Location, Firms, and Consumption Inequality in the US

Itzik Fadlon, University of California, San Diego and NBER

Ezra Karger, Federal Reserve Bank of Chicago

Shanthi P. Ramnath, Federal Reserve Bank of Chicago

September 2025

Preliminary Draft

Abstract: We study the role of location and firms in determining American households' consumption behavior and its nationwide dispersion using transaction-level bank data. We use household migrations across locations in a movers' design to study how differences in average outcomes across origin and destination locations pass through into movers' own outcomes. We find that cross-state job moves exhibit a passthrough rate of 26 percentage point, so that about a quarter of the nationwide consumption differentials can be causally attributable to the "location" composite. Decomposing this total location effect, we find that 14 pp are attributable to an individual's firm (as captured by job moves within ZIP codes) and 12 pp are attributable to an average effect of location itself. Whereas the entire firm effect can be explained by changes in income, 40% of the average location effect remain after income is accounted for. This "excess consumption" component implies that household preferences exhibit "location dependence," inclusive of social incentives (such as norms) for which we find some empirical support. Our findings have implications for placed-based policies, predicting that, on average, every $1 increase in location-level income (e.g., via improved employment and earnings opportunities) will produce an 82 cent increase in local consumption.

1. Introduction

Economic inequality is a central topic of research and policy debate in the US. A rich and active body of work aims to identify the root causes of income and earnings dispersion, with an emphasis on the role of locations and the employment opportunities they provide (e.g., Chetty and Hendren 2018a,b, Card et al. 2018, Song et al. 2019, Card et al. 2023, Fadlon et al. 2025). Recent work, however, underscores the empirical importance of the concept that analyzing consumption, a utility input, more accurately identifies well-being than analyzing income, a budget input (e.g., 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, they account for the insurance protection of the tax-and-transfer system, and they encompass crowd-out of unobservable consumption-inducing actions (such as support from family members).[1]

In this paper, we connect these two lines of work and study the role of location and firms in determining American households' consumption behavior and its nationwide dispersion. We use transaction-level bank data spanning bank accounts and credit cards of millions of individual US households. The richness of the data allows us to construct a high-frequency panel of employer-employee matches with detailed information on income flows, expenditures, and geographic location. For identification, we use quasi-experimental variation that leverages migration across locations, following the pioneering work by Finkelstein, Gentzkow, and Williams (2016).[2]

Specifically, we treat jobs as the underlying micro-environment of an individual (in line with Card et al. 2023), and we define the difference in average employee expenditure across jobs as the "treatment intensity" in a dose-response design. We estimate the overall role of location using event-study designs that are based on how the treatment intensity passes through into the mover's own expenditure following cross-state migrations. The identifying assumption is that movers with differential treatment intensity would exhibit similar trends, allowing for systematic selection across movers and non-movers. As developed in Finkelstein, Gentzkow, and Williams (2016), under this identifying assumption, the passthrough rate translates to the causal role of the "units" across which the individual moves (e.g., states), in terms of the share of the observed differentials that they explain. A useful interpretation comes in the form of counterfactuals: the passthrough measures the share that the gap in consumption between origin and destination units would fall if households were randomly reallocated between them. We further decompose the total cross-state effect into an average location effect and a (within-location) firm effect by leveraging

[1] In practical terms, analyzing consumption also reflects all available sources of household wealth, information that is typically not comprehensively collected in administrative datasets.

[2] In related earlier work in the context of consumption, Bronnenberg, Dubé, and Gentzkow (2012) study the long-run evolution of brand preferences using the Nielsen Homescan panel by analyzing how current purchases depend on both where consumers live currently and where they lived in the past.

comparisons across classes of moves. That is, we augment the analysis of cross-state migrations (that are inclusive of the total effect) with job moves that are within one's origin 5-digit ZIP code (that are exclusive of a location effect). Importantly, the difference-in-differences style decomposition allows for systematic selection across these classes of moves (which we show to be empirically present).

Our analysis provides the following set of new findings. First, we find that cross-state moves exhibit a 26 percentage point (pp) passthrough rate. With empirical support for the identifying assumption (via comparability of pre-trends), it implies that about 26% of the nationwide consumption differentials can be causally attributable to states. Decomposing the total effect based on comparisons across classes of moves, we find that 14 pp are attributable to an individual's job, and that 12 pp are attributable to location itself.

Second, we characterize the channels that drive these effects. Through the lens of a household budget constraint, the effects of moves can come from changes in consumption technology (via local market prices that convert income to consumption), changes in income (that increase the purchasing capacity), and changes in consumption choices *conditional* on prices and income. The latter channel implies changes in the "upper level" of the household's optimization problem, which translates to households' preferences being location-dependent (e.g., via norms).

For within-location moves, we find that the entire passthrough is explained by changes in income upon the job switch. Not surprisingly, it implies that the firm effect portion can be fully explained by earnings. For cross-state moves, we find that relative prices have no explanatory role; the coefficients show miniscule changes when we control for Regional Price Parities (RPPs).[3] We then show that, out of the 12 pp explained by location itself, about 40% cannot be accounted for by income. This implies a meaningful "excess consumption" channel in the effects of location on consumption with location-dependence in household preferences over consumption choices. We explore categories of consumption to learn more about this channel, which can incorporate social incentives and considerations such as norms. Consistent with this view, we find a large passthrough for spending on "traveling" with null effects on daily necessities as captured by spending on "groceries."

Finally, we analyze the welfare and policy implications of our findings by characterizing the connection between income changes and consumption changes in household migration. We define the "premia MPC" to be the ratio between the location's consumption effect to the location's income effect,

[3] Regional price parities (RPPs), produced by the US Bureau of Economic Analysis, measure the differences in price levels across states and metropolitan areas for a given year and are expressed as a percentage of the overall national price level. See details at: https://www.bea.gov/data/prices-inflation/regional-price-parities-state-and-metro-area.

which allows us to investigate two related points. The first point is the importance of measuring consumption, as compared to measuring earnings or total income, in assessing the implications of spatial economic dispersion on household welfare. The degree to which the premia MPC deviates from the benchmark of one quantifies the empirical relevance of this distinction. We find a premia MPC of 0.82 nationwide, implying we would be off by 18 cents on a dollar in income-based assessments of household welfare.

The second point relates to the implications of our findings to the efficiency and equity of placed-based policies, which are regaining significant attention in academic research and policy discussions. Common place-based policies can be mapped to changes in a location income premia, as they have shown to induce increases in location-level earnings or income (e.g., Busso et al. 2013). Thus, we can use the premia MPC—that maps location income premia to location consumption premia—to evaluate the implications of place-based policies, both in terms of normative effects of changes to household welfare (through a utility input) and in terms of positive implications for market-level demand (through the location-level policy). Analyzing moves across varying levels of geography, we find that the efficacy of place-based policies in terms of improving welfare is similar across the "reach" of the policy: the premia MPC is 0.82 in cross-state moves, 0.82 in cross-county moves, and 0.83 in moves across 5-digit ZIP codes. With respect to equity, we find that location premia tend to benefit lower-income families to a higher degree (or, alternatively, that individual-specific components, such as education, are a stronger driver of higher-income families' economic outcomes). This opens up further opportunities for redistributive gains through place-based policies, that is, not only across locations but also within.

Our work builds on and contributes to three strands of the literature. The first is the rich work on inequality in economic outcomes in the US (particularly, in income, earnings, and wealth) and its underling determinants (e.g., Saez and Zucman 2016, Chetty et al. 2017, Chetty et al. 2020, Hubmer et al. 2021). The second is the active and growing body of work highlighting the central role of "place" in shaping household outcomes, including education, earnings, health, and mobility (e.g., Chetty and Hendren 2018a,b, Finkelstein et al. 2016, 2021, Card et al. 2023). The third related line of work includes a series of influential papers on the role of employers in generating earnings inequality (e.g., Song et al. 2019; Kline et al. 2022; and summarized in Card et al. 2018). We contribute to these important strands of the literature by providing, to the best of our knowledge, the first quasi-experimental analysis of the effects of location and firms in driving household consumption in the US, using administrative transaction-level data on a nationwide scale.

2. Data and Sample

Data. We utilize 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 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 of a household's 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.

Sample. We construct a sample of job movers as follows. First, for precision in the analysis of firms as the underlying unit, we keep employee-month observations at employers with at least 20 employees across all months they are in the sample. For creating a meaningful panel, we also exclude individuals with less than 3 observations. We identify individuals who earn stable income from a single employer for at least six months leading up to a job transition, and who then earn stable income for at least six month from a different employer, allowing a gap period of employment of no more than 2 months in between. We define stable income as deposits that total to at least $1,250 on a monthly basis. This is the monthly equivalent of an annual income of at least $15,000, which corresponds to the approximate level of earnings from full-time work at a minimum wage of $7.25 per hour. In Table 1 we provide details on the monthly panel of job movers. In our analysis below, we will aggregate the panel to a quarterly level to improve precision. Overall, we identify 130,349 unique movers with 1,449,029 person-month observations.

Table 1: Summary Statistics of the Sample of Movers

	Mean	S.D.
Log Expenditure	8.35	0.88
Expenditure	6,466.29	10,674.83
Log Income	8.38	0.84
Income	6,586.87	10,180.22
Year of Observation	2,016.22	2.30
Number of Unique Individuals	130,349	
Number of Person-Month Observations	1,449,029	
Number of Month-Level Observations Per Person	11.12	
Number of Employers	6,224	

Notes: This table provides baseline statistics for the sample of all job moves identified in the data.

3. Analytical Framework

Our research design relies on event studies that compare movers across differentially "intense" units and analyze how unit-level intensity passes through into movers' own outcomes. The estimating equation will take the following straightforward "dose-response" event-study structure:

$$(1) \qquad y_{i,r} = \rho_i + \sum_{r\neq -1; r=-4}^{r=4} \gamma_r \times I_r + \sum_{r\neq -1; r=-4}^{r=4} \lambda_r \times \Delta_{\sigma(i)} \times I_r + \varepsilon_{i,t},$$

where r is time relative to the move (in quarters), I_r are relative time indicators, and ρ_i are individual-level fixed effects. The key right-hand side variable is $\Delta_{\sigma(i)}$, the "treatment intensity" variable. It measures, for each mover i, the observed gap in the average of a given outcome across i's origin unit and i's destination unit. We will characterize the "cohort" of a mover i, denoted by $\sigma(i)$, as the combination of 'origin unit,' 'destination unit,' and 'time of move,' so that $\Delta_{\sigma(i)}$ varies at the cohort level. λ_r captures the parameter vector of interest: the degree to which the observed gap, $\Delta_{\sigma(i)}$, "passes through" into the mover's outcome, $y_{i,r}$. We now describe the mapping between these reduced-form parameters in equation (1) to location and firm premia in a canonical two-way fixed effects framework. This establishes the interpretation of λ_r as the average share of the observed differentials that can be causally attributed to the "location" composite, as originally developed by Finkelstein, Gentzkow, and Williams (2016) and applied to our setting.

Consumption Determination Process. Our starting point is the canonical two-way fixed effects model (Abowd, Kramarz, and Margolis 1999, Finkelstein, Gentzkow, and Williams 2016). We treat jobs as the underlying unit for the "micro environment" of an individual and characterize a reduced-form statistical equation for the evolution of household outcomes. We consider individual i who works in firm j in year t,

and we use $J(i,t)$ to index the employer of individual i in year t. We let the generating process of economic outcomes, $y_{i,t}$, take the following structure:

$$y_{i,t} = \alpha_i + \gamma_t + x_{i,t}\beta + \theta_j + \varepsilon_{i,t}, \tag{2}$$

where α_i is a person fixed effect, γ_t's are calendar period fixed effects, $x_{i,t}$ are potential time-varying controls, 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." This is the total reduced-form effect on those employed in firm j, *inclusive* of the average location effect of where the job is located.

Total Effect of Moves. We define cohorts of movers, indexed by $\sigma(i)$, according to their origin firm $o(\sigma)$, destination firm $d(\sigma)$, and move period $m(\sigma)$. We consider movers who are in their origin firm for periods $t < m(\sigma)$, in their destination firm for periods $t > m(\sigma)$, and may be in either their origin firm or their destination firm in period $t = m(\sigma)$. To focus on event time when studying movers, we let $r(\sigma(i),t) \equiv t - m(\sigma(i))$ index time relative to the period of the move.

With this parameterization, we can rewrite equation (2) for movers in cohort σ as:

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

where $\rho_i \equiv \alpha_i + \theta_{o(\sigma)}$, and $\mathbb{I}\{\cdot\}$ is the indicator function. As we describe below, assumptions about underlying time patterns—the cohort-specific time trend, $\gamma_{\sigma,r}$—are at the core of identification. We can generalize the model to allow for dynamics in job effects, which is done in the empirical analysis.

Treatment Effect and Identification. We define the treatment effect of the move as:

$$\tau_\sigma \equiv \theta_{d(\sigma)} - \theta_{o(\sigma)},$$

and we decompose the total reduced-form firm premia (θ_j) to a geographic average and a firm's deviation from that average:

$$\theta_j \equiv \mu_{l(j)} + \tilde{\theta}_j,$$

where $l(j)$ indexes the firm's geographic location, and $\mu_{l(j)} \equiv E_j\left[\theta_j | l(j) = l\right]$ (the average for all firms in a given location). This leads to the following decomposition of the full treatment effect:

$$\tau_\sigma = \underbrace{\left[\mu_{l(d)} - \mu_{l(o)}\right]}_{\substack{\textbf{Location} \\ \textbf{Effect}} \equiv \tau_\sigma^L} + \underbrace{\left[\tilde{\theta}_d - \tilde{\theta}_o\right]}_{\substack{\textbf{Firm} \\ \textbf{Effect}} \equiv \tau_\sigma^F} = \tau_\sigma^L + \tau_\sigma^F.$$

Allowing for dynamics in the treatment effect, it follows that we can rewrite the transition equation as:

$$\begin{aligned} y_{i,r} &= \rho_i + \gamma_{\sigma,r} + x_{i,r}\beta + \tau_{\sigma,r}\, \mathbb{I}\{r(\sigma(i),t) > 0\} + \varepsilon_{i,r} \\ &= \rho_i + \gamma_{\sigma,r} + x_{i,r}\beta + \left[\tau_{\sigma,r}^L + \tau_{\sigma,r}^F\right] \mathbb{I}\{r(\sigma(i),t) > 0\} + \varepsilon_{i,r}. \end{aligned}$$

Estimating equation. The empirical implementation of this equation differs across movers' design applications and AKM applications. The latter aims to identify premia for each separate unit (i.e., a firm)

and study the unit-level premia distribution, e.g., for explaining second moments such as the variance of wage premia and covariances across firm and person effects for identifying sorting patterns. Instead, movers' design applications such as ours, focus on identifying the overall share out of the observed dispersion in economic outcomes that can be causally explained by the unit of interest (such as the role of states in explaining the geographic variation in healthcare as in the original Finkelstein et al. 2016 paper). To do so, the movers' design transforms the outcome determination process from its firm-specific formulation to its representation as shares out of the observed gaps as we now develop.

We denote the difference in average outcomes of all workers across mover i's destination job and origin job by:

$$\Delta_{\sigma(i)} \equiv E_k\left[y_{k,t} | J(k,t) = d(\sigma(i))\right] - E_k\left[y_{k,t} | J(k,t) = o(\sigma(i))\right],$$

which we term the "treatment intensity." We then scale to define the share of differentials causally attributable to firms, so that: $\lambda_r \equiv \frac{\tau_{\sigma,r}}{\Delta_\sigma}$, $\lambda_r^L \equiv \frac{\tau_{\sigma,r}^L}{\Delta_\sigma}$, and $\lambda_r^F \equiv \frac{\tau_\sigma^F}{\Delta_\sigma}$, with $\lambda_r = \lambda_r^L + \lambda_r^F$.

Taken together, this implies for individuals from mover cohort σ that:

$$y_{i,r} = \rho_i + \gamma_{c,r} + x_{i,r}\beta + \lambda_r \, \mathbb{I}\{r(\sigma(i),t) > 0\}\Delta_{\sigma(i)} + \varepsilon_{i,r}, \tag{3}$$

This is our estimating event-study equation. It aggregates across cohorts σ within a class of moves as the underlying level of variation. It takes a log-linear (or, a constant-elasticity) form with an overall passthrough parameter across cohorts. A key appeal of the movers' design formulation is that it produces easily-interpretable parameters: λ_r's capture the average share of economic dispersion attributable to the analysis unit of interest out of the total differentials observed in the data. Equation (3) is equivalent to equation (1) above, which is written in its more familiar dynamic event-study formulation.

Identification. We define classes of cohorts based on their geographic-relocation distance: job moves across states (S), and job moves within 5-digit ZIP codes (Z). We define the estimand for the total effect (inclusive of both location and firm effects) as:

$$T_{\sigma',r} \equiv E_{i\in\Omega_{\sigma'}}\left[y_{i,r} - y_{i,-1}\right] - E_{i\in\Omega_{\sigma''}}\left[y_{i,r} - y_{i,-1}\right],$$

where σ' is a "treatment" cohort and σ'' is a counterfactual "control" cohort with zero treatment intensity $\Delta_\sigma = 0$, with both $\sigma', \sigma'' \in S$ (i.e., in the cross-state class of moves). Our estimating equation in (3) is a log-linear aggregation of all such pairwise difference-in-differences comparisons within a class of moves. It exploits the full range of the sample variation in Δ_σ, in lieu of "clean" hypothetical counterfactuals (such as cohort σ''), using a functional form.

Our formulation implies the following mapping between the estimand of within-class moves to the parameters in equation (3):

$$T_{\sigma',r} = \underbrace{\lambda_r \Delta_{\sigma'}}_{\substack{\textbf{Scaled Treatment} \\ \textbf{Effect}}} + \underbrace{\{\gamma_{\sigma',r} - \gamma_{\sigma'',r}\}}_{\textbf{Time Effect}}.$$

It immediately leads to the following identifying assumption:

Assumption 1: [Conditional Parallel Time Trends] For mover cohort pairs σ' and σ'', the following holds: $\gamma_{\sigma',r} - \gamma_{\sigma'',r} = 0$.

This assumption requires that, conditional on fixed effects, outcomes would run parallel in the absence of the treatment across mover cohorts within a class of moves. That is, taking into account a common relative time factor, γ_r, there remain no mean systematic differences. Importantly, it still allows for systematic selection across movers and non-movers. It is simply implemented in practice by replacing $\gamma_{\sigma,r} = \gamma_r$ in equation (3) within a class of moves. This assumption delivers the usual testable implication of parallel pre-trends for non-selection across within-class move cohorts: $\lambda_{\sigma,r} = 0$ for $r < 0$.

We then decompose the total treatment effect into an average location effect and a within-location firm effect. With our structure, an additional difference identifies the average location effect:

$$T_{\sigma' \in S,r} - T_{\sigma' \in Z,r} = \tau_c - \tau_F = \tau_L,$$

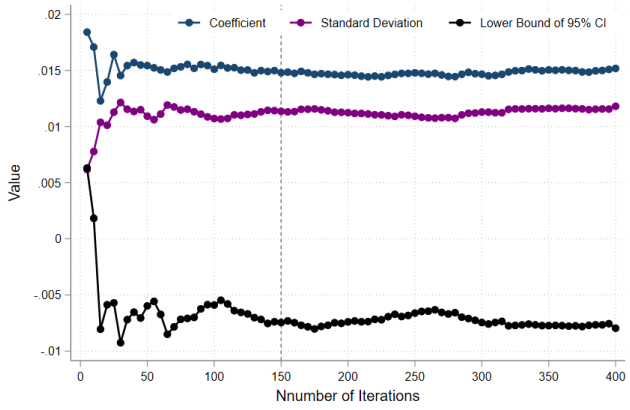
when each estimator is estimated separately for each class of moves. This setup can accommodate an "extra constant" within each of the differences on the left-hand side. Specifically, it allows for systematic selection across classes of moves, $g \in \{S, Z\}$. To see this, we can explicitly decompose the error term into a conditional sorting component and a residual:

$$\varepsilon_{g,i,r} = \kappa_{g,i} \times \mathbb{I}\{r(\sigma(i), t) > 0\} + u_{g,i,r},$$

where $\kappa_{g,i}$ comes from projecting $\varepsilon_{g,i,r}$ onto a constant post-move, so that $E_i(u_{g,i,r}) = 0$. This allows for systematic differences across classes of moves, such that $E_i(\varepsilon_{g,i,r}) = \kappa_g \times \mathbb{I}\{r(\sigma(i), t) > 0\}$ with $\kappa_g \equiv E_{i \in g}(\kappa_{g,i})$.

Allowing for such systematic differences is easily implemented by letting the estimating equation fully interact with an indicator for the class of moves. We later show that this "bias correction" by using separate estimations for each class of moves is needed (since the post-move baseline time patterns, as captured by $\gamma_{g,r}$ for $r > 0$, exhibit systematic differences across g). This is consistent with optimal mobility choices so that, for a given level of treatment intensity (Δ_σ) and the associated expected change in outcomes, you are willing to move further if the match-specific returns are higher ($\kappa_{g,i}$)

Figure 1: Estimation Statistics as a Function of Number of Bootstrap Iterations



Notes: This figure illustrates how the parameter estimates change as a function of the number of bootstrap iterations in the instrumenting stage of the estimation procedure. It provides statistics for the average post-move effect as captured by the average of λ_r for periods 1-3 in the estimation of equation (3). We estimate a version with income controls, where we expect the treatment effect to be lower, so that the number of iterations we choose arrives at stable estimates for specifications where it is hardest to detect an effect.

Estimation Procedure. We estimate equation (3) using a "split-sample" approach to address sampling error in Δ_σ. While the dataset is large, it is only partial to the population, providing noisy estimates for the "population" Δ_σ. We proceed in two steps. First, we randomly split the sample, calculate two vectors of Δ_σ, and use one to instrument for the other in an estimation of equation (3). Second, we loop this step for N times to avoid noise from one arbitrary split (capturing procedural noise from the need to instrument). Our estimates for the average treatment effect and its standard error are, respectively, $\hat{\lambda}_r = \frac{1}{N}\sum_{n=1}^{N}\hat{\lambda}(n)_r$ and $\hat{\sigma}_r = \frac{1}{N}\sum_{n=1}^{N}\hat{\sigma}(n)_r$, where $\hat{\lambda}(n)_r$ and $\hat{\sigma}(n)_r$ are the estimates from iteration n.

In practice, we choose $N = 150$ based on the investigation provided in Figure 1. We see that the in-sample noise stabilizes long before this threshold, where the noise is concentrated on lower numbers. We note that in this investigation we use the standard deviation of the estimated coefficients from the bootstrap iterations, which accounts only for the procedural noise. We do not use the standard-error estimates used for inference, which need to also account for the "higher-level" sampling error (out of the US population).

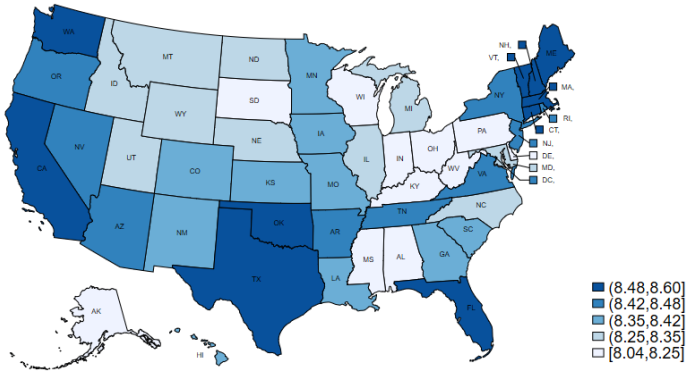
4. Empirical Evidence

4.1. Baseline Statistics

Figure 2 plots the spatial distribution of consumption. 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 follows the logic of jobs as the underlying "unit" and state as the main "location" level we consider. 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: Dispersion of Consumption across the US



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.

As an initial step in the analysis of moves, we characterize the distribution of treatment intensity in consumption, $\Delta_{\sigma(i)}$. We also provide the respective statistics for log income for comparisons. Table 2 displays statistics for averages, standard deviation, and percentiles among the sample of all moves, cross-state moves, and moves within 5-digit ZIP codes. We see that most transitions in our analysis involve "upward" mobility, that is, to jobs with higher earnings and employee consumption levels. From an analytical standpoint, we also see that we have meaningful variation (of 47 log points in consumption and 45 log points in income) with wide a coverage of negative and positive treatment intensity values.

Table 2: Treatment Intensity (Δ_σ)

	All Moves	Cross-State Moves	Moves within ZIP Code
Δ_σ in Log Expenditure:			
Mean	0.0742	0.0870	0.0668
SD	0.4460	0.4326	0.4630
1st Percentile	-0.9384	-0.9033	-0.9888
25th Percentile	-0.1701	-0.1411	-0.1973
50th Percentile	0.0260	0.0424	0.0144
75th Percentile	0.2672	0.2751	0.2619
99th Percentile	1.5674	1.4725	1.5684
Δ_σ in Log Income:			
Mean	0.0897	0.0951	0.0900
SD	0.4528	0.4415	0.4671
1st Percentile	-0.9930	-0.9473	-1.0327
25th Percentile	-0.1486	-0.1426	-0.1544
50th Percentile	0.0542	0.0602	0.0554
75th Percentile	0.2975	0.3158	0.2876
99th Percentile	1.4556	1.3573	1.6073

Notes: This table characterizes the distribution of the "treatment intensity," Δ_σ, within our sample of moves. We provide statistics for variation in log consumption and in log income.

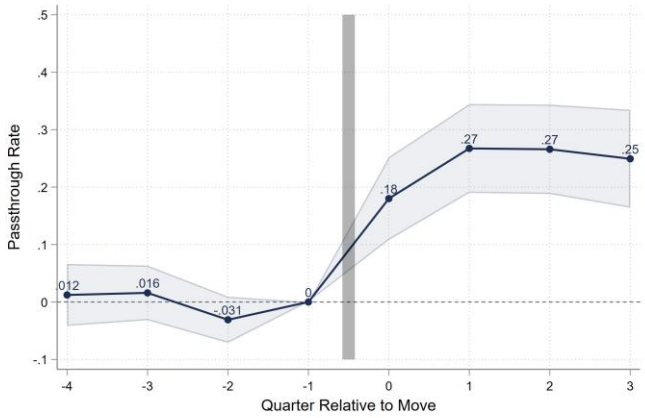
4.2. Dynamic Effects of Cross-State Moves

We now turn to the main analysis and estimate equation (3), using its empirical analogue in equation (1), where the outcome is the log of household expenditure. The results for all cross-state moves are provided in Figure 3. In this and the subsequent figures, the x-axis represents time relative to the move, and on the y-axis we plot estimates for the passthrough rates over time, λ_r, along with their 95-percent confidence intervals. We note that quarter 0 is a transition period during which individuals are only partially exposed to the treatment of the move, so quarter 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 in the pre-period across cohorts who differ in treatment intensity, 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, this finding implies meaningful location and firm effects in determining household consumption choices. They imply that about a quarter of the entire nationwide consumption differentials are attributed

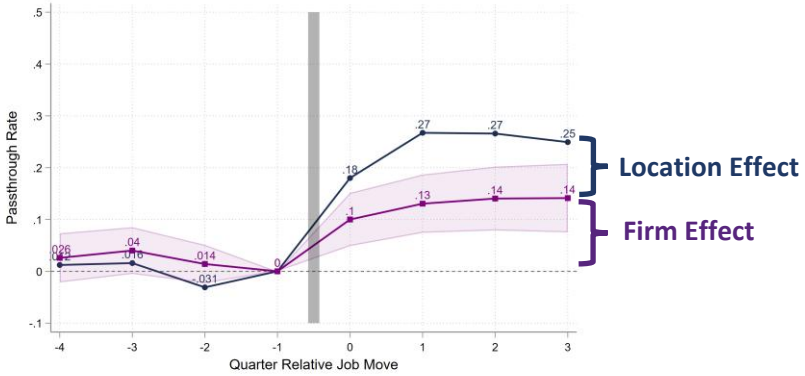
to location and firms. To clarify the economic relevance of this share, we note that the differentials in the denominator (Δ_σ) use raw measures of household consumption (i.e., with no controls), so they include all heterogeneity dimensions: in skills, preferences, education, wealth, life circumstances, and so on.

Figure 3: Dynamic Effects of Cross-State Moves on Household Consumption



Notes: This figure plots estimates of the effects of cross-state moves on the dynamics of household consumption. The *x*-axis represents time relative to the move in quarters, and on the *y*-axis we plot estimates for the passthrough rates over time, λ_r from equation (3), along with their 95-percent confidence intervals.

Figure 4: Decomposition of the Effects on Consumption: Location and Jobs



Notes: This figure plots estimates of the effects of moves on the dynamics of household consumption. The *x*-axis represents time relative to the move in quarters, and on the *y*-axis we plot estimates for the passthrough rates over time, λ_r from equation (3). We provide two sets of estimates: one for cross-state moves that identify the total effect (in blue), and one for moves within 5-digit ZIP codes that identify the firm effect (in purple), along with the corresponding 95-percent confidence intervals. The difference between the two plots captures the average location effect.

Next, we decompose the total treatment effect into the average location component and the within-location firm component. In Figure 4, we overlay on top of the results from Figure 3 for cross-state moves, estimates of equation (3) for job moves within 5-digit ZIP codes. This exercise suggests that, out of the total effect of the move (26 pp), 14 pp (54%) are attributed to the individual's job and 12 pp (46%) are attributed to location in itself.

4.3. Mechanisms

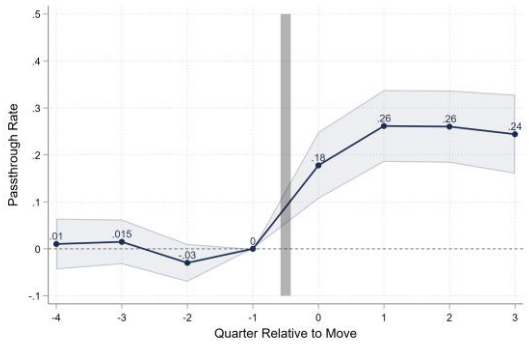
It is useful to organize the discussion of the ways in which expenditure may change following a move using the household's within-period budget constraint for a bundle of M goods:

$$\sum_{m=1}^{M} \underbrace{p_{m(l)}}_{\text{Price vector in location } l} \times \underbrace{c_{m,i}}_{\text{Choice of each product } m} = \underbrace{I_i}_{\text{Income}} + \underbrace{Y_i}_{\text{Baseline wealth}} - \underbrace{\epsilon_i}_{\text{Any residual activity (such as savings)}}.$$

There are three possible channels of change: (1) prices, $p_{m(l)}$, that capture the technical conversion of income to consumption; (2) household income, I_i, that captures changes in consumption capacity; and (3) consumption choices, $c_{m,i}$: both in terms of the relative shares of components of the bundle and in terms of levels conditional on income and prices, which would capture "location-dependence" in preferences.

We begin with Figure 5 that estimates a version of equation (3) which includes a quadratic in RPP (Regional Price Parity) for the state in which the individual resides during the observation's period. The estimates are almost identical to those in Figure 3, suggesting that prices play a minor role.

Figure 5: Effects of Cross-State Moves with Regional Price Parity (RPP) Controls

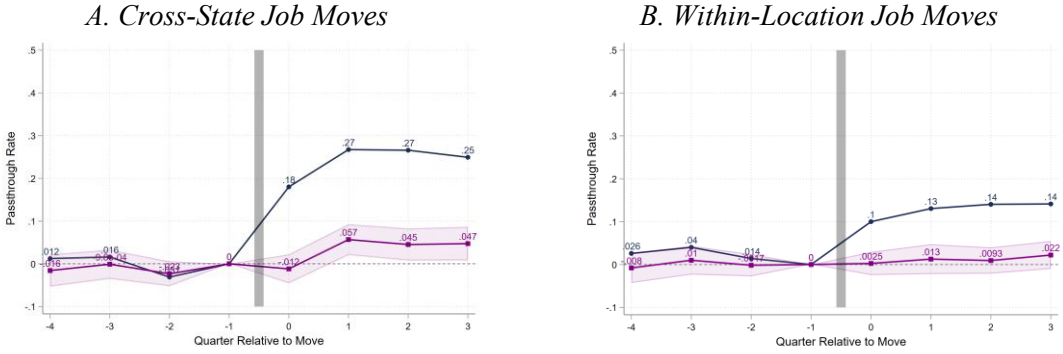


Notes: This figure plots estimates of the effects of cross-state moves on the dynamics of household consumption. The *x*-axis represents time relative to the move in quarters. On the *y*-axis we plot estimates for the passthrough rates over time, λ_r, from a specification of equation (3) that includes a quadratic in Regional Price Parities (RPPs), along with their 95-percent confidence intervals.

In Figure 6, we then estimate equation (3) where we include a quadratic in log income in the period of observation. In panel A on the left, we provide the estimates for cross-state moves, with and without income controls. We see a persistent effect that cannot be accounted for by income, which averages to approximately 5 pp in the post-move periods (quarters 1-3). We term this share as "excess consumption." In panel B on the right, we instead provide estimates for within-location job moves, with and without income controls. We see that the effect on consumption disappears when we control for income, suggesting that differences in consumption premia across firms are fully explained by differences in pay. Though intuitive, it also precludes coworker peer effects in consumption "levels" ("excess consumption") from the perspective of a nationwide average.

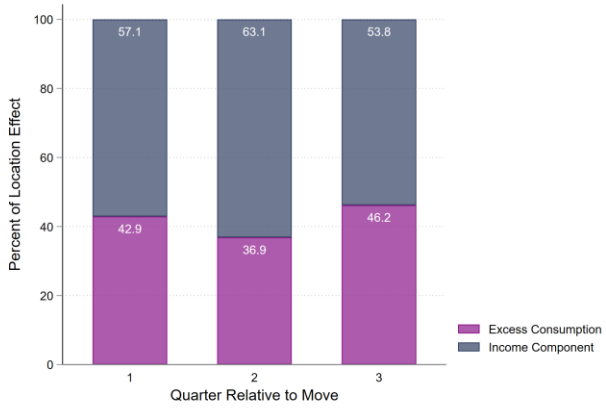
In combination, the two panels of Figure 6 imply that the 5 pp "excess consumption" come from the average location effect (exclusive of the firm effect). Accordingly, Figure 7 plots the decomposition of the average location effect (i.e., the differential between the total effect and the within-location firm effect) into an income component and an "excess consumption" component. On average, about 40% of the average location effect on consumption is driven by increases in consumption conditional on prices and income, that is, come from excess consumption. Through the lens of a standard consumption choice model, such changes in choice trace back to changes outside the budget set. That is, a meaningful share of the change in consumption is driven by "location-dependence" in preferences, inclusive of social incentives. A natural exercise for investigation is analyzing categories of consumption, specifically comparisons across spending categories that are more "visible" to those that are more reflective of daily necessities.

<u>Figure 6:</u> Mechanisms: Income versus Excess Consumption



Notes: This figure plots estimates of the effects of moves on the dynamics of household consumption. The *x*-axis represents time relative to the move in quarters, and on the *y*-axis we plot estimates for the passthrough rates over time, λ_r from equation (3). Panel A studies job moves across states, and panel B studies job moves within 5-digit ZIP codes. In each panel, we provide two sets of estimates: one without income controls (in blue), and one that includes a quadratic in log income in the period of observation (in purple), along with the corresponding 95-percent confidence intervals.

Figure 7: Mechanisms in Consumption Effect of Location: Income versus Excess Consumption



Notes: This figure displays in a dynamic decomposition of the effect of location on consumption in household moves, split into an income component and an "excess consumption" component that cannot be explained by changes in income.

Table 3: Average Effects on Consumption Categories

Consumption categories:	Cross-State Moves (1)	Moves within ZIP codes (2)
Total consumption	0.2291	0.1018
	(0.0347)	(0.0258)
Travel	0.3508	0.0459
	(0.0724)	(0.0602)
Recreation	0.1445	0.0435
	(0.0615)	(0.0534)
Dining	0.1212	0.0355
	(0.0519)	(0.0401)
Groceries	0.0293	-0.0204
	(0.0793)	(0.0502)

Notes: This table provides estimates for the average passthrough rates on different consumption categories following a household move. It provides estimates for λ_r from a specification of equation (3) that aggregates quarters 1-3 into the post-move period.

Table 3 provides evidence consistent with this view. We estimate specifications of equation (3) for the average passthroughs by aggregating quarters 1-3 into the post-move period. This allows us to increase precision in the category breakdown, but we should note the decreased overall precision in the analysis. We estimate this specification when the outcome variable varies by the consumption category under consideration (dining, grocery, recreation, travel), analyzing cross-state moves and moves within-location for a comparison. In column 1, we see that the cross-state effect is highest for "travel," remains meaningful among other potentially visible categories ("recreation" and "dining"), and is then null for "groceries," a necessities category. Moreover, there are no detectable passthroughs in these more-visible categories in within-location moves, which did not exhibit excess consumption. In combination, the category analysis is consistent with the view that the "location-dependence" in preferences upon moves could be shaped by social incentives.

5. Welfare and Policy Implications

In this final section, we study the implications of our findings by analyzing the connection between income changes and consumption changes in household moves. To do so, we re-estimate equation (3), but where the treatment intensity variable is in terms of income, i.e., the differential in income across individuals in origin and destination units. We run two versions of this regression: one when using consumption as the outcome variable (producing passthrough estimates that we denote by λ^C), and one when using income as the outcome variable (producing passthrough estimates that we denote by λ^I). From the mapping of λ's (passthroughs) to τ's (premia) in Section 3, the ratio of the former identifies the ratio of the latter. Hence, the premia MPC can be estimated by the ratio of the passthroughs across the two equations: λ^C/λ^I.

Figure 8 provides the following analysis. In panel A on the left, we plot the dynamic estimates of specifications of equation (3) with variation in income as the treatment intensity, with either own consumption or own income as the outcome variable. To arrive at one estimate for each of the passthrough rates λ^C and λ^I, we also run average regressions that aggregate over post-move periods 1-3. We calculate the MPC as the ratio of these average estimates, as reported in column 1 of Table 4. In panel B of Figure 8, we plot location-level consumption (on the y-axis) against location-level income (on the x-axis). We first plot means of the raw data to describe the cross-section, and we then layover the portion of the relationship that is "casual": the premia MPC. We find that the MPC in cross-state moves is 0.82 (with $\lambda^C = 0.24$ and $\lambda^I = 0.29$). The cross-sectional slope is 0.95, so that 86% (=0.82/0.95) of the observed spatial relationship between income and consumption is due to households' marginal propensity to consume.

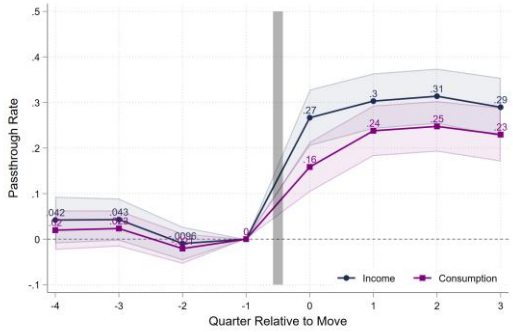
This has two sets of implications. First, by its definition, the premia MPC characterizes and quantifies the differences in the conclusions we would arrive at if we were to analyze income measures

instead of consumption. Assessing the implications of spatial economic dispersion on household welfare requires translating a dollar of income to consumption terms (as a utility input). An MPC of 0.82 implies that any $1 attributed to improvements in well-being in the analysis of income is over-estimated by 18 cents in comparison to the analysis of consumption.

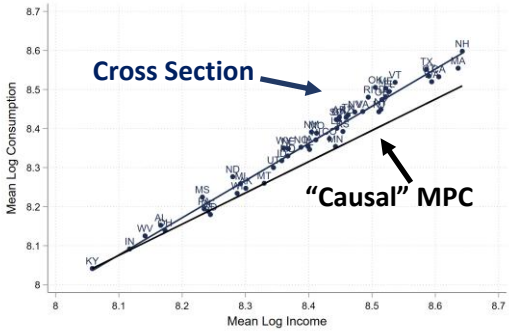
Second, our estimated MPC has implications for placed-based policies, in both efficiency and equity terms. Place-based policies essentially affect the location income effect. The premia MPC (that maps location income premia to location consumption premia) can therefore map the effects of place-based policies into welfare via household consumption. Since place-based policies can target location at differing levels of aggregation, we consider MPCs in moves across states, moves across counties, and moves across 5-digit ZIP codes. Table 4 shows that the MPC is very similar along the distance of moves. For welfare predictions, our estimates imply that every $1 of an "efficacious" place-based policy should translate, on average, to an 82-83 cent increase in household consumption. Lastly, considering the equity implications of such policies, we study heterogeneity in passthrough rates across the household baseline income rank, by estimating the average regressions separately for households in the bottom, middle, and top thirds of the income distribution. The results are provided in Table 5. We find that location premia tend to benefit lower-income families to a higher degree, so that placed-based policies (that vary the location premia) can encompass further redistributive gains, i.e., not only across locations but also within.

Figure 8: Premia MPC

A. Cross-State Moves with Income as Treatment Intensity



B. Cross-Section State-Level Relationship and Premia MPC



Notes: Panel A provides the dynamic estimates of specifications of equation (3) with variation in income as the treatment intensity, with either own consumption or own income as the outcome variable. Panel B plots location-level consumption against location-level income. We plot means of the raw data to describe the cross-section, and we layover the portion of the relationship that is "casual," the premia MPC.

Table 4: Premia MPC for Varying Levels of Location

Outcome:	State (1)	County (2)	ZIP Code (3)
Expenditure	0.24	0.22	0.18
	(0.03)	(0.02)	(0.02)
Income	0.29	0.27	0.21
	(0.03)	(0.03)	(0.02)
MPC	0.82	0.82	0.83

Notes: This table provides estimates for average passthrough rates from specifications of equation (3) that aggregate quarters 1-3 into the post-move period. In these regressions, we use differences in income across origin and destination units as the treatment intensity, with either own consumption or own income as the outcome variable. Column 1 includes moves across states, column 2 includes moves across counties, and column 3 includes moves across 5-digit ZIP codes. The MPC is calculated as the ratio of the coefficient on expenditure and the coefficient on income.

Table 5: Treatment Effect of Moves by Household Income Rank

	Consumption as Treatment Intensity		Income as Treatment Intensity	
Class of Move:	Cross-State	Cross-County	Cross-State	Cross-State
Outcome:	Consumption	Consumption	Income	Consumption
	(1)	(2)	(3)	(4)
Income Group:				
Bottom Third	0.34	0.33	0.34	0.29
	(0.05)	(0.05)	(0.05)	(0.04)
Middle Third	0.23	0.20	0.25	0.21
	(0.05)	(0.05)	(0.04)	(0.04)
Top Third	0.17	0.13	0.20	0.15
	(0.04)	(0.04)	(0.03)	(0.03)

Notes: This table provides estimates for average passthrough rates from specifications of equation (3) that aggregate quarters 1-3 into the post-move period. We provide different versions that vary the treatment intensity variable (using either consumption or income differentials), the distance of the move (across states or counties), and the outcome variable (consumption or income).

6. Conclusion

This paper uses transaction-level bank data to study the role of location and firms in determining American households' consumption behavior and its nationwide dispersion. Based on quasi-experimental variation that leverages migration across locations in a movers' design, we find that about a quarter of the cross-state differences in consumption would fall if households were randomly reallocated between them. We find that about 54% of this effect are attributable to firms, and 46% are attributed to an average location effect. While the firm consumption effect is fully accounted for by income changes, 40% of the average location effect is not tied to income. This "excess consumption" implies, in turn, location-dependence in preferences, with some evidence consistent with social incentives. Our findings have implications for placed-based policies, providing predictions for their efficacy in improving welfare via household consumption, and pointing to potential additional distributional gains of benefitting families at the lower-end of the income distribution to a higher extent.

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.

Bronnenberg, Bart J., Jean-Pierre H. Dubé, and Matthew Gentzkow (2012). "The Evolution of Brand Preferences: Evidence from Consumer Migration." *American Economic Review*, 102(6): 2472-2508.

Busso, Matias, Jesse Gregory, and Patrick Kline (2013). "Assessing the Incidence and Efficiency of a Prominent Place Based Policy." *American Economic Review*, 103(2): 897-947.

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.

Chetty, Raj, David Grusky, Maximilian Hell, Nathaniel Hendren, Robert Manduca, and Jimmy Narang (2017). "The Fading American Dream: Trends in Absolute Income Mobility since 1940." *Science*, 356 (6336): 398-406.

Chetty, Raj, Nathaniel Hendren, Maggie R. Jones, and Sonya R. Porter (2020). "Race and Economic Opportunity in the United States: An Intergenerational Perspective." *Quarterly Journal of Economics*, 135(2): 711-783.

Fadlon, Itzik, Briana Sullivan, and Vedant Vohra (2025). "The Racial Penalty in Job Ladder Transitions." NBER Working Paper w34058.

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.

Finkelstein, Amy, Matthew Gentzkow, and Dean Li (2025). "What Drives Risky Prescription Opioid Use? Evidence from Migration." *Quarterly Journal of Economics* (2025).

Hubmer, Joachim, Per Krusell, and Anthony A. Smith Jr. (2021). "Sources of US Wealth Inequality: Past, Present, and Future." *NBER Macroeconomics Annual*, 35(1): 391-455.

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.

Saez, Emmanuel, and Gabriel Zucman (2016). "Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax Data." *Quarterly Journal of Economics*, 131(2): 519-578.

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.