additional findings for income as a benchmark. Table 1 displays statistics for Δ_c when the outcome is log consumption in column 1 and when the outcome is log income in column 2. We find that most transitions in our analysis are up the job ladder, that is, to jobs with higher earnings and employee consumption levels. From an analytical standpoint, we also see that we have meaningful variation (of 57 log points in consumption and 56 log points in income) with wide coverage of negative and positive treatment intensity values.

***Table 2:* Job Transition Treatment Intensity (Δ_c)**

	Log Consumption (1)	Log Income (2)
Mean	.1097	.1255
Standard Deviation	.5683	.5583
Percentiles:		
1%	-1.1245	-1.1392
25%	-.1933	-.1688
50%	.0452	.0828
75%	.3317	.3567
99%	2.0275	1.9741

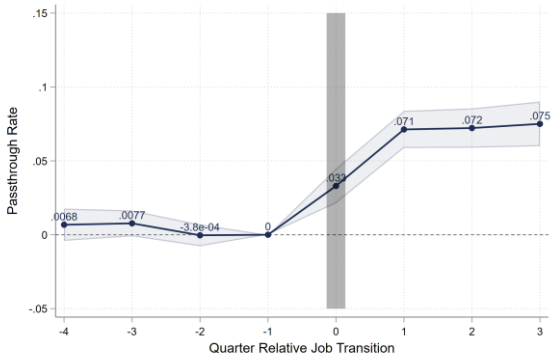
Notes: This table characterizes the distribution of the "treatment intensity" of job transitions, Δ_c, in our analysis sample. Column 1 provides statistics for log consumption, and column 2 provides statistics for long income.

4.2. Firm Effects on Household Consumption

Dynamic Effects on Consumption. We now turn to the main analysis and estimate equation (2) where the outcome is the log of household expenditure. The results are provided in Figure 3. In this and the subsequent figure, the *x*-axis represents time relative to the job move, and on the *y*-axis we plot estimates for the passthrough rates over time, λ_r, along with their 95-percent confidence intervals. Period 0 is a transition period during which employees are only partially exposed to the treatment of the move, and we accordingly grey it out in the figures for visual convenience. Period 1 is the first period in which outcomes reflect the treatment in its full capacity.

We can first see that there are no systematically different trends across cohorts with differential treatment intensity in the pre-period, in support of the identifying assumption. There is then a clear jump in passthroughs following the move, which displays persistence over our analysis horizon. As laid out in Section 3.2 under the identifying assumption, the findings imply meaningful firm effects in determining household consumption choices. They imply that over 7 percent of the *entire* cross-job consumption differentials are causally attributable to firms. To clarify the economic magnitude of this amount, we note that the variation denominator includes all variation in household consumption, such as that coming from heterogeneity in skills, preferences, education, age, wealth, gender, race, and life circumstances, e.g., fertility and geographic location.

***Figure 3:* Dynamic Effects of Job Mobility on Consumption**

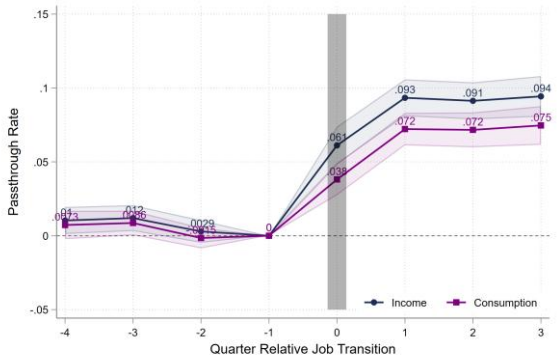


Notes: This figure studies how household consumption evolves following job transitions. The *x*-axis represents time relative to the job move, and on the *y*-axis we plot estimates for the passthrough rates over time, λ_r from equation (2), along with their 95-percent confidence intervals.

Comparison to Income. In Figure 4 we provide two specifications of equation (2) in which the treatment intensity of the move, Δ_c, is in terms of log income instead of log expenditure. The two specifications differ by the outcome variable, $y_{i,r}$: in the first specification the outcome is log consumption, and in the second specification the outcome is log income. When the outcome is log consumption, the

passthrough coefficients describe the reduced-form effect of how a move to a job with higher average income translates to increases in individual consumption. The estimates imply that a move to a job with a $1 higher average income (among all workers) increases a moving household's consumption by 7 cent.

<u>Figure 4:</u> Dynamic Effects of Mobility across Jobs with Differential Average Income



Notes: This figure studies how household economic outcomes evolve following job transitions. The *x*-axis represents time relative to the job move, and on the *y*-axis we plot estimates for the passthrough rates over time, λ_r from equation (2), along with their 95-percent confidence intervals. In these estimations, Δ_c is in log income and the outcomes studied are log income and log consumption.

When the outcome is log income instead, the passthrough rate captures the share of the variation in log income that is causally determined by firms. We find that firms explain a somewhat higher share of the variation in income (the blue lines in Figure 4) than the variation in consumption (in Figure 3). For example, in the first quarter following the move, the relative differences are on the order of 31% (=9.3/7.1). This implies that, when analyzing income, we inflate every $1 of variation in economic outcomes that is ascribed to firms by 31% as compared to when we study consumption.

4.3. Decomposition of Consumption Effects

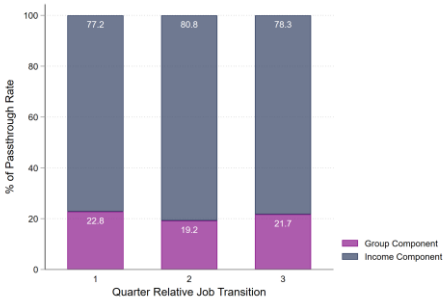
In our last set of preliminary findings, we investigate the degree to which the average passthrough rate in firms' effects on household consumption is driven by changes in income, changes in location, and a residual of group-level interactions. We specifically assess the degree to which each component explains a $1 change in consumption that is due to a job switch.

Figure 5 displays a dynamic decomposition of the effect into income changes versus non-income components. We find that income changes explain about 80% of changes in household consumption, so that a significant share of 20% still remains unexplained.

Accordingly, we further decompose the non-income component into a location component and a residual group-level component. In this draft, location is at a state level, and we will later create a finer-level geographic match to study the degree of locality at which consumption norms might take effect. We

increase precision by aggregating over periods, where we include quarters -4 to -1 in the pre-move periods and quarters 1 to 3 in the post-move periods. The results are provided in Table 3. We find that basically none of the non-income components in the effects of jobs on consumption is due to differences in the household's state.

***Figure 5:** Dynamic Decomposition of Consumption Effects*



Notes: This figure displays in a dynamic decomposition of the effect of job mobility on consumption into income changes versus non-income components.

***Table 3:** Decomposition of Average Consumption Effects*

	Baseline Specification (1)	Income Controls (2)	Location Controls (3)
Passthrough Rate	.06758	.01402	.01395
	(.00728)	(.00377)	(.00378)
% Out of Total	100	20.75	20.64
Income Controls		X	X
State Fixed Effects			X
Number of Observations	2,003,667	2,003,667	2,003,667

Notes: This table provides a decomposition of the effects of firms on household consumption.

Our findings so far can be summarized as follow. Our quasi-experimental estimates of the role of firms in determining American households' consumption patterns imply that 7 percent of the entire variation in consumption nationwide is casually attributable to firms. Decomposing these consumption effects, we find that for a $1 increase in consumption from firm effects, 80 percent is due to income increases, whereas 20 percent are driven by workplace-level factors, such as peer effects, that cannot be explained by geographic location per se in terms of state.

References

Abowd, J., Kramarz, F., and Margolis, D. (1999). "High Wage Workers and High Wage Firms." *Econometrica*, 67(2): 251-333.

Attanasio, Orazio, and Luigi Pistaferri (2014). "Consumption Inequality Over The Last Half Century: Some Evidence Using The New PSID Consumption Measure." *American Economic Review*, 104(5): 122-126.

Card, David, Ana Rute Cardoso, Joerg Heining, and Patrick Kline (2018). "Firms and Labor Market Inequality: Evidence and Some Theory." *Journal of Labor Economics*, 36(S1): S13-S70.

Finkelstein, A., Gentzkow, M. and Williams, H. (2016). "Sources of Geographic Variation in Health Care: Evidence from Patient Migration." *Quarterly Journal of Economics*, 131(4): pp.1681-1726.

Kline, Patrick, Raffaele Saggio, and Mikkel Sølvsten (2020). "Leave-out Estimation of Variance Components." *Econometrica*, 88(5): 1859-1898.

Meyer, Bruce D., and James X. Sullivan (2023). "Consumption and Income Inequality in the United States since the 1960s." *Journal of Political Economy*, 131(2): 247-284.

Song, J., Price, D.J., Guvenen, F., Bloom, N., and Von Wachter, T. (2019). "Firming Up Inequality." *Quarterly Journal of Economics*, 134(1): pp.1-50.

Walters, Christopher (2024). "Empirical Bayes Methods in Labor Economics." In Handbook of Labor Economics, vol. 5, pp. 183-260. Elsevier, 2024.