Optimal Place-Based Transfers using Measured Geographic Variation in the Marginal Utility of Income

Md Moshi Ul Alam — Clark University
Morris A. Davis — Rutgers University
Jesse Gregory — University of Wisconsin

December 23, 2025

Preliminary and Incomplete – Do Not Cite

Abstract

We address a fundamental identification challenge in optimal place-based redistribution policy design: characterizing how marginal utility of income (MU) varies across locations. We show that MU is not identified from location choice data alone, even with full-support location-specific instruments that generate arbitrarily large positive or negative income shifts in every location. To overcome this challenge, we design and implement a nationally representative survey that directly elicits location-specific MUs and recovers how respondents would substitute across locations when relative net incomes change, using respondent-specific feasible choice sets. We characterize optimal place-based transfers using both a sufficient-statistics approach with an iterative algorithm for marginal policies, and a structural model with heterogeneous types and non-separable preferences for non-marginal policies. Although location is often treated as a tag for types (such as skill groups), we find that nearly half of the cross-metro variation in the planner's welfare weights—and thus the motive to redistribute—comes from within-person and within-skill-group differences in marginal utility across locations, identified only through our survey. We find that optimal policy is substantially more nuanced than a simple redistribution from high-wage to low-wage locations.

1 Introduction

Place-based transfers across locations have the potential to raise average welfare by correcting externality-driven misallocation and redistributing toward places where the marginal utility of additional income is high. A growing literature has made considerable theoretical progress characterizing optimal spatial transfers motivated by agglomeration externalities (Fajgelbaum and Gaubert, 2020, 2025; Donald et al., 2025), knowledge spillovers across occupations and locations (Rossi-Hansberg et al., 2022), and equity concerns (Gaubert et al., 2021). These papers show that optimal taxes and transfers depend on the magnitude of agglomeration and productivity spillovers, elasticities of location choice with respect to after-tax income, and variation across locations in the marginal utility of income, both across worker types that differ in earnings capacity (location as a tag) and within-type differences across locations. The empirical literature provides quantitative evidence on the size of agglomeration and productivity externalities and on the responsiveness of location choices to financial incentives. However, little direct evidence exists on how the marginal utility of income varies across locations. This paper develops and implements a new survey-based approach that measures within-person and within-type differences in the marginal utility of income across locations, and, aggregating these comparisons, studies the consequences of that variation for optimal redistribution-motivated spatial transfers.

We first derive optimal redistribution formulas for location-specific transfers, which make clear that variation in the marginal utility of income across places is a key driver of redistribution-motivated transfers. To isolate the identification and measurement issues at the center of the paper, we set aside place-based externalities and focus on the redistribution component of the planner's problem. We then turn to our main contribution: measuring this variation directly and using it to study redistribution-motivated transfers across space. The formulas highlight that a planner who chooses transfers across locations, subject to a budget constraint, needs two sets of sufficient statistics. The first is the full set of marginal utilities of income by location and worker type. This determines where an extra dollar has the largest welfare payoff. The second is the location demand system, summarized by a matrix of location choice elasticities describing how people re-sort when transfers change. This characterizes the efficiency cost of transfers to any given location in terms of deadweight loss. We show that the marginal utilities in this system are not identified by data in which we only observe how location choices respond to changes in financial incentives, without imposing a priori assumptions on the separability of preferences, the role of unobserved amenities, or the properties of idiosyncratic location attachments. This non-identification motivates moving beyond standard revealed-preference approaches and directly measuring how the value of additional income varies across locations.

Our solution is to measure the planner's key inputs directly. We field a nationally representative survey that elicits respondent-specific location choice sets which we henceforth refer to as menus, stated choice probabilities under experimentally varied scenarios, and within-person measures of the marginal utility of income across locations where each respondent could plausibly live in the future. For each respondent, we elicit four metropolitan areas that they realistically could choose to live in several years ahead: their hometown and three additional unique destinations. Respondents then report their expected income in each location, describe contextual factors such as the presence of family and friends and how well the local environment fits their interests, and state the probabilities with which they would choose each location. We then elicit within-person comparisons of the value of a permanent income increase across these locations. In particular, respondents first indicate which location-specific version of themselves would benefit most from a permanent $5,000 raise. We then ask, for each of the other locations, what permanent income increase there would deliver the same increase in well-being as the $5,000 raise in the chosen location. Repeating this set of comparisons across experimentally varied hypothetical scenarios that shift incomes and contextual factors allows us to recover how each respondent's marginal utility of income differs across locations and over a range of income levels.

An alternative, more traditional approach would be to infer cross-location differences in the marginal utility of income from revealed preferences–using how location choices respond to differences in income opportunities across places. Under a separable random-utility model, marginal utilities are identified from cross-location wage responses: how the probability of choosing one location changes when wages change in other locations. Identification therefore relies on ratios of these cross effects, not on the own-wage elasticities that applied work typically targets. Measuring the required cross effects would in turn require location-specific, plausibly exogenous wage variation and enough precision to detect how shocks in one location shift choice probabilities elsewhere for *all* locations, which is difficult in practice. Instead, quantitative work typically relies on functional-form assumptions or effectively imposes that marginal utility is directly tied to the location's income. Moreover, the separability condition on which this revealed-preference argument rests is not supported in our data. These practical and conceptual limitations necessitate an alternative approach, which is the role played by our survey-based comparisons.

The second crucial input to the optimal transfer formulas is the location-demand system, summarized by own- and cross-location elasticities of location choice with respect to location-specific after-tax income. A planner needs to know not only how sensitive each city's demand is to its own net income, but also the substitution pattern across cities: when one city becomes less attractive, which other cities gain the displaced residents and by how much. A natural analog is the discrete-choice product-demand literature, which recovers rich substitution patterns across alternatives. In that literature, heterogeneity in preferences for observable attributes–modeled through random coefficients and interactions with observed characteristics–generates realistic substitution: when the price of one product rises, close substitutes absorb a disproportionate share of the switching, and these switching patterns identify the matrix of own- and cross-price elasticities.The difficulty is that the data requirements needed to estimate these elasticities are hard to satisfy for location demand. In product-demand applications, one typically observes similar sets of products sold in many markets with different price vectors, so that the same product is repeatedly observed under many relative-price regimes. There is no close analog for cities: we do not repeatedly observe the full set of U.S. cities under many different configurations of wages and living costs, and plausibly exogenous variation in relative location-specific net incomes is limited. As a result, most spatial applications calibrate a small number of elasticities or rely on restrictive functional forms rather than estimating a full matrix of own- and cross-location elasticities.

Our survey provides a way to bring the logic of discrete-choice demand to location decisions. At the individual level, we observe stated choice probabilities over a menu of cities, which reveals which cities are commonly considered together and thus likely to receive one another's outflows when relative net incomes change. Across experimentally varied scenarios, we shift incomes in specific cities and track how each respondent's choice probabilities move, yielding individual-level substitution patterns between locations in individual's menu. Aggregating these responses produces the analog of the cross-price effects that define substitution patterns in discrete-choice demand systems. This ability to estimate a flexible location-demand system is a second central contribution of the paper. Finally, we supplement our survey data with multiple other publicly available data sources to provide extensive validation checks of our survey data.

Our survey yields several key empirical patterns. First, our estimated average location choice elasticity is 1.69, which falls within the range of 1.2 to 3.1 documented in prior work using revealed-preference methods (Bound and Holzer, 2000; Serrato and Wingender, 2014; Serrato and Zidar, 2015; Notowidigdo, 2011; Diamond, 2016), providing reassurance that stated-preference responses align with observed behavior. The experimental variation across scenarios also allows us to estimate the distribution of individual-level elasticities, which we recover using an empirical Bayes shrinkage procedure (Efron, 2016).

Second, our marginal utility measures exhibit substantial geographic variation that is not explained by income differences alone. Within-person regressions show that location-specific marginal utility is weakly but positively correlated with income, while marginal utility is lower in high-housing-cost locations holding income fixed. Locations with better schools and stronger labor markets (measured by non-wage job amenities) exhibit higher marginal utility, while high-crime locations exhibit lower marginal utility. We summarize this geographic variation using an AKM-style decomposition (Abowd et al., 1999) that isolates metro fixed effects from the within-person comparisons, and these metro effects serve as a key input to the optimal transfer formulas.

Third, our menu-based design, which allows for individual-specific location choice sets, reveals rich substitution patterns across space. Tracing how stated choice probabilities shift in response to location-specific income shocks, we find strong regional clustering.[1] This pattern reflects both geographic proximity and the tendency for respondents to list culturally and economically similar metros in their choice sets.

Combining these inputs, we compute optimal metro-specific taxes and transfers for over more than 600 U.S. metropolitan areas. The distribution of optimal taxes exhibits substantial dispersion and is notably bimodal: one group of metros optimally receives subsidies while another group faces positive taxes, with relatively few metros near the break-even point.[2] This pattern does

[1]For example, marginal movers from Louisiana would relocate predominantly to Texas, Mississippi, and Florida, with minimal substitution to Northern or Western states.

[2]Geographically, coastal metros–particularly along the Eastern Seaboard and in California–tend to face higher

not simply mirror a tax-from-high-wage-to-low-wage logic. Instead, it reflects measured marginal utility across locations, which can depend on local amenities, cost of living, and the presence of family networks. We decompose cross-metro variation in optimal taxes using our formula, which expresses each metro's optimal tax as the sum of two terms: the migration-weighted average tax of substitute metros, and an adjustment that depends on the metro's marginal utility deviation from the national mean. Allowing for heterogeneity by skill type, three variables–the metro-level marginal utility effect, the local college share, and the average tax of substitutes–explain sixty percent of the variation in optimal taxes. The bimodality and geographic clustering arise because metros that are close substitutes tend to share similar characteristics generating groups of metros that are taxed or subsidized together. Most importantly, the planner's motive to redistribute within-person and within-type—identified through our survey—accounts for half of the variation in MU differences across locations required to obtain location-specific optimal taxes. Finally, we show that the scope for spatial redistribution depends on migration elasticities. Doubling location-choice elasticities compresses the optimal tax distribution by approximately 25%, as higher mobility raises the fiscal cost of tax differentials across space.

Next, we build a structural model to compare welfare under sufficient statistics derived optimal redistribution across cities with other policies. Following our results on non-identification, our model allows for imperfect substitutability between consumption and the idiosyncratic shocks. This is not identified by the location choice data alone, but is identified by both location choice data supplemented with direct measurement of marginal utility of income which is necessary for optimal policy. The imperfect substitutability allows marginal entrants to a given location to have different marginal utilities of income than inframarginal entrants. Our model is identified from the exogenous variation that we induce through our survey by exogenously shifting individual's location-specific expected income and contextual factors. We estimate the model using Indirect Inference (Gourieroux et al., 1993) and find that the model is able to replicate the empirical patterns in the data and is able to match non-targeted moments like average choice probability by the rank of locations. Using the estimated parameters of the model we can simulate counterfactual policies and compare welfare.

This paper contributes to several literatures. First, we contribute to the growing literature on optimal spatial taxes and transfers by providing new evidence on a central sufficient statistic—how the marginal utility of income varies across locations—and by showing that this object is not identified from migration responses to financial incentives alone without strong preference restrictions. Second, we contribute to the empirical literature on spatial equilibrium and migration by developing a survey-based approach that elicits within-person comparisons of the value of additional income across plausible future locations, yielding location-specific marginal utility measures that do not rely on separability assumptions. Third, we contribute to the literature on location choice by providing new evidence on which cities households treat as substitutes: using individualized choice menus and experimentally varied scenarios, we recover substitution patterns across cities

taxes, while metros in the Southeast and parts of the Midwest warrant subsidies.

and construct the matrix of own- and cross-location responses required for optimal policy. Finally, we contribute to the stated-preference and survey-measurement literature by showing how a menu-based, within-person design can elicit welfare-relevant objects—marginal utilities and substitution patterns—that are difficult to recover from revealed-preference data.

The rest of the paper is structured as follows. In Section 2, we derive the optimal redistribution formula and its components which form the central object of interest of our paper. In Section 3, we discuss the strong assumptions on which identification rests in canonical models with separability, and yet is infeasible to estimate. Next we show that the components are not identified with location choice data under non-separability with examples illustrating observational equivalence. Section 4 provides an overview of our survey and describes how we elicit the components required for optimal taxes and construct measures of the corresponding sufficient statistics. Section 5 provides provides a battery of evidence consistent with existing literature which provide validation checks for our survey. Next in Section 6 and 7, we describe the empirical patterns in the marginal utility of income and in the geographic substitutions respectively. Section 8 computes optimal place-based transfers using the Gauss-Sidel iterative algorithm. We discuss results on optimal transfers in Section 9. In section 10 and 11 we build the structural model and discuss parameter estimates and model fit. In Section 12 we (plan to) use the model to compare welfare across various policies.

2 Optimal Redistribution Formula and Matrix Setup

In this section we characterize optimal place–based redistribution when people can move across metropolitan areas. The government chooses a schedule of metro–specific head taxes and transfers, and households respond by sorting across metros based on after–tax incomes and their idiosyncratic location preferences. Because taxes change both consumption and where people live, optimal policy trades off redistribution against the fiscal externalities created by migration.

The main message is that optimal policy depends on two ingredients. The first is the marginal value of an additional dollar of income for the people who live in each metro (after accounting for the planner's Pareto weights). This object determines where a marginal dollar delivers the largest welfare gain. The second is the location demand system: how sensitive metro populations are to metro–specific taxes, including where movers go when a given metro becomes more or less attractive. These mobility responses determine how costly it is to raise revenue from a particular metro, because changing a tax shifts the tax base both locally and in other metros.

We proceed as follows. We first set up the environment and the planner's problem. We then derive a sufficient-statistics characterization of optimal taxes and provide intuition using a two–metro special case. We next turn to the identification question that motivates our survey-based measurement of marginal utilities. We return to the matrix representation and iterative solution method later when we implement the policy calculations.

2.1 Environment

There is a unit mass of individuals indexed by i. Individuals belong to finitely many observable types $t(i) \in \{1, \ldots, T\}$. Let N_t denote the population share of type t, so $\sum_t N_t = 1$. Locations are indexed by $j \in \{1, \ldots, J\}$ and correspond to metropolitan areas.

For a given vector of metro–specific head taxes (or subsidies when negative)

$$\tau = (\tau_1, \ldots, \tau_J),$$

type t individuals face the indirect utility vector $\{u^t_{ij}(w_{ij} - \tau_j, \epsilon_{ij})\}_{j=1}^J$ over locations and choose their location optimally. Let

$$P^t_j(\tau) = \Pr(\text{type } t \text{ chooses metro } j \mid \tau)$$

denote the resulting choice probability, and define the total population share in metro j as

$$L_j(\tau) = \sum_t N_t P^t_j(\tau).$$

Each individual in metro j pays the same head tax τ_j. A marginal change in τ_j changes consumption one–for–one for all individuals who choose metro j, so the relevant object for welfare is the expected marginal utility of income in that metro, averaged over the individuals who actually live there. For type t, let μ^t_j denote the expected marginal utility of income conditional on choosing metro j. Specifically,

$$\mu^t_j \;\equiv\; \mathbb{E}\big[u^t_w(w_{ij} - \tau_j, \epsilon_{ij}) \mid j \text{ chosen}\big]\,, \tag{1}$$

2.2 The planner's problem

The planner's Pareto objective is

$$W(\tau) \;=\; \sum_{t=1}^T \Pi_t N_t U_t(\tau),$$

where Π_t is the Pareto weight on type t and $U_t(\tau)$ is expected utility for a type–t individual under taxes τ. The planner maximizes the Pareto objective by choosing τ subject to a balanced budget constraint

$$\max_\tau \; W(\tau) \quad \text{s.t.} \quad \sum_{j=1}^J \tau_j L_j(\tau) = G,$$

where G is exogenous net spending, possibly zero.

Proposition 2.1 (Sufficient statistics for optimal metro head taxes)**.** *Suppose (i) a metro–specific head tax changes consumption one–for–one for residents of that metro, (ii) location choice prob-*

abilities are differentiable in the tax vector τ, and (iii) a common marginal increase in all metro head taxes leaves location choices unchanged in a neighborhood of the optimum.[3] *Then any interior optimum τ that satisfies the balanced budget constraint $\sum_{j=1}^{J} \tau_j L_j(\tau) = G$, must satisfy, for each metro j,*

$$\frac{\mu_j - \bar{\mu}}{\bar{\mu}} = \frac{1}{L_j(\tau)} \sum_{j'=1}^{J} \tau_{j'} \frac{\partial L_{j'}(\tau)}{\partial \tau_j}. \tag{2}$$

Here μ_j is the Pareto–weighted average marginal utility of income among residents of metro j, and $\bar{\mu} \equiv \sum_{j=1}^{J} L_j(\tau)\mu_j$ is the Pareto–weighted average marginal utility of income in the population. Appendix A provides the derivation.

Equivalently, by rearranging equation (2), it follows that the optimal tax for city j can equal the sum of the average tax of substitute metros and a term capturing the deviation of MU of metro j from the average

$$\tau_j = \underbrace{\frac{\sum_{j' \neq j} \frac{\partial \mathcal{P}_{j'}}{\partial \tau_j} \tau_{j'}}{\sum_{j' \neq j} \frac{\partial \mathcal{P}_{j'}}{\partial \tau_j}}}_{\text{Average tax of substitute metros}} + \underbrace{\frac{1}{\frac{\partial \mathcal{P}_j}{\partial \tau_j}} \cdot \frac{\mu_j - \bar{\mu}}{\bar{\mu}}}_{\text{MU-based adjustment}} \tag{3}$$

The second term in (3) captures the Baily–Chetty intuition: metros with above-average Pareto-weighted marginal utility of income should be net recipients of transfers (i.e., face lower taxes), while metros with below-average marginal utility should be net contributors (i.e., face higher taxes). The magnitude of this adjustment is governed by the responsiveness of the metro's population share to changes in its tax, $\partial \mathcal{P}_j / \partial \tau_j$: when population flows are more elastic, redistribution through metro-specific taxes is more distortionary, so the marginal-utility adjustment is attenuated.

The first term captures the fiscal-externality channel created by migration. When metro j is taxed more heavily (or subsidized less), marginal residents who are induced to leave substitute toward other metros and pay the taxes prevailing in those destinations. Because residents leaving a given metro may substitute more toward some metros than others, the average tax paid by those exiting j will in general differ across j. The weighted average of those destination taxes is exactly the first term, and can be interpreted as the average tax faced by the marginal outflow from j.

Together, the formula shows that optimal policy starts from a "baseline" determined by the taxes in substitute metros and then shifts taxes up or down depending on whether μ_j is below or above the population average. In particular, if $\mu_j = \bar{\mu}$, the optimal policy generally does *not* set $\tau_j = 0$; instead, it sets τ_j equal to the average tax faced by the marginal movers who would leave j in response to a marginal increase in τ_j.

Equations (2) and (3) highlight two takeaways. First, (2) makes clear the objects the planner needs: the Pareto-weighted marginal utilities $\{\mu_j\}$ and the matrix of migration responses $\{\partial L_{j'} / \partial \tau_j\}$. Second, (3) rewrites the same condition in a form that is closer to a policy prescrip-

[3] This holds in standard random utility models in which utilities depend on after–tax consumption and idiosyncratic location shifters, because a common tax change shifts all location utilities by the same amount. Appendix A formalizes the argument.

tion: a metro's optimal tax equals the average tax faced by the marginal outflow from that metro, plus an adjustment that depends on whether μ_j is above or below the population average, scaled by the own migration response. In the multi-metro setting these conditions define a system in which each τ_j depends on the full vector τ, so computing the optimal taxes requires solving a fixed point. In our computations below, we implement this by iterating on the mapping implied by (3) until convergence.

Equation (2) is a spatial analog of the sufficient-statistics formula for optimal unemployment insurance: the left-hand side captures the marginal welfare gain from raising net resources in metro j relative to the population average, while the right-hand side captures the marginal fiscal cost created by migration-induced changes in the tax base across metros. The two-city example in the next subsection makes this analogy transparent in the simplest setting, where there is only one margin of re-sorting.

2.3 Two–city example (intuition)

To provide intuition, consider a simplified environment with only two locations, indexed by $j = 1, 2$, and a single individual type. The government wishes to redistribute income from metro 2 to metro 1 via a head transfer b to everyone in metro 1, financed by a head tax t on everyone in metro 2, under a balanced–budget constraint.

Let W_1 and W_2 denote pre–tax incomes in the two metros. Under the policy, consumption is $W_1 + b$ in metro 1 and $W_2 - t$ in metro 2. Individuals choose between these two options, subject to idiosyncratic shocks (ϵ_1, ϵ_2), so indirect utility is

$$V(b,t) = \mathbb{E}_{\epsilon_1,\epsilon_2}\left[\max\left\{u(W_1 + b, \epsilon_1),\ u(W_2 - t, \epsilon_2)\right\}\right].$$

Let $L_1(b)$ denote the equilibrium share of the population in metro 1 when the transfer is b. With a balanced budget, total revenue from metro 2 must equal total transfers to metro 1:

$$L_1(b)\, b = [1 - L_1(b)]\, t(b), \qquad \text{so} \qquad t(b) = \frac{L_1(b)}{1 - L_1(b)}\, b.$$

Differentiating this expression yields

$$\frac{dt}{db} = \underbrace{\frac{L_1}{1 - L_1}}_{\text{mechanical}} + \underbrace{\frac{b}{(1 - L_1)^2}\,\frac{dL_1}{db}}_{\text{"leaky bucket"}},$$

where the first term is the tax increase required to finance a higher transfer holding the population fixed, and the second term captures the additional tax increase required because the policy induces migration that expands the number of transfer recipients and shrinks the tax base.

Let μ_1 and μ_2 denote the expected marginal utilities of income for households who choose metros 1 and 2, respectively, under the policy $(b, t(b))$. The optimal transfer b satisfies

$$\frac{\mu_1 - \mu_2}{\mu_2} = \frac{1}{1 - L_1(b)} \left(\frac{dL_1(b)}{db} \frac{b}{L_1(b)} \right). \tag{4}$$

Appendix A.3 provides details of the derivation.

The left–hand side of (4) is the marginal benefit of shifting one dollar from metro 2 to metro 1, measured relative to the marginal utility in metro 2. The right–hand side is the marginal cost of raising that dollar, which depends on how strongly the transfer affects the equilibrium population share $L_1(b)$ and on the size of the tax base in metro 2.

Equation (4) is a special case of Proposition 2.1 when there are only two metros and one type. In the multi–metro setting, the scalar response dL_1/db is replaced by the matrix of cross–metro derivatives $\partial L_{j'}/\partial \tau_j$, and the optimality conditions jointly determine the entire vector of taxes, implying a fixed-point problem. Because our main focus in the next section is identification of the planner's key inputs, we defer the matrix notation and numerical solution method to the policy-calculation section.

3 Non-identification

Section 2 shows that redistribution-motivated place-based transfers depend on the average marginal utility of income among residents of each metro,

$$\mu_j \equiv \mathbb{E}[u_w(w_{ij}, \epsilon_{ij}) \mid j \text{ chosen}],$$

together with the matrix of migration responses to metro-specific taxes. This section explains why $\{\mu_j\}$ is not identified from location choice data alone without strong assumptions on preferences, even in a hypothetical environment with rich or even full-support location-specific wage shifters. The core issue is that income shifts reveal how much additional income changes a location's attractiveness for people who are close to moving—and therefore reveal marginal utilities of income for these marginal movers—whereas the planner's first-order conditions depend on the *average* marginal utility of income among the residents of each location under the policy being evaluated.

3.1 A benchmark revealed-preference argument under separability

To fix ideas, consider a standard separable random-utility model in which utility in location j takes the form

$$U_{ij} = \alpha_j \log(w_{ij}) + a_j + \varepsilon_{ij},$$

where w_{ij} is income (net of any head tax), $\alpha_j > 0$ governs the marginal utility of income in location j, a_j is a common amenity component, and ε_{ij} is an idiosyncratic location shifter that enters additively. In such models, cross-location choice responses encode information about marginal utilities. In particular, under a logit structure (and in equivalent canonical Fréchet-based formulations), cross derivatives take the form

$$\frac{\partial P_j}{\partial w_{j'}} \propto -P_j P_{j'} \frac{\alpha_{j'}}{w_{j'}}, \qquad j \neq j'.$$

That is, the effect of increasing income in j' on the probability of choosing j is proportional to the marginal utility of income in j'.

Crucially, relative marginal utilities are obtained from *cross derivatives in both directions*. For any pair of distinct locations j and j',

$$\frac{\partial P_j}{\partial w_{j'}} \propto -P_j P_{j'} \frac{\alpha_{j'}}{w_{j'}}, \qquad \frac{\partial P_{j'}}{\partial w_j} \propto -P_{j'} P_j \frac{\alpha_j}{w_j},$$

and therefore

$$\frac{\partial P_j / \partial w_{j'}}{\partial P_{j'} / \partial w_j} = \frac{\alpha_{j'} / w_{j'}}{\alpha_j / w_j}.$$

This benchmark illustrates the revealed-preference logic: comparing how the probability of choosing j responds to income changes in j' with the reverse cross response identifies the relative marginal utility of income across the two locations (up to the mechanical $1/w$ scaling under log income utility). Appendix B summarizes the canonical cases and the mapping from cross-location responses to marginal-utility ratios under separability.

Two points are important for our purposes. First, this identification argument relies on *cross* responses—how shocks to income in one location shift choice probabilities in other locations—not on the own-location elasticity that applied work typically targets. In practice, available quasi-experimental variation is typically used to estimate a single average elasticity with respect to a location's own income, and even that is often at the limit of what the data can support. For example, traditional Bartik-style instruments implemented with decennial Census data yield one predicted shifter per location per decade, which is far from the variation needed to recover the full matrix of cross-location responses required by the benchmark identification argument.

Second, this benchmark relies on separability as an identifying assumption: conditional on income, the marginal utility of income in a location does not depend on the idiosyncratic location shifter that drives selection into that location. Location-choice responses reveal how income changes shift a location's attractiveness for people who are close to switching across locations, and therefore speak most directly to marginal utilities for these marginal movers. Separability is what lets one treat these marginal- mover marginal utilities as representative of the average marginal utility among residents, μ_j.

3.2 Non-identification of policy-relevant μ_j without separability, even with full-support instruments

The more fundamental problem is that once one relaxes separability between income and idiosyncratic location shifters (entertaining the possibility that idiosyncratic factors that affect location choices may also shift the marginal utility of income), location-choice data do not identify the policy-relevant objects $\{\mu_j\}$. To illustrate, consider two locations. Individual i chooses location 1 if

$$u(w_{i1}, e_{i1}) \;\geq\; v(w_{i2}, e_{i2}),$$

and chooses location 2 otherwise. The planner-relevant average marginal utilities among residents are

$$\mu_1 = [u_w(w_{i1}, e_{i1}) \mid u(w_{i1}, e_{i1}) \geq v(w_{i2}, e_{i2})]\,, \qquad \mu_2 = [v_w(w_{i2}, e_{i2}) \mid u(w_{i1}, e_{i1}) < v(w_{i2}, e_{i2})]\,.$$

Suppose we observe location-specific instruments (Z_1, Z_2) that shift wages so that the econometrician can trace out location choice probabilities and cross-location responses as $(w_{i1}+Z_1, w_{i2}+Z_2)$ vary over a wide support. Let P_1 and P_2 denote the resulting choice probabilities for locations 1 and 2. These data are sufficient to identify ratios of marginal utilities at the *margin* in terms of location choice: cross responses identify objects tied to the set of individuals who are indifferent at a given configuration of perturbed wages. Formally, the analog of the separable identification ratios recovers ratios of derivatives evaluated at indifference (up to a common density term that cancels in the ratio),

$$\left.\frac{\partial P_2/\partial w_1}{\partial P_1/\partial w_2}\right|_{(Z_1,Z_2)} = \frac{\mathbb{E}[u_w(w_{i1}+Z_1, e_{i1}) \mid u(w_{i1}+Z_1, e_{i1}) = v(w_{i2}+Z_2, e_{i2})]}{\mathbb{E}[v_w(w_{i2}+Z_2, e_{i2}) \mid u(w_{i1}+Z_1, e_{i1}) = v(w_{i2}+Z_2, e_{i2})]}.$$

Even with full-support instruments, which under weaker assumptions allow the researcher to observe circumstances where residents who are inframarginal when $Z_1 = Z_2 = 0$ are made to be marginal, this delivers information about *their* marginal utilities of income at *perturbed* wages, not the average marginal utilities μ_1 and μ_2 among inframarginal residents at baseline wages.

3.3 Examples illustrating observational equivalence

A complementary way to understand the non-identification result is through observational equivalence. Without a separability assumption, one can construct distinct utility specifications that generate the same location choice behavior—and therefore the same predicted responses of choice probabilities to financial incentives—but that imply different values for the policy-relevant conditional averages μ_j (or, in a two-location case, μ_1 and μ_2). The examples below illustrate this logic in simple settings.

3.3.1 Example 1: Complementarity between income and shocks

To fix ideas, consider J locations and a utility index that is non-separable in income w_j and an idiosyncratic shock ϵ_{ij}. Consider:

$$u_{ij} \;=\; w_j \epsilon_{ij} \qquad \text{versus} \qquad u_{ij} \;=\; w_j \tilde{\epsilon}_{ij}, \quad \text{where} \quad \tilde{\epsilon}_{ij} = \epsilon_{ij} \times D(\epsilon_{i1}, \ldots, \epsilon_{iJ}),$$

and $D(\epsilon_{i1}, \ldots, \epsilon_{iJ}) > 0$ is an arbitrary strictly positive function of the entire vector of shocks for individual i. Importantly for the point being illustrated, the realized value of $D()$ for individual i does not vary across locations j, so it does not affect choices but it does affect marginal utilities.

A useful way to interpret this construction is to think of ϵ_{ij} as capturing idiosyncratic factors that affect the attractiveness of location j for individual i—for example, proximity to family, a spouse's job prospects, health considerations of oneself or one's relatives, access to a specialized school or medical provider, or other life circumstances that make one place especially appealing. In the baseline specification $u_{ij} = w_j \epsilon_{ij}$, these same idiosyncratic factors mechanically scale the marginal utility of income in that location, since $\partial u_{ij}/\partial w_j = \epsilon_{ij}$. In other words, the very forces that make a location more likely to be chosen also raise the marginal utility of income *in that location.*

The $D(\cdot)$ device breaks this tight link by allowing the factors that shift location choice to also be correlated with a separate shifter of marginal utility that is not location-specific. For instance, one convenient parameterization is

$$D(\epsilon_{i1}, \ldots, \epsilon_{iJ}) = \prod_{k=1}^{J} \epsilon_{ik}^{\phi_k},$$

so that the marginal utility in any location is scaled by $\prod_{k=1}^{J} \epsilon_{ik}^{\phi_k}$. If $\phi_k < 0$, then realizations in which ϵ_{ik} is high are precisely those in which this shifter is low (and hence marginal utility is low), whereas if $\phi_k > 0$ then high ϵ_{ik} realizations coincide with high marginal utility.

This captures the idea that the circumstances that make someone especially likely to choose a particular location might coincide with periods when additional income is especially valuable regardless of where the person lives (e.g., financial stress, caregiving responsibilities, or temporary liquidity needs), or with periods when additional income is less valuable regardless of location. Depending on the direction of this correlation, conditioning on the event that location j is chosen can select individuals into systematically higher- or lower-marginal-utility situations, even though the choice behavior itself is unchanged.

For each i, the second model scales the entire vector of utilities by a positive factor and therefore implies the same argmax and the same location choices. However, marginal utilities differ. In the first model,

$$\mu_j^{(1)} \equiv \left[\frac{\partial u_{ij}}{\partial w_j} \,\middle|\, j \text{ chosen} \right] = \left[\epsilon_{ij} \,\middle|\, j \text{ chosen} \right],$$

while in the second model,

$$\mu_j^{(2)} \equiv \left[\frac{\partial u_{ij}}{\partial w_j} \,\middle|\, j \text{ chosen} \right] = \left[\tilde{\epsilon}_{ij} \,\middle|\, j \text{ chosen} \right] = \left[D(\epsilon_{i1}, \ldots, \epsilon_{iJ})\, \epsilon_{ij} \,\middle|\, j \text{ chosen} \right].$$

Unless $D(\cdot)$ is almost surely constant, $\mu_j^{(2)}$ generically differs from $\mu_j^{(1)}$. Thus location choice data alone do not identify the level or pattern of marginal utility of income across locations. Put differently, exogenous wage variation is sufficient to identify location-demand responses (how choice probabilities shift with financial incentives), but it is not sufficient to pin identify the shape or properties of the $D(\cdot)$ function—that is, how the forces that shift location choice are correlated with marginal utility of income.

3.3.2 Example 2: Shocks that are perfect substitutes for income

A second class of non-separable specifications treats income and idiosyncratic location factors as (perfect) substitutes by combining them inside a common index of "income-equivalent" resources. In this class, the idiosyncratic factors that drive location choice enter utility in the same way as income, so choice behavior is governed by which location delivers the highest value of this composite index. The example below shows that, even when this index is pinned down by observed choices, location choice data need not reveal how the idiosyncratic component of the index maps into the policy-relevant conditional marginal utilities μ_j.

Let both locations have the same deterministic income level w, and consider the pair of models:

$$\text{Model A:} \quad u_{i1} = u\big(w + \epsilon_{i1} - \epsilon_{i2}\big), \qquad u_{i2} = u(w),$$

$$\text{Model B:} \quad u_{i1} = u(w), \qquad u_{i2} = u\big(w + \epsilon_{i2} - \epsilon_{i1}\big),$$

where $u(\cdot)$ is strictly increasing.

A useful interpretation is that $\epsilon_{i1} - \epsilon_{i2}$ is the individual's net, income-equivalent advantage of location 1 relative to location 2. For example, it could capture differences in childcare costs due to nearby family, commute-time savings, school-district differences that would otherwise require private-school tuition, spouse-job considerations, medical spending driven by provider access, or local tax and insurance differences. The key assumption in this example is that these forces enter utility as an additive dollar-equivalent term inside the same composite index as income.

In Model A, individual i chooses location 1 if and only if

$$u\big(w + \epsilon_{i1} - \epsilon_{i2}\big) \;\geq\; u(w) \quad \Longleftrightarrow \quad \epsilon_{i1} \;\geq\; \epsilon_{i2},$$

where the equivalence uses only that $u(\cdot)$ is increasing. In Model B, the choice rule for location 1 is

$$u(w) \;\geq\; u\big(w + \epsilon_{i2} - \epsilon_{i1}\big) \quad \Longleftrightarrow \quad \epsilon_{i1} \;\geq\; \epsilon_{i2}.$$

Thus the two models generate identical choices and identical choice probabilities for every distribution of $(\epsilon_{i1}, \epsilon_{i2})$: observationally, both models say that what matters for location choice is whether the income-equivalent advantage of location 1 relative to location 2 is positive.

Yet the models imply different marginal utilities across locations. The key distinction is where the idiosyncratic component of the income-equivalent index is "loaded." In Model A, the index varies in location 1 while location 2 delivers the fixed value w; in Model B, the index varies in location 2 while location 1 delivers w. This matters whenever marginal utility depends on the realized value of the index. In particular, if $u(\cdot)$ is concave, then $u'(\cdot)$ is decreasing, so the location whose composite index is shifted upward for the chosen individuals will tend to have *lower* marginal utility of income.

Formally, in Model A we have

$$\mu_1^A = \left[u'\big(w + \epsilon_{i1} - \epsilon_{i2}\big) \,\Big|\, 1 \text{ chosen}\right], \qquad \mu_2^A = u'(w),$$

while in Model B the roles reverse:

$$\mu_1^B = u'(w), \qquad \mu_2^B = \left[u'\big(w + \epsilon_{i2} - \epsilon_{i1}\big) \,\Big|\, 2 \text{ chosen}\right].$$

Both models fit the same location choice behavior, but they can imply opposite conclusions about which location has higher marginal utility of income. Put differently, observing where people locate can identify which location tends to deliver the higher income-equivalent index, but without additional restrictions it does not identify how the idiosyncratic component of that index maps into the conditional marginal utilities μ_j that enter optimal policy.

3.4 Implications

The non-identification results in this section clarify what can, and cannot, be learned from location choice data for designing redistribution-motivated place-based transfers. Even under separability, the revealed-preference link from wage variation to marginal utilities requires a rich set of *cross*-location responses that is rarely feasible to estimate in practice. More fundamentally, once separability is relaxed, income shifters identify how wages change location *attractiveness* for individuals at the margin of moving, but they do not identify the policy-relevant conditional averages $\{\mu_j\}$ among all residents, including the inframarginal, that enter the planner's first-order conditions. The observational-equivalence examples show that this is not merely a power or instrument-availability issue: without additional restrictions, distinct preference specifications can fit the same location-choice responses while implying different $\{\mu_j\}$.

These results motivate moving beyond revealed preference with location choice data. Our survey directly elicits within-person comparisons of the value of additional income across plausible future locations (to measure $\{\mu_j\}$), while still collecting information on how location choices probabilities respond to experimentally varied location incomes that is appropriate to identify location elasticities and cross-location substitution patterns.

4 Data

4.1 Survey overview and structure

This section describes the survey data used throughout the paper and the external data sources used for validation exercises. Our main dataset is a nationally representative survey that we designed to measure, at the individual level, both stated location choice behavior and within-person marginal-utility comparisons across locations. Table 1 provides a schematic overview of the survey modules and the outputs used in the analysis. The survey proceeds in three steps. First, respondents specify a respondent-specific menu of four metropolitan areas where they could plausibly live in five years and report baseline attributes for each location, including expected household income and whether family or friends would be present. Second, respondents report baseline choice probabilities across the four locations and answer a marginal-utility elicitation block that asks respondents to compare the impact on well-being of permanent income increases experienced in different locations. Third, respondents repeat the choice-probability and marginal-utility blocks under eight experimentally varied scenarios that shift location-specific income and family/friends presence while instructing respondents to hold all other location attributes fixed. These scenario responses identify (i) location demand elasticities and substitution patterns and (ii) how marginal utility varies across locations within individuals and as incomes change, which we aggregate to construct the metro-level objects used in our optimal-transfer computations.

Table 1: Survey flow and what each module provides

Module	What respondents do	What it provides / how used
Background (current/past)	• Current MSA; hometown MSA • Current income (indiv/HH)	• Descriptives + validation inputs • Baseline income measures for the choice menu
Future location menu (4 MSAs)	• Choose 4 feasible MSAs: – hometown – 3 distinct destinations (not current)	• Individual menu $\mathcal{M}_i$ • Which locations are considered together (substitution)
Baseline attributes by MSA	• For each $j \in \mathcal{M}_i$: – expected HH income (5y) – family/friends present (0/1)	• Baseline covariates (validation + correlates) • Inputs to scenario perturbations
Rankings by MSA	• Rank the 4 MSAs by: – support from family/friends – fun / hobbies fit – expensiveness (COL + desired housing)	• Validation vs external data: – expensiveness vs housing prices/costs – fun vs amenity indices (and crime)
Baseline choice probabilities	• Report p_{ij0} over 4 MSAs (sum to 1)	• Baseline choice weights for aggregation • Menu-based substitution information
Baseline MU elicitation	• Pick MSA where $5k is most valuable • Set compensating differentials in other MSAs	• Within-person MU comparisons across MSAs • Relative MU (up to person scale)
Scenario block (8 scenarios)	• Instruction: hold all else fixed within each MSA • Randomly vary within menu: – income by MSA – family/friends by MSA • In each scenario, collect: – choice probabilities p_{ijk} – MU block (location of $5k, compensating differentials)	• Location-demand responsiveness: – income elasticity – family/friends effect • Own/cross substitution patterns across locations • Panel MU responses to income and family/friends shifts
Demographics module	• Age, gender, race/ethnicity, education, household, employment	• Sample description; ACS comparison; Pareto weighting by skill group

Table 1 summarizes the survey modules and the objects we use in the analysis. The remainder of this section provides additional detail on the sample and on the key variables generated by each module. In particular, we first describe the analysis sample and its representativeness, then define the main variables and experimental variation that identify our location-demand and marginal-utility estimates.

4.1.1 Eliciting choice probabilities and the marginal utility

The key part of our survey designed to elicit choice probabilities and MUC is divided into their two broad respective parts. Before this, each respondent is asked to report their hometown where they grew up and three other locations that they think they are most likely to live in the next 5 years. Next, respondents are asked to report their location-specific expected annual income, presence of family and friends, their ranking of perceived expenses and ease of pursuing their hobbies and habits, and the probability of choosing each location in the next 5 years. We call this to be the baseline scenario.

Then we present respondents with a series of scenarios where we exogenously shock their reported location-specific expected income and presence of family and friends, and ask them to report their choice probabilities across the four locations. This completes the first part of the survey. In the second part of the survey, we ask the respondent to imagine different counterfactual worlds where they have chosen to live in each of the four locations. Then we ask them that suppose they receive an unanticipated permanent income shock, in which location would that permanent jump in income be the most valuable to them. Once they choose a location, we ask them to report the compensating differential of a permanent income shock that would make them just as happy as they would be in the chosen location. We repeat this exercise over the same set of exogenously varied scenarios as before and ask them to report their compensating differentials.

Thus, for each respondent, we have a panel data with 24 linearly independent observations on choice probabilities and 24 linearly independent observations on compensating differentials in response to exogenous shocks to their location-specific incomes and presence of family and friends, in addition to their baseline responses.

4.1.2 Key highlights of our survey design

One major difference that our survey design has compared to the literature using hypothetical choice surveys for identification is that our identification exploits variation within each location between scenarios for each individual and not within scenarios which imposes stronger assumptions on unobservable location-specific heterogeneity. Specifically, we ask individuals to assume that given a location, any unobserved factor remains the same across scenarios. Thus, we can allow for unobserved differences across the locations which are held constant across scenarios. This is weaker than assuming that everything else between the locations within each scenario is held fixed, thereby instructing the respondents to make choices based on the differences in the observed characteristics of the locations. Additionally, by allowing for unobserved differences across locations, we reduce the cognitive load on respondents, and make the identification strategy more robust to variations in unobserved factors. For completeness, we are able to do this because the choice sets involve MSAs which are well-defined unlike other surveys that use hypothetical choice experiments to elicit preferences for objects in choice-sets which can have varied definitions and thus can be more difficult to define and compare, for example hypothetical jobs.[4]

We also differ in the way we induce experimental variation in our survey. Instead of providing quasi-randomly drawn incomes, we draw shocks following the distribution in income documented in Blundell et al. (2008) and use those shocks to exogenously vary their expected location-specific income across different scenarios. To exogenously vary the expected presence of family and friends we quasi-randomly draw zeros and ones for each of the four locations.[5] The advantage of this approach is that the exogenously shocked incomes are specific to the individual-location pair and are thus more realistic than the quasi-randomly drawn incomes for each of the four locations.

4.2 Key variables and scenario design

For each respondent i, the survey elicits an individualized choice set (menu) of four metropolitan areas, $\mathcal{M}_i = \{j_{1(i)}, j_{2(i)}, j_{3(i)}, j_{4(i)}\}$ consisting of the respondent's hometown and three additional destinations. The survey then collects (i) stated location choice probabilities under a baseline scenario and eight experimentally varied scenarios, and (ii) within-person marginal-utility comparisons under the same scenarios.

4.2.1 Baseline variables

For each i and each location $j \in \mathcal{M}_i$, the baseline module records:

1. Expected household income five years ahead, w_{ij0}.
2. An indicator for whether family or close friends would be present, $f_{ij0} \in \{0, 1\}$.
3. Stated choice probabilities over the menu, $\{p_{ij0}\}_{j \in \mathcal{M}_i}$, with $\sum_{j \in \mathcal{M}_i} p_{ij0} = 1$.
4. Subjective rankings of each location by expensiveness, fun / hobbies fit, and support from family/friends.

4.2.2 Scenario variation

Respondents then complete eight additional scenarios indexed by $k = 1, \ldots, 8$ (with $k = 0$ denoting the baseline). In each scenario, we present modified values of income and family/friends presence for each $j \in \mathcal{M}_i$,

$$w_{ijk}, \qquad f_{ijk} \in \{0, 1\},$$

[4] For example Wiswall and Zafar (2018) use hypothetical choice experiments to identify preferences over hypothetical job attributes, Alam et al. (2021) embed information experiments within hypothetical choice surveys to separately identify beliefs on gender gaps in manager quality, and Andrew and Adams (2025) identify beliefs on future marriage market matches.

[5] Note that the variations we induce are quasi-random and not completely random. This is done to avoid draws in scenarios which can lead to choices being strictly dominant in all attributes thereby rendering variations in those scenarios to be ineffective following Wiswall and Zafar (2018) and Alam et al. (2021).

and instruct respondents to hold all other aspects of each location fixed relative to baseline. Respondents report updated choice probabilities $\{p_{ijk}\}_{j\in\mathcal{M}_i}$ under each scenario. The scenario variation in w_{ijk} is generated by applying multiplicative income shocks to each individual–location baseline value, and the variation in f_{ijk} is generated by quasi-randomly toggling the presence of family/friends across locations and scenarios.[6]

4.2.3 Marginal-utility elicitation and constructed objects

In each scenario $k = 0, 1, \ldots, 8$, respondents answer two marginal-utility questions. First, they indicate the location in which an unexpected and permanent \$5,000 increase in household income would be most valuable, conditional on having chosen to live in that location. Second, for each of the other three locations, they report the permanent income increase that would deliver the same change in well-being as the \$5,000 increase in the chosen location. Let b_{ijk} denote this compensating differential for location j in scenario k (with $b_{ijk} = 5{,}000$ for the chosen location by construction). We define

$$muc_{ijk} \equiv \frac{1}{b_{ijk}},$$

which is proportional to the respondent's marginal utility of income in location j in scenario k, up to a respondent–scenario-specific scale factor. Our empirical specifications exploit within-respondent comparisons across locations and within-respondent variation across scenarios to eliminate these scale factors.

4.3 External data sources for validation tests

We supplement our survey data with three other external data sources for validation checks and to construct location-specific covariates used in various analyses. The first source is the 2022 American Community Survey (ACS), providing demographic and economic indicators including median house prices, wage indices, and housing cost measures. The amenity indices developed by Diamond (2016), which quantify location-specific amenity indices including school quality, retail access, crime rates, congestion measures, environmental quality, and non-wage job attributes, are the second source. The third source is the Zillow Home Value Index (ZHVI), which provides monthly data on median home prices for each MSA.

Using these additional sources of data we conduct multiple validation tests in Section 5.2.1 to check the internal consistency of our survey responses. We first check whether the expected location-specific income reported by respondents is positively correlated with a skill adjusted location-specific wage index that we construct from the ACS. We also examine the correspondence between subjective rankings of location "expensiveness" and median house prices, and subjective rankings of location "fun-ness" and the amenity indices.

[6] We use quasi-random assignment rules to avoid scenarios with strictly dominated options, following the stated-preference design literature. Appendix C provides the full details of scenario construction and randomization.

5 Survey validation and location-demand estimates

5.1 Analysis sample and representativeness

Our survey design elicits information about location preferences and the marginal utility of income across different locations. We collected responses from 2,193 individuals with non-missing data on key questions, combining data from our second pilot and final data collection.

Table 2 compares the composition of our analysis sample to the American Community Survey (ACS) along several core dimensions. The survey sample is broadly similar to the ACS along Census divisions and age bins, with modest differences in gender and race shares. The most notable differences are in education, where the survey sample has a higher share of four-year college graduates and a lower share in "some college," as well as a smaller share in the lowest education category. These comparisons are meant to be a transparent diagnostic of representativeness rather than a requirement for identification, since our key estimates exploit within-respondent variation across experimentally varied scenarios. Applying survey weights that balance the joint distribution of our sample to the ACS benchmark has negligible effects on our estimated location-choice elasticity and on our survey-based measures of marginal utility differences across locations.

5.1.1 Spatial distribution of respondents and their future potential locations

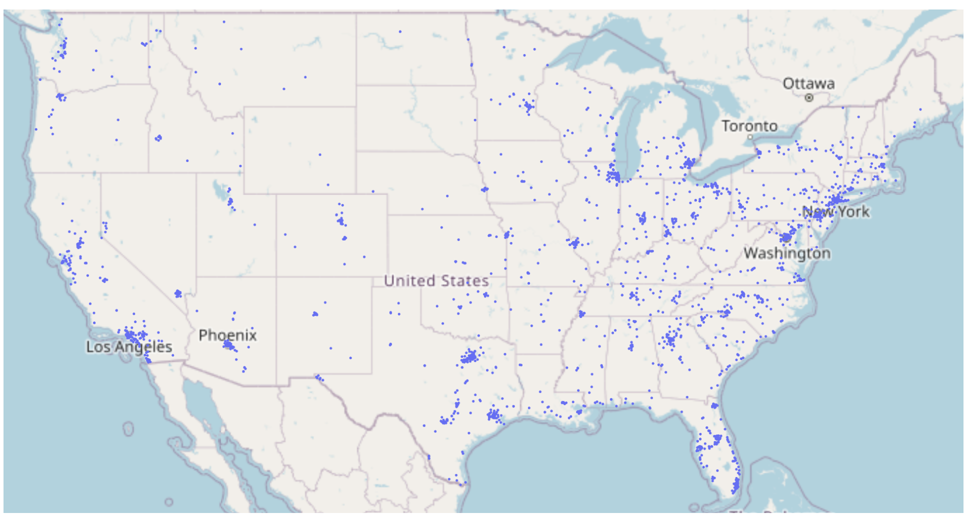

Figure 1: Location of Respondents

The geographic distribution of respondents covers a wide range of metropolitan areas across the United States as shown in Figure 1. Respondents listed a diverse set of potential future metropolitan areas in Figure 2 with darker shaded MSAs indicating locations that were chosen more than those shaded lighter. The spatial distribution of these destinations reflects both the geographic dispersion of our sample shown in the previous figure and their intended migration patterns. Major coastal metropolitan areas feature prominently, though respondents also identify a substantial number of interior MSAs.

Table 2: Sample composition relative to the ACS

	ACS	**Survey**
Census division		
Pacific Division	0.159	0.136
Mountain Division	0.076	0.051
West South Central Div.	0.122	0.121
East South Central Div.	0.058	0.057
South Atlantic Division	0.206	0.196
West North Central Div.	0.064	0.066
East North Central Div.	0.141	0.172
Middle Atlantic Division	0.127	0.158
New England Division	0.047	0.043
Education		
Doctoral Degree	0.015	0.022
Professional Degree (JD, MD, MBA)	0.021	0.018
Master's Degree	0.092	0.075
4-year College Degree	0.207	0.281
2-year College Degree	0.085	0.072
Some College	0.205	0.147
High School degree / GED	0.271	0.264
Some High School	0.060	0.112
Primary education or less	0.043	0.008
Gender		
Male	0.510	0.433
Female	0.490	0.562
Hispanic origin		
Hispanic	0.176	0.156
Not Hispanic	0.824	0.844
Race		
Other	0.069	0.036
Mixed race	0.111	0.024
Native American	0.010	0.019
Asian / Pacific Islander	0.064	0.057
Black	0.118	0.146
White	0.629	0.719
Age		
65+ years old	0.226	0.209
55–64 years old	0.159	0.192
45–54 years old	0.154	0.187
35–44 years old	0.171	0.191
25–34 years old	0.173	0.113
18–24 years old	0.116	0.107

Notes: Entries report shares. The ACS column reports shares from the American Community Survey. The MUC column reports shares in the survey analysis sample used for the marginal-utility and location-demand estimates. Shares within each block sum to 1.

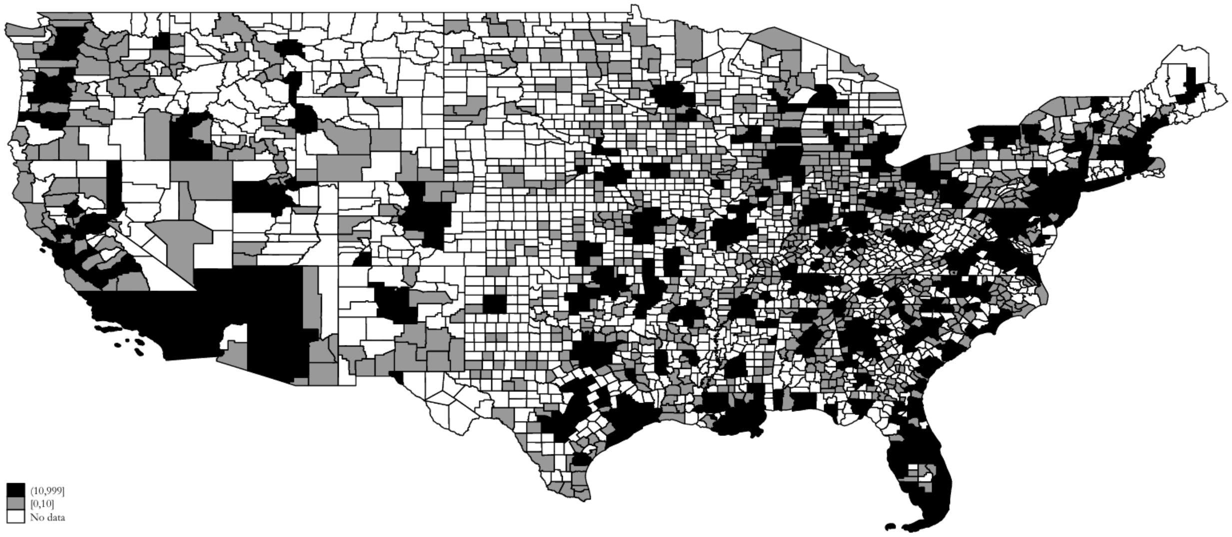

Figure 2: Potential Future Locations

5.2 Patterns in our Survey Data

5.2.1 Validation of survey responses

We conduct several validation exercises to confirm that respondents provided substantively meaningful information.

First, we examine the relationship between reported location-specific expected income and metro-level wage index. Specifically, using respondents' baseline responses on their location-specific expected income, we estimate

$$ln(ExpInc_{ij}) = \alpha_0 + \alpha_1 CensusWageIndex_j + \mu_i + e_{ij}$$

where μ_i is a respondent fixed effect and the Census wage index is the skill-adjusted hourly wage index for the location j that we estimate from the American Community Survey (ACS).[7]

[7] To construct a location-specific wage index adjusted for skill, we first regress log hourly wages on demographic characteristics (age, education, race, sex, and marital status) and metropolitan area fixed effects using the 2022 ACS 5-year sample, restricting to full-time workers (working more than 35 hours per week). The resulting metropolitan area fixed effects represent skill-adjusted average hourly wages that account for demographic composition differences across locations, providing a measure of location-specific wage premiums.

Dependent Variable: Log(expected household income)	
Variable	**Coefficient**
Census Wage index (ACS)	0.421
	(0.017)
Constant	4.330
	(0.001)
Observations	64,503
R^2	0.842
Individual FEs	Yes

Robust standard errors in parentheses.

Our estimates reveal a statistically significant positive relationship between expected income and skill-adjusted metro hourly-wages. Indeed, respondents anticipate higher incomes in metropolitan areas with higher wage indices. This indicates that respondents anticipate systematic income differentials across metropolitan areas in patterns that align with observed wage variation, suggesting that income expectations align with regional labor market conditions.

Next, we show that reported probabilities of any move in the next five years strongly align with established life-cycle patterns, with moving probabilities declining monotonically with age shown in the binned scatterplot of Figure 3. Young adults (18-24) report approximately 60% probability of moving within five years, declining systematically to approximately 20% for respondents over 65. This pattern is consistent with existing literature on migration patterns, confirming that our survey captures expected life-cycle dynamics in location choice.

Respondents' rankings of metropolitan area characteristics also show strong correlation with objective measures. Binned scatterplots in Figure 4 reveal that when respondents rank candidate locations by "expensiveness," these rankings show a strong positive correlation with median house prices from ACS data, with median house prices monotonically decreasing from the most to least expensive ranked locations.[8] Similarly, binned scatterplots of rankings of metropolitan "fun-ness" in Figure 5 demonstrate positive association with objective measures of amenities, such as environmental quality indices that incorporate per-capita park spending and air quality measures obtained from Diamond (2016).[9] These same "fun-ness" rankings correlate negatively with metropolitan crime indices obtained from Diamond (2016) and shown in Figure 6, indicating that respondents perceive locations with lower crime rates as more enjoyable.

[8] We ask "Given the sorts of goods and services you like to buy and given the kind of housing unit you wish to live in, which locations do you think would be most expensive for you?" and then ask the respondent to rank the likely future locations from 1 (most expensive) to 4 (least).

[9] We ask "Thinking about your hobbies and things you like to do for entertainment, where do you think you will have the most fun?" and then ask the respondent to rank the likely future locations from 1 (most fun) to 4 (least).

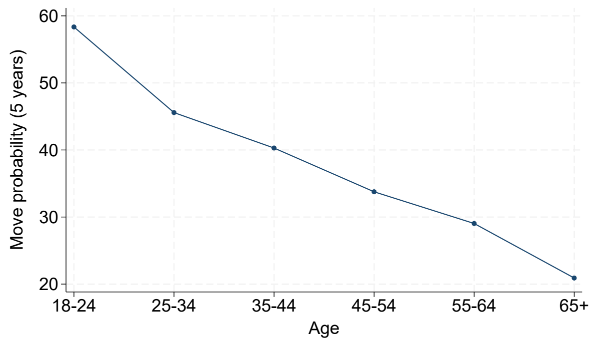

Figure 3: Reported chance of any move in the next 5 years falls with age

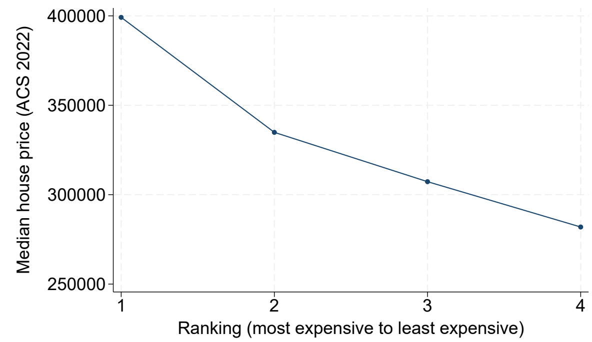

Figure 4: Reported rankings align with metro median house prices

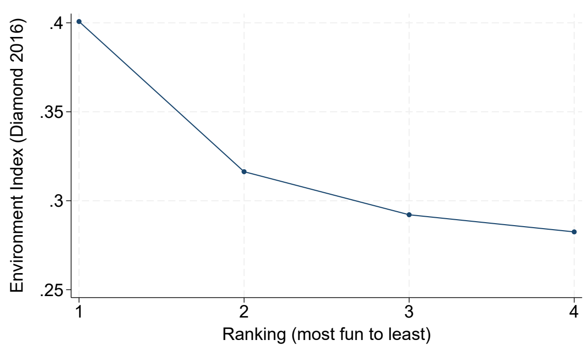

Figure 5: Reported rankings align with park spending / air quality

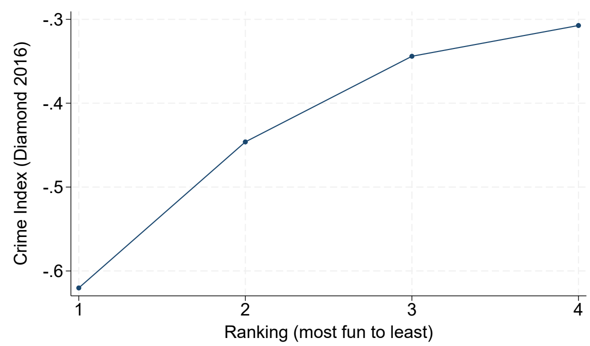

Figure 6: Reported rankings align negatively with metro crime rates

5.3 Elasticities of location choice

The experimental variation we induce over multiple scenarios in our survey allows us to identify the elasticity of location choice with respect to income. Specifically, our instructions in the survey guide respondents to hold location specific unobservable constant across scenarios, while we exogenously vary the expected income and the presence of family and friends in each of the four locations across eight scenarios. Thus, conditional on an individual-location fixed effect and an individual-scenario fixed effect, variation in expected income and the presence of family and friends across scenarios allows us to identify the elasticity of location choice with respect to income and the presence of family and friends.

We note that choice probabilities are typically responded in multiples of five and thus have measurement error associated with them; we use a least absolute deviation estimator (Blass et al., 2010).[10] Specifically, we consider the following reduced form specification:

$$\log\left(p_{ijk}\right) = \nu_1 \log\left(w_{ijk}\right) \;+\; \nu_2 f_{ijk} \;+\; \phi_{ik} \;+\; \phi_{ij} \;+\; u_{ijk} \tag{5}$$

[10] The estimator is not sensitive to small perturbations of extreme probabilities of 0 and 1.

where p_{ijk} is the location choice probability of household i in location j in scenario k, log income is $\log(w_{ijk})$), f_{ijk} is an indicator variable for the presence of family and friends, ϕ_{ik} are person-scenario fixed effects and ϕ_{ij} are person-location fixed effects.

Taking into account measurement error in reported choice probabilities, and using experimental variation of our survey allows identification through the assumption of $\mathbb{M}[u_{ijk} \mid w_{ijk}, f_{ijk}, \phi_{ik}, \phi_{ij}] = 0$. However, since the conditional median independence assumption conditions on fixed effects as well, we estimate the above equation using a two-step quantile regression approach following Canay (2011). In the first step we estimate the fixed effects ϕ_{ik} and ϕ_{ij} using a simple fixed effect estimator for equation (5), and residualize the estimates $\widehat{\phi}_{ik}$ and $\widehat{\phi}_{ij}$ from the log choice probabilities $log(p_{ijk})$. In the second step we estimate the quantile regression of the residualized log choice probabilities on the experimentally induced variation in income and the presence of family and friends. The key assumption in this two-step estimator is that the fixed effects are simply shifters of all quantiles by the same amount. Then as the number of scenarios approaches infinity, the quantile regression of the residualized log choice probabilities on the experimentally varied income and the presence of family and friends yields consistent estimates of the elasticities of location choice with respect to income and the presence of family and friends. We block bootstrap the standard errors at the individual level.

We estimate an average location choice elasticity with respect to income (ν_1) of 1.69 (SE = 0.004), indicating substantial responsiveness of location decisions to income differences.[11] Our estimate is squarely in line with existing estimates from the literature, which typically range in between 1.2 and 3.1 depending on context and methodology as shown below.

To estimate the empirical distribution of the elasticities across respondents, we run the above estimation routine separately for each individual i. As scenarios for each individual approach infinity, the individual location choice elasticities are identified. However, since estimation uses 8 scenarios which gives us 24 linearly independent observations for each individual, the standard errors will be very high. Thus, to estimate the distribution of location choice elasticity we implement a parametric empirical Bayes (also known as a shrinkage estimator) approach which shrinks the individual estimates towards the mean of the distribution based on estimated signal-to-noise ratio. We explain the details of shrinkage procedure with the two-step estimator described above in Appendix D.[12] The graph below plots the shrinked distribution of location choice elasticities after implementing the empirical Bayes approach. The distribution is centered around 1.73, with a standard deviation of 0.55 which yields a standard error for the mean equal to 0.011.

[11] Our estimate of the location elasticity with respect to the presence of family and friends is $\nu_2 = 0.20$ with a standard error of 0.007, quantifying the non-monetary value of social connections in location choice.

[12] Also see Walters (2024) for a comprehensive review on parameteric and non-parametric empirical Bayes methods and application in labor economics.

Figure 7: Comparison of our estimate to the literature estimates

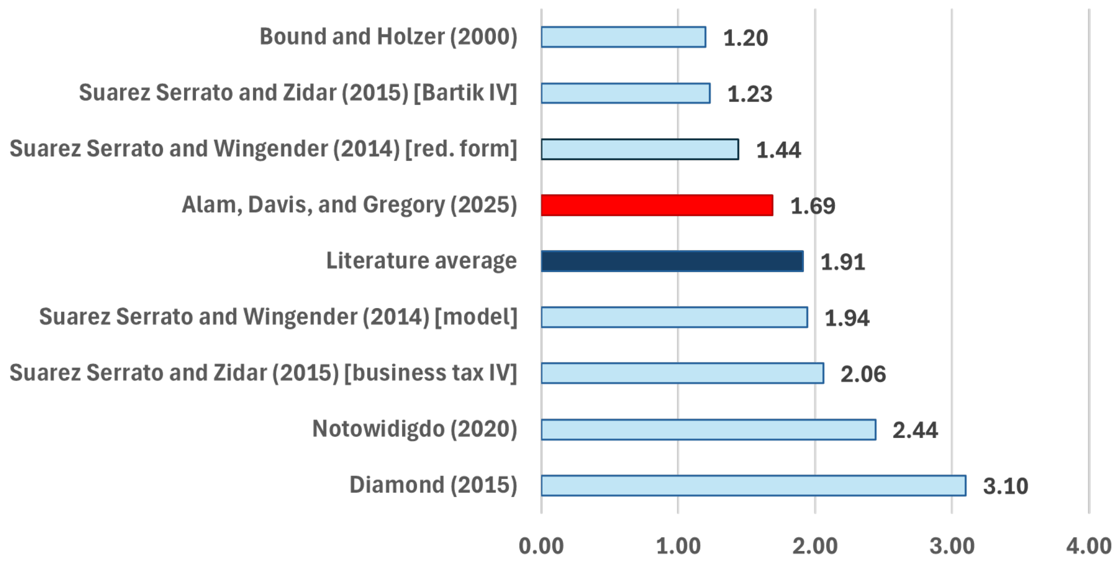

Notes: The estimates documented in the literature synthesized in Fajgelbaum et al. (2019) are from Bound and Holzer (2000), Serrato and Wingender (2014), Serrato and Zidar (2015), Notowidigdo (2011), and Diamond (2016) and highlighted in light blue. Their average is highlighted in dark blue. Our estimate is highlighted in red.

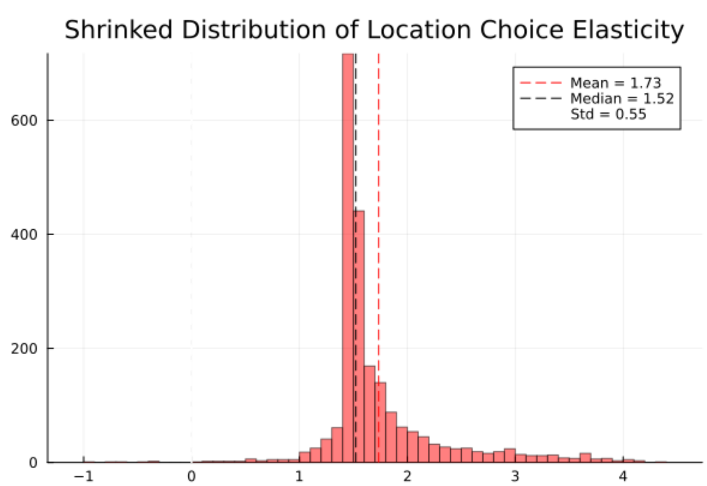

Figure 8: Distribution of location choice elasticity estimates

6 Empirical Patterns in the Marginal Utility of Income

In order to explain the patterns in the marginal utility of income across locations, we first show how the survey elicits data that can be mapped to the MUC within the context of a generic choice model. Hence, before we describe the empirical patterns in the location-specific MUC, we first provide a micro-foundation for construction of variables which can be used as measures of location-specific MUC.

6.1 Variable construction to utilize MUC

For the purposes of this section, abstract from the presence of family and friends. Denote the utility for household i from location j as

$$U_{ij}\left(C_{ij}, X_j, \epsilon_{ij}\right)$$

where C_{ij} is consumption, X_j are amenities of the location common to all people, and ϵ_{ij} are person-specific amenities.[13] The marginal utility of income is

$$MU_{ij} \quad = \quad \frac{dU_{ij}\left(C_{ij}, X_j, \epsilon_{ij}\right)}{dC_{ij}}$$

When asked about location where MUC is highest, we envision respondents comparing expected MUCs conditional on choice:

$$E\left[MU_{ij} \mid j \text{ chosen}\right] \quad = \quad E_{\epsilon_{ij}}\left[\frac{dU_{ij}\left(C_{ij}, X_j, \epsilon_{ij}\right)}{dC_{ij}} \mid u_{ij} = \max\left\{u_{ij'}\right\}\right]$$

The average of the marginal utility of income depends on the (a) the level of consumption C_{ij}, (b) the complementary of consumption with X_j and with ϵ_{ij}, and (c) the mix of ϵ_{ij} for which j is the chosen location. The correlation between the average of the marginal utility of income and the level of consumption in any given location can therefore be positive or negative.

When respondents answer this question:

> *In which location would you benefit the most from an unexpected and permanent increase of \$5,000 in household income, had you chosen to live there?*

we assume the respondent is making comparisons of the expected marginal utility of incom, conditional on choice. In other words, we assume the answer (j^*) satisfies this condition:

$$E\left[\frac{dU_{ij^*}\left(C_{ij^*}, X_{j^*}, \epsilon_{ij^*}\right)}{dC_{ij^*}} \mid j^* = \text{argmax}\right] \quad > \quad E\left[\frac{dU_{ij}\left(C_{ij}, X_j, \epsilon_{ij}\right)}{dC_{ij}} \mid j = \text{argmax}\right] \quad \forall j \neq j^*$$

Then, when respondents answer this question,

[13] The presence of X_j allows the marginal utility of income to vary across locations that offer the same consumption.

Now ... indicate the size of the jump in household income that they would need living in those locations 5 years from now such that their change in happiness after receiving this jump in household income is the same as the change in happiness from the $5,000 jump in household income ... in [location of previous answer].

we assume the response for each location $j \neq j^*$, which we denote as b_j, satisfies the following condition:

$$\mathbb{E}\left[\frac{dU_{ij}\left(C_{ij}, X_j, \epsilon_{ij}\right)}{dC_{ij}} \mid j = \text{argmax}\right] b_{ij} \quad = \quad \mathbb{E}\left[\frac{dU_{ij^*}\left(C_{ij^*}, X_{j^*}, \epsilon_{ij^*}\right)}{dC_{ij^*}} \mid j^* = \text{argmax}\right] \$5,000$$

For each household i, for all $j = 1, \ldots, 4$, and for any given scenario $k = 0, \ldots, 8$ (where $k = 0$ is the baseline) this enables us to compute for each j

$$\frac{1}{b_{ijk}} \quad = \quad \zeta_{ik}\ \mathbb{E}\left[\frac{dU_{ij}\left(C_{ij}, X_j, \epsilon_{ij}\right)}{dC_{ij}} \mid j = \text{argmax}\right] \tag{6}$$

where ζ_{ik} is a person- and scenario-specific constant. In other words, for each person and for each scenario, we can identify the marginal utility of income in three of the likely locations relative to the fourth. This implies that if we define $muc_{ijk} = 1/b_{ijk}$, then if we regress $log\left(muc_{ijk}\right)$ on person-by-scenario fixed effects, we can uncover how the marginal utility of consumption varies by location and directly estimate how income and other factors correlate with the variation.

6.2 Cross-Sectional Analysis of Marginal Utility at Baseline Wages

To estimate the variation in marginal utility of income (MUC) across locations, we implement two complementary approaches using our cross-sectional baseline data. The first approach non-parametrically estimates metropolitan-specific MUC values in an AKM framework Abowd et al. (1999), while the second examines systematic relationships between MUC and observable location characteristics using data from Diamond (2016) and the American Community Survey (ACS).

For our non-parametric estimation, we specify:

$$\ln(MU_{ij}) = \theta_i + \theta_{metro} + e_{ij} \tag{7}$$

where $\ln(muc_{ij})$ represents the log of the inverse required income boost in location j (relative to the highest-MUC location) for respondent i, θ_i captures individual fixed effects, and θ_{metro} represents metropolitan area fixed effects that constitute our parameters of interest. This approach, analogous to the AKM (Abowd, Kramarz, and Margolis 1999), allows us to obtain metropolitan-specific MUC estimates while controlling for respondent heterogeneity. Using estimates of θ_{metro}, we regress them on skill adjusted average hourly wage indices from the ACS described above to obtain a measure of the marginal utility of income in each metropolitan area.

Our second approach examines the relationship between MUC and observable metropolitan characteristics:

$$\log(MU_{ij}) = X'_{ij}\beta + \gamma_i + \varepsilon_{ij}, \tag{8}$$

where Z'_j represents a vector of metropolitan characteristics. This specification enables us to analyze systematic patterns in how MUC varies with location attributes.

6.3 Within-person association between marginal utility and metro characteristics

We study how marginal utility varies across locations within individuals. Specifically, for each individual i and location j in their choice set, we estimate the following within-person specification:

$$\log(MU_{ij}) = X'_{ij}\beta + \gamma_i + \varepsilon_{ij}, \tag{9}$$

where MU_{ij} denotes the elicited marginal utility of income in location j, X_{ij} is a vector of location characteristics evaluated at baseline, and γ_i is an individual fixed effect.

Table 3 reports the results. MU is weakly but positively correlated with location specific income. Specifically, all else equal, a 0.1 increase in log expected shocked income, corresponding to roughly a 10 percent increase in income is associated with an increase in marginal utility by a factor of $\exp(0.1 \times 0.053) \approx 1.0053$ implying a 0.5 percent increase in marginal utility. Similarly, marginal utility is negatively correlated with local housing costs. Holding other characteristics fixed, a 0.1 increase in the log house price index is associated with a reduction in marginal utility by a factor of $\exp(0.1 \times (-0.142)) \approx 0.986$, corresponding to a 1.4 percent lower marginal utility. Social proximity is also strongly associated with marginal utility. All else equal, the presence of family or friends in a location is associated with an increase in marginal utility by a factor of $\exp(0.134) \approx 1.143$, implying approximately a 14 percent higher marginal utility relative to otherwise comparable locations.

We also find locations ranked as less "fun"—based on respondents' hobbies and interests are associated with lower marginal utility relative to those locations which rank the highest. We also find that locations ranked as less "fun"—based on respondents' stated hobbies and interests—are associated with lower marginal utility relative to locations in the highest-ranked category. Relative to fun rank 1, locations in fun rank 2 are associated with marginal utility lower by a factor of $\exp(-0.173) \approx 0.84$. This gradient steepens monotonically across categories: fun rank 3 locations are associated with marginal utility lower by a factor of $\exp(-0.251) \approx 0.78$, while fun rank 4 locations are associated with marginal utility lower by a factor of $\exp(-0.327) \approx 0.72$. It is important to highlight that the non-wage location specific aspects are substantially correlated with marginal utility.

As a complementary specification, we also estimate a two-way fixed effects model with metro fixed effects to characterize geographic variation in marginal utility:

$$\underbrace{\ln(MU_{ij})}_{\ln(1/b_{ij})} = \underbrace{\theta_i}_{\text{Individual FE}} + \underbrace{\theta_{\text{metro}}}_{\text{Metro FE}} + e_{ij}, \tag{10}$$

Table 3: Within-person association between marginal utility and metro characteristics

	Log(MU)
Log(Expected shocked income)	0.053*
	(0.023)
Log(House price index)	-0.142***
	(0.019)
Presence of family/friends	0.134***
	(0.018)
log-Distance from hometown	-0.016***
	(0.003)
Fun rank: 1 (omitted)	
Fun rank: 2	-0.173***
	(0.019)
Fun rank: 3	-0.251***
	(0.018)
Fun rank: 4	-0.327***
	(0.018)
Constant	-7.482***
	(0.251)
Observations	8,350
Individual fixed effects	Yes

Notes: The table reports coefficients from a within-person regression of log marginal utility on metro characteristics. Individual fixed effects absorb time-invariant heterogeneity. Standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

which isolates persistent location-specific components of marginal utility net of individual-level heterogeneity.

Our non-parametric estimates reveal substantial heterogeneity in average MUC across metropolitan areas. We show the geographic variation in the estimated metropolitan fixed effects θ_{metro} in Figure 9 below.

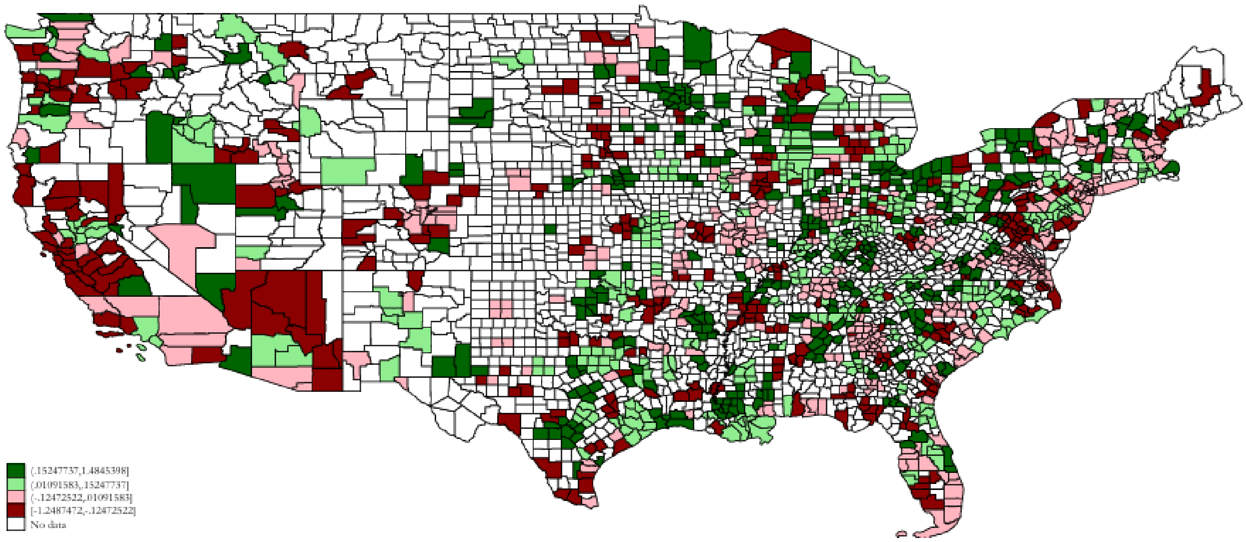

Figure 9: Estimated metropolitan fixed effects θ_{metro} from AKM model

Coastal metros and parts of California show low MU. Many metros in the south and southeast, parts of the Midwest have high MU with Florida and Texas being substantially heterogeneous having both high and low MU metros.

6.4 Constructing Comparable Marginal Utilities for Aggregation

Our survey elicits within-person comparisons of marginal utility across locations, but optimal policy requires a metro-level object μ_j that is comparable across respondents and incorporates the planner's Pareto weights. Two issues arise. First, within a respondent and scenario, our compensating-differential answers identify marginal utility only up to a respondent–scenario-specific scale factor. Second, respondents report different menus of feasible locations, so aggregation must account for the fact that the within-person comparisons are made over different choice sets.

We address these issues in two steps. Step 1 uses a within-respondent normalization that removes the respondent-specific scale. Step 2 uses the AKM metro effects to attach a common (cross-respondent) level scale to each respondent's menu, and then applies Pareto weights when aggregating to metros.

Step 1: Within-person normalization (removing the respondent-specific scale). Let $\mathcal{M}_i$ denote respondent i's menu. Let p_{ij0} denote respondent i's baseline stated choice probability for location $j \in \mathcal{M}_i$, and let b_{ij0} denote the corresponding baseline compensating differential (so

$1/b_{ij0}$ is proportional to marginal utility in j for respondent i). We define the within-person normalized marginal utility as

$$MU_{ij}^{within} = \frac{1/b_{ij0}}{\sum_{\ell \in \mathcal{M}_i} p_{i\ell 0} \left(1/b_{i\ell 0}\right)}. \tag{11}$$

By construction, this normalization sets the choice-probability-weighted mean within respondent i to one:

$$\sum_{\ell \in \mathcal{M}_i} p_{i\ell 0} MU_{i\ell}^{within} = 1.$$

Intuitively, MU_{ij}^{within} captures how high marginal utility is in location j relative to the other locations in i's menu, holding fixed the respondent-specific scale.

Step 2: Putting respondents on a common level scale using AKM metro effects. The normalization above delivers only relative marginal utilities within a menu. To aggregate across respondents with different menus, we attach a level scale to each respondent by using the AKM metro effects from (10). Let θ_ℓ^{metro} denote the estimated metro fixed effect, and define the implied metro-level marginal-utility index in levels as

$$\widetilde{MU}_\ell \equiv \exp\left(\theta_\ell^{metro}\right).$$

We then define a respondent-specific menu-level scale factor as the baseline choice-probability-weighted average of these metro indices:

$$MU_i^{loc} = \sum_{\ell \in \mathcal{M}_i} p_{i\ell 0} \, \widetilde{MU}_\ell. \tag{12}$$

This object captures the average level of marginal utility faced by respondent i across the locations they consider, on a common scale anchored by the estimated metro effects.

Metro-level aggregation with Pareto weights. Let $\Pi_{g(i)}$ denote the planner's Pareto weight for respondent i's type $g(i)$ (e.g., skill group). The planner-relevant marginal utility for respondent i in location j is proportional to

$$\Pi_{g(i)} \times MU_{ij}^{within} \times MU_i^{loc}.$$

Finally, the metro-level average marginal utility entering the optimal tax formula is constructed as the p_{ij0}-weighted average across respondents:

$$\mu_j = \frac{\sum_i p_{ij0} \left(\Pi_{g(i)} \times MU_{ij}^{within} \times MU_i^{loc}\right)}{\sum_i p_{ij0}}. \tag{13}$$

This aggregation uses p_{ij0} as the weight because the planner-relevant μ_j is an average marginal utility among residents of metro j, and in our stated-preference framework p_{ij0} is the analogue of the probability that respondent i resides in j under baseline conditions.

7 Geographic Substitution Patterns

Understanding how individuals substitute across locations in response to economic shocks is fundamental to designing optimal place-based policies. Traditional approaches to identifying substitution patterns rely on observational variation in wages or amenities, which are often confounded by unobserved factors that simultaneously affect multiple locations. Our survey design overcomes this challenge by creating exogenous within-person variation in location-specific attributes, allowing us to trace out the full matrix of substitution elasticities across metros. In this section, we first develop the theoretical framework that links our experimental design to substitution parameters, then present empirical estimates of these patterns, and finally examine specific geographic case studies that reveal how social attachments and cultural similarity shape migration responses.

7.1 Substitution Patterns from Menu-Based Scenarios

Our survey design allows us to identify substitution patterns through within-person variation in location-specific shocks. By varying income and social attachments across the respondent's feasible choice set, we observe how choice probabilities respond to location-specific changes. This approach addresses a fundamental identification challenge: in observational data, shocks to one location's attractiveness are often correlated with shocks to other locations through common factors such as regional business cycles, industry trends, or policy changes. Our experimental variation breaks these correlations by independently manipulating attributes across locations within each scenario.

The key innovation in our approach is the use of restricted choice sets or "menus" that vary across individuals. Each respondent evaluates scenarios involving their current location (or hometown if not currently residing there) and three alternative metros that represent their most likely relocation destinations based on demographic and geographic factors. This menu restriction serves two crucial purposes. First, it dramatically reduces the cognitive burden on respondents, who need only evaluate four locations rather than hundreds. Second, and more importantly for identification, it generates variation in which locations are treated as potential substitutes for each origin.

7.2 Computing Cross-Location Substitution Elasticities

A key contribution of our survey design is the ability to compute cross-location substitution elasticities using the individual-level choice probabilities and elasticities we recover from the experimental scenarios given the menu design which allows for different choice sets across individuals.

Menus and zero-probability convention. Each respondent i reports a menu of feasible locations $\mathcal{M}_i \subseteq \{1, \dots, J\}$ and a vector of stated choice probabilities $\{p_{ij}\}_{j \in \mathcal{M}_i}$ over that menu. For notational convenience, we extend $\{p_{ij}\}$ to the full set of metros by defining

$$p_{ij} = 0 \quad \text{for all } j \notin \mathcal{M}_i.$$

With this convention, sums/integrals can be taken over all $j \in \{1, \ldots, J\}$ without explicitly restricting to subgroups based on $\mathcal{M}_i$.

Let the aggregate population share in metro j be

$$\mathcal{P}_j = \int p_{ij}\, di.$$

where p_{ij} denotes individual i's baseline probability of choosing location j from their menu $\mathcal{M}_i$. When wages in metro j' change, the effect on metro j's population depends on both direct effects (for individuals who have both j and j' in their menu) and the individual-specific migration elasticities.

To connect the individual-level elasticities we estimate in the scenarios to the aggregate cross-metro responses, we use the standard logit-style approximation that, for a given individual i, the effect of raising income in location j' on the probability of choosing j is proportional to $p_{ij}p_{ij'}$. In particular, under a constant-elasticity-in-income approximation,

$$\frac{\partial p_{ij}}{\partial w_{ij'}} \;\approx\; -\,p_{ij}\,p_{ij'}\,\frac{\beta_i}{w_{ij'}}, \qquad j \neq j', \tag{14}$$

where β_i is individual i's (semi-)elasticity of location choice with respect to log income. Under the zero-probability convention ($p_{ij} = 0$ for $j \notin \mathcal{M}_i$), this derivative is automatically zero whenever j' is not on i's menu. Consequently, the share of metro j's population changes approximately by:

$$\frac{d\mathcal{P}_j}{dw_{j'}} \approx \int -p_{ij}\,p_{ij'}\,\frac{\beta_i}{w_{ij'}}\,di. \tag{15}$$

where β_i is individual i's location choice elasticity and $w_{ij'}$ is the wage individual i expects in location j' and $\mathcal{M}_i$ represents individual i's menu of locations. Under the zero-probability convention ($p_{ij} = 0$ for $j \notin \mathcal{M}_i$), individuals who do not consider j' automatically drop out of the integral because $p_{ij'} = 0$.

The cross-location substitution shown in Equation (15) depends on three factors: (1) the baseline probability of choosing the destination j, (2) the baseline probability of choosing the origin j', and (3) the individual's responsiveness to wage changes as captured by β_i. Importantly, if location j' is not in individual i's menu, then $p_{ij'} = 0$ and that individual contributes nothing to the cross-derivative, reflecting the fact that shocks to locations outside one's consideration set have no effect on choices.

To build intuition, consider a special case in which wages are uniform within metros ($w_{ij'} = \bar{w}_{j'}$), individuals share a common wage sensitivity ($\beta_i = \bar{\beta}$), but baseline choice probabilities vary due to heterogeneous attachments θ_{ij}. In this case, equation (15) simplifies to

$$\frac{d\mathcal{P}_j}{dw_{j'}} \approx -\left[\mathcal{P}_j\mathcal{P}_{j'} + \text{cov}(p_{ij}, p_{ij'})\right] \frac{\bar{\beta}}{\bar{w}_{j'}}, \tag{16}$$

The terms $\mathcal{P}_j$ and $\mathcal{P}_{j'}$ are the average probabilities of choosing locations j and j' respectively, and the term $\mathcal{P}_j\ \mathcal{P}_{j'}$ captures the cross-elasticity that would arise if individual location probabilities were common across individuals (as in a non-mixed logit model), or, weaker, independent across individuals. The second term, $\text{cov}(p_{ij}, p_{ij'})$, captures the additional substitution that arises from preferences for j and j' being correlated across individuals (a given person tends to either have a high preference for both or a low preference for both).[14]

Our menu approach combined with the experimental scenarios provides all the inputs needed to compute these elasticities. From the baseline survey responses, we obtain each individual's choice probabilities p_{ij} for all locations in their menu. From the experimental scenarios that vary wages and family/friend presence, we estimate individual-specific elasticities. Together, these allow us to construct the full matrix of cross-location substitution elasticities using equation (15).

Menus play a crucial role. By eliciting individual-specific choice sets, our design captures the heterogeneity in which locations are considered substitutes. Locations that rarely appear together in menus will have near-zero cross-elasticities, not because we assume they are poor substitutes, but because the data shows that few individuals consider both locations simultaneously. By limiting each individual to four locations, we also reduce computational burden while also generating realistic substitution patterns.

To illustrate how these cross-location elasticities vary geographically, Figure 10 presents the substitution patterns when Louisiana metros face a wage decrease (or equivalently, a tax increase). The map shows which states would receive marginal movers from Louisiana, with darker shading indicating higher substitution elasticities calculated using our formula from equation (15) and calculated as $\frac{d\mathcal{P}_{state'}/dw_{state}}{-d\mathcal{P}_{state}/dw_{state}}$ where $state =$ Louisiana. The pattern reveals strong regional clustering in substitution. Texas, Mississippi, and Florida receive the bulk of potential out-migrants, while there is minimal substitution to Northern or Western states except New York and California.

8 Computing Optimal Place-Based Transfers

The previous sections have established two key empirical findings: first, that the average marginal utility of income is positively correlated with metro income levels; and second, that substitution patterns across metros vary substantially. In this section, we show how to use these

[14] When this covariance is positive, the metros are stronger substitutes than their average choice probabilities $\{\bar{p}_j, \bar{p}_{j'}\}$ would suggest. In generic discrete choice product demand, this can occur when goods (here, locations) share common characteristics that appeal to similar types of people, such as cultural amenities, industry specializations, or social networks. In location demand, proximity to one's birthplace, college, or extended family is another "characteristic" that pairs of metros that are close together either both possess or both do not, which will tend cause geographically close-together metros to be closer substitutes.

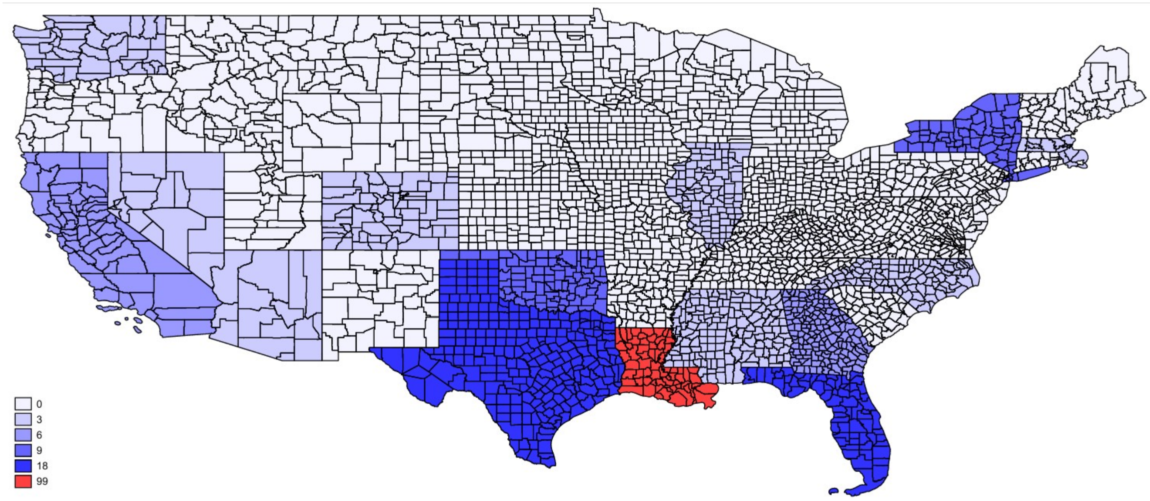

Figure 10: Substitution Patterns from Louisiana Metros

Notes: Shading indicates the share of marginal movers from Louisiana metros that would relocate to each destination state following a hypothetical tax increase. Darker shading represents higher substitution elasticities. The pattern shows strong regional clustering in the South, with Texas, Mississippi, and Alabama receiving the bulk of out-migrants. Calculated as: $\frac{d\mathcal{P}_{state'}/dw_{state}}{-d\mathcal{P}_{state}/dw_{state}}$ where $state$ = Louisiana.

sufficient statistics to feasibly compute the optimal place-based transfers at scale for all 623 metros in our sample. Recall, from Proposition 2.1 that optimal taxes solve

$$\tau_j = \underbrace{\frac{\sum_{j' \neq j} \frac{\partial \mathcal{P}_{j'}}{\partial \tau_j} \tau_{j'}}{\sum_{j' \neq j} \frac{\partial \mathcal{P}_{j'}}{\partial \tau_j}}}_{\substack{\text{Avg. metro tax faced} \\ \text{by those leaving } j}} + \underbrace{\frac{1}{\frac{\partial \mathcal{P}_j}{\partial \tau_j}} \frac{\mu_j - \bar{\mu}}{\bar{\mu}}}_{\substack{\text{MU deviation} \\ \text{scaled by mobility}}}$$

The optimal tax in metro j consists of two components. The first term is the weighted average tax faced by metros that would receive migrants if metro j raised its tax, capturing the fiscal externality of migration. The second term adjusts for differences in marginal utility: metros with above-average marginal utility receive subsidies (negative taxes), while those with below-average marginal utility face positive taxes. The adjustment is scaled by the elasticity of location choice, which governs the efficiency cost of redistribution.

8.1 Iterative Solution Method

While equation (3) characterizes the optimal taxes implicitly, solving for the tax vector $\boldsymbol{\tau} = (\tau_1, ..., \tau_J)'$ requires accounting for the interdependence across metros. Each metro's optimal tax depends on the taxes of potential substitute metros, which in turn depend on the original metro's tax. We develop an iterative algorithm that efficiently computes the fixed point of this system.

The system of first-order conditions can be written in matrix form as:

$$\mathbf{A}(\boldsymbol{\tau})\boldsymbol{\tau} = \mathbf{b} \tag{17}$$

where the elements of matrix $\mathbf{A}$ depend on the migration elasticities evaluated at the current tax vector, and vector $\mathbf{b}$ contains the MUC adjustments.

Specifically, the diagonal elements of $\mathbf{A}$ are:

$$A_{jj} = 1 + \frac{\sum_{k \neq j} \frac{\partial \mathcal{P}_k}{\partial \tau_j}}{\frac{\partial \mathcal{P}_j}{\partial \tau_j}} \tag{18}$$

and the off-diagonal elements are:

$$A_{jk} = -\frac{\frac{\partial \mathcal{P}_k}{\partial \tau_j}}{\frac{\partial \mathcal{P}_j}{\partial \tau_j}} \quad \text{for } k \neq j \tag{19}$$

The vector $\mathbf{b}$ has elements:

$$b_j = \frac{1}{\frac{\partial \mathcal{P}_j}{\partial \tau_j}} \cdot \frac{\mu_j - \bar{\mu}}{\bar{\mu}} \tag{20}$$

Our iterative algorithm proceeds as follows:

Algorithm: Iterative Computation of Optimal Taxes

- **Initialize:** Set $k = 0$ and $\boldsymbol{\tau}^{(0)} = \mathbf{0}$.
- **Repeat until convergence:**
  - Compute after-tax incomes $y_j^{(k)} = w_j - \tau_j^{(k)}$.
  - Compute location choice probabilities $\mathcal{P}_{ij}^{(k)}$ and population shares $L_j^{(k)} = \sum_i \mathcal{P}_{ij}^{(k)}$
  - Evaluate migration elasticities $\frac{\partial \mathcal{P}}{\partial \boldsymbol{\tau}}\big|_{\boldsymbol{\tau}^{(k)}}$
  - Update taxes by solving $\mathbf{A}(\boldsymbol{\tau}^{(k)})\boldsymbol{\tau}^{(k+1)} = \mathbf{b}$ using Gauss–Seidel:

$$\tau_j^{(k+1)} = \frac{1}{A_{jj}} \left(b_j - \sum_{m<j} A_{jm}\tau_m^{(k+1)} - \sum_{m>j} A_{jm}\tau_m^{(k)} \right).$$

  - Set $k \leftarrow k + 1$.
- **Stop** when $\|\boldsymbol{\tau}^{(k)} - \boldsymbol{\tau}^{(k-1)}\| < \epsilon$

Several implementation details ensure the algorithm's stability and economic interpretability. First, we impose income floors to prevent implausible poverty traps: pre-tax income must exceed \$30,000 for non-college graduates and \$60,000 for college graduates, while post-tax income must exceed \$10,000 for all individuals. These ensure that the optimal taxes do not generate extreme hardship while maintaining sufficient variation for meaningful redistribution.

9 Optimal Tax Results

9.1 Distribution of Optimal Metro Tax across Metros

Figure 11 presents the distribution of optimal metro-specific taxes across the 301 metropolitan areas in our sample. The distribution exhibits substantial dispersion, with a standard deviation of approximately \$17,200. Optimal taxes range from subsidies of approximately \$30,000 to positive taxes exceeding \$40,000.

Figure 11: Distribution of Optimal Taxes

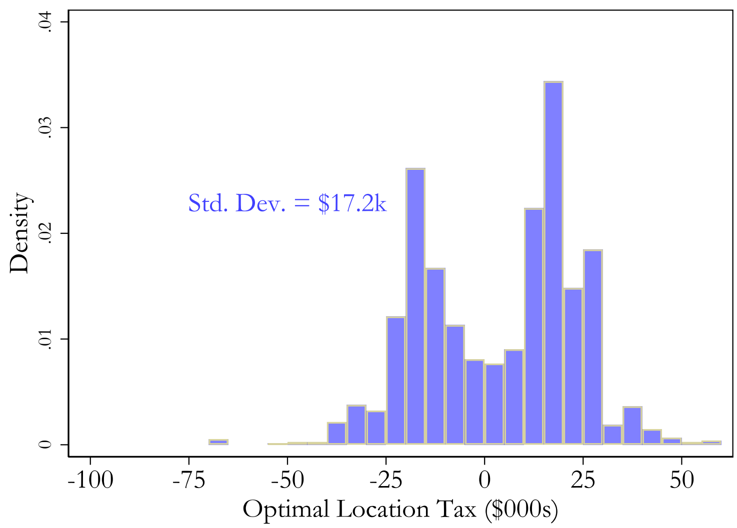

A notable feature of the distribution is its bimodality. One mode appears in the range of modest subsidies (approximately −\$25,000 to −\$10,000), while a second, more prominent mode emerges in the range of positive taxes (approximately \$10,000 to \$25,000). This bimodal structure suggests that location specific MU patterns and individual menus are such that there are broad groups of high and low MU metros, with relatively fewer metros near the break-even point of zero. To see this recall that the optimal tax for metro j can be expressed as:

$$\tau_j = \underbrace{\frac{\sum_{j'\neq j} \frac{\partial \mathcal{P}_{j'}}{\partial \tau_j} \tau_{j'}}{\sum_{j'\neq j} \frac{\partial \mathcal{P}_{j'}}{\partial \tau_j}}}_{\substack{\text{Avg. metro tax faced} \\ \text{by those leaving } j}} + \underbrace{\frac{1}{\frac{\partial \mathcal{P}_j}{\partial \tau_j}} \frac{\mu_j - \bar{\mu}}{\bar{\mu}}}_{\substack{\text{MU deviation} \\ \text{scaled by mobility}}}$$

The observed dispersion and bimodality therefore arise from two sources: (i) variation in the average marginal utility of income (μ_j) across metros relative to the population-weighted mean ($\bar{\mu}$), and (ii) variation in the taxes levied on substitute locations, weighted by migration responses. In the subsections that follow, we decompose this variation by examining how optimal taxes correlate with metro-level marginal utility, demographic composition, geographic location, and the characteristics of substitute metros. As we demonstrate, the variation and the bimodal structure can be largely attributed to systematic differences in marginal utility across metros—driven in part by differences in population composition by education—and the spatial clustering of similar metros that generates correlation between own taxes and the taxes of close substitutes.

9.2 Geographic variation in optimal taxes

Figure 12 maps the geographic distribution of optimal metro-level taxes across the continental United States.

Figure 12: Geographical variation in optimal taxes

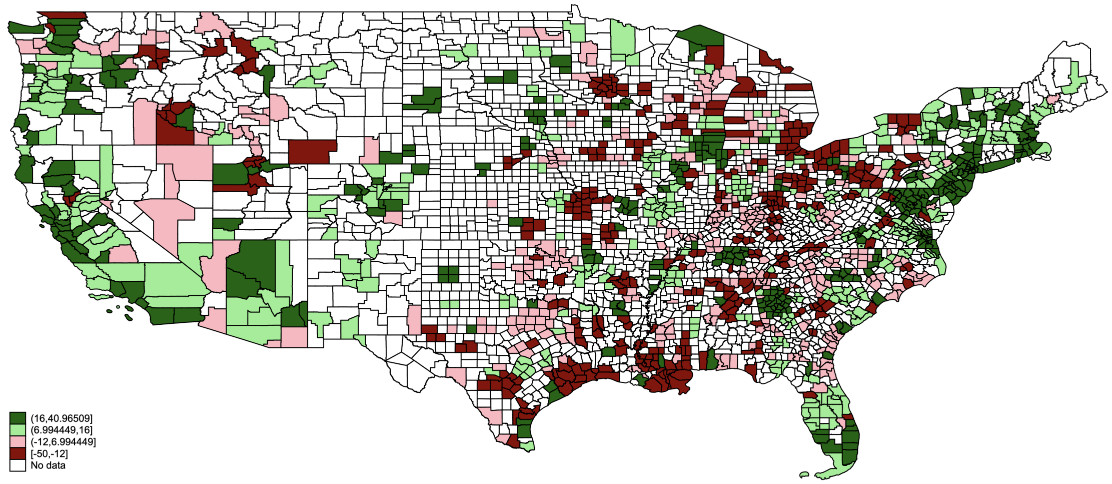

Notes: The color scheme distinguishes metros receiving subsidies (dark red: $[−50, −12] thousand; light red: $(−12, 7] thousand) from those facing positive taxes (light green: $(7, 16] thousand; dark green: $(16, 41] thousand). White areas indicate metros for which we lack sufficient survey coverage to compute optimal taxes.

The map reveals different geographic patterns that correspond to the bimodal distribution reported in Figure 11. Coastal metropolitan areas, particularly along the Eastern Seaboard from Washington, D.C. through New England, as well as parts of California and the Pacific Northwest, tend to face higher optimal taxes (green shading). In contrast, many metros throughout the South, parts coastal Southeast, portions of the Midwest, warrant subsidies (red shading). Florida and Texas exhibits heterogeneity, with both subsidized and taxed metros in close geographic proximity.

Importantly, this geographic pattern does not simply mirror the spatial distribution of wages or population density. Several large, high-wage metros in the South receive subsidies, while some smaller metros in the Northeast face positive taxes. This divergence from a naive "tax the rich cities" prescription reflects the primary contribution of our approach: optimal place-based transfers depend on measured marginal utility of income conditional on location choice, not merely on income levels. As documented in Section 6.2, marginal utility varies with local amenities, cost of living, and the presence of family and friends, individual perceived entertainment index—factors that need not move in lockstep with wages.

The spatial clustering visible in the map—whereby neighboring metros tend to exhibit similar tax assignments—foreshadows an important mechanism in our optimal tax formula. As we demonstrate in Section 7, metros that serve as close substitutes for one another tend to have correlated optimal taxes, generating the regional patterns observed here. The following subsections systematically decompose the sources of geographic variation by examining how optimal taxes relate to metro-level marginal utility, demographic composition, and the characteristics of substitute locations.

9.3 Optimal Metro Tax and (AKM) "Metro Effect" for MU

Figure 13 displays the relationship between optimal metro taxes and the metro fixed effects estimated from our two-way fixed effects specification: $\ln(MU_{ij}) = \theta_i + \theta_j^{\text{metro}} + e_{ij}$ where θ_i captures individual-specific factors affecting marginal utility and θ_j^{metro} isolates the location-specific component—the contribution of metro j to log marginal utility, holding individual characteristics constant.

Figure 13: Log-MU Index and Optimal Taxes

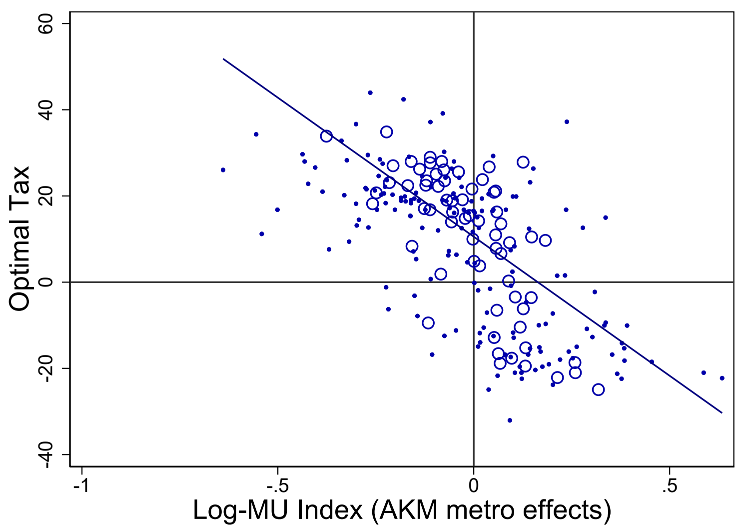

The scatter plot reveals a negative relationship between MU and optimal taxes. Metros with higher marginal utility effects receive lower optimal taxes, while metros with lower marginal utility effects face higher taxes. This pattern reflects the relative strength of the second term in our optimal tax formula:

$$\tau_j = \underbrace{\frac{\sum_{j' \neq j} \frac{\partial \mathcal{P}_{j'}}{\partial \tau_j} \tau_{j'}}{\sum_{j' \neq j} \frac{\partial \mathcal{P}_{j'}}{\partial \tau_j}}}_{\text{Avg. tax of substitutes}} + \underbrace{\frac{1}{\frac{\partial \mathcal{P}_j}{\partial \tau_j}} \frac{\mu_j - \bar{\mu}}{\bar{\mu}}}_{\text{MU deviation term}}$$

Since $\frac{\partial \mathcal{P}_j}{\partial \tau_j} < 0$ (raising taxes in a metro reduces its population share), the second term is negative when $\mu_j > \bar{\mu}$ and positive when $\mu_j < \bar{\mu}$. Thus, the MU deviation term pushes optimal taxes *below* the average tax of substitutes for high-MU metros and *above* the average tax of substitutes for low-MU metros. Whether a metro ultimately receives a subsidy or faces a positive tax depends on both its marginal utility deviation and the taxes levied on its substitute locations.

The negative slope reflects the redistributive intuition of optimal place-based policy. Holding individual characteristics constant, locations where residents derive greater marginal utility from additional income, represent opportunities for welfare-improving transfers. The AKM metro effects, by isolating location-specific contributions to marginal utility after controlling for individual unobserved heterogeneity through our survey design, thus serve as a key sufficient statistic for optimal spatial redistribution.

This relationship also helps explain the bimodal distribution documented in Figure 11 and the geographic clustering observed in Figure 12. In our data, we have high-MU locations (warranting subsidies) and low-MU locations (warranting taxes), with the strength of this separation generating the two modes. The geographic patterns arise because the determinants of marginal utility, including local amenities, cost of living, and access to family, exhibit spatial correlation.

The scatter plot also displays residual variation around the fitted line. This variation indicates that metro-level MU is not the sole determinant of optimal taxes. The first term of the formula i.e., the average tax faced by substitute metros also contributes. We now turn to examining this linkage across metros enabled by the menu-design of our survey which captures metro substitutes.

9.4 Optimal Metro Tax and Average Tax of Substitute Metros

Figure 14 plots the relationship between a metro's optimal tax and the average tax faced by its substitute locations. The average substitute tax is the first term in the optimal tax formula. It is the migration-weighted mean of taxes across all other metros, where the weights obtained from our full substitution matrix reflect how individuals leaving metro j would reallocate across alternative destinations based on their estimated substitution patterns from our survey.

Figure 14: Own optimal tax and Average tax of substitute metros

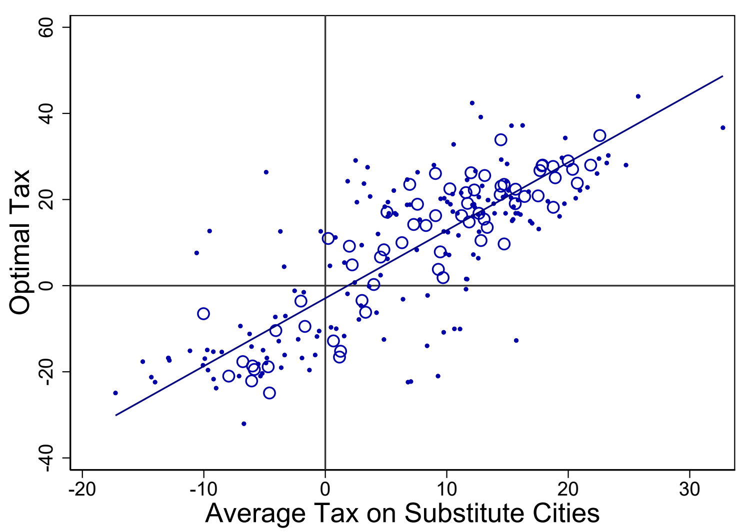

A key question for understanding the role of this term is whether the migration-weighted average

Figure 15: Optimal Metro Tax and College Share

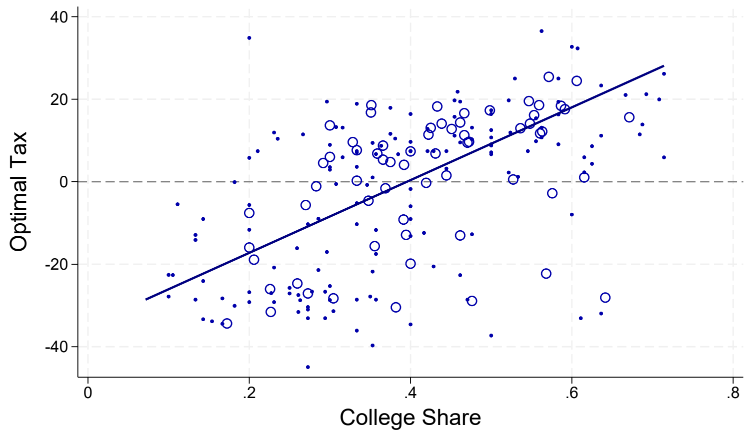

tax faced by the marginal outflow varies substantially across origin metros. Under a benchmark non-mixed logit model, marginal movers from any origin substitute toward destinations in proportion to aggregate market shares. In that case, the "average tax of substitute metros" would be (nearly) the same for every origin, and this component would explain essentially none of the cross-metro variation in optimal taxes. In contrast, our menu-based design reveals strong geographic clustering in substitution patterns, so the implied destination mix differs sharply across origins, generating substantial dispersion in the average tax faced by movers leaving different metros.

Consistent with the optimal-tax formula, Figure 14 shows that a metro's optimal tax comoves with the migration-weighted average tax of its substitutes. The positive relationship is therefore not itself the main empirical takeaway; rather, the fact that the average tax of substitute metros varies meaningfully across metros implies that fiscal externalities through migration are an important source of cross-metro differences in optimal taxes.

The scatter also displays considerable variation around the fitted line, reflecting the independent contribution of the MU deviation term. Metros lying above the fitted line have lower marginal utility than the national average (relative to their substitute-tax baseline) and thus face higher taxes than the substitute-weighted average would imply. Metros below the line have higher marginal utility and receive lower taxes. Together, these patterns illustrate how the two components of the formula—the substitute-tax baseline and the MU-based adjustment—jointly determine each metro's optimal tax.

9.5 Optimal Metro Tax and College Share

Figure 15 plots optimal metro taxes against the share of residents holding a college degree. We find that metros with higher college shares face higher optimal taxes, while those with lower college shares tend to receive subsidies. This relationship reflects the interaction between demographic composition and the planner's redistributive objectives. Recall that the average marginal utility in metro j is given by:

$$\mu_j = \frac{\sum_t N^t \mathcal{P}_j^t \left(\Pi^t \mu_j^t\right)}{\sum_t N^t \mathcal{P}_j^t}; \qquad (\text{with} \quad \bar{\mu} = \sum_j L_j \mu_j)$$

where Π^t denotes the Pareto weight assigned to type t and μ_j^t is the average marginal utility of income for type t individuals who choose metro j. In our baseline specification, we assign a Pareto weight of 2 to non-college workers relative to college workers.

Under this weighting scheme, metros populated predominantly by college-educated workers have lower Pareto-weighted average marginal utility, pushing the MU deviation term in the positive direction and raising optimal taxes. Conversely, metros with higher shares of non-college workers—who receive greater weight in the planner's objective—exhibit higher weighted marginal utility and thus warrant more favorable tax treatment. The positive slope in Figure 15 captures this compositional effect.

The college-share geographical variation also contributes to the geographic patterns documented in Figure 12. Coastal metros and major knowledge-economy hubs—which tend to have high concentrations of college-educated workers—face higher optimal taxes, while metros in regions with lower college attainment rates receive subsidies.

9.5.1 Sensitivity to Non-College Pareto Weights

Table 4: Sensitivity to Pareto Weights

	(1)	(2)	(3)
	Pareto Wgt = 2	Pareto Wgt = 4	Pareto Wgt = 8
MU "Metro Effect" (AKM)	-34.08***	-24.91***	-22.00***
	(3.747)	(3.938)	(4.193)
College Share	16.43***	24.36***	29.12***
	(3.360)	(3.685)	(4.023)
Avg. Tax of Substitutes	0.803***	0.852***	0.867***
	(0.0725)	(0.0765)	(0.0769)
Constant	-5.573***	-8.782***	-10.89***
	(1.514)	(1.665)	(1.818)
Observations	301	301	301
R^2	0.599	0.579	0.581

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 4 examines the sensitivity of optimal taxes to the planner's redistributive preferences by varying the Pareto weight assigned to non-college workers. We re-solve optimal taxes under three specifications: Pareto weights of 2, 4, and 8 for non-college relative to college workers (columns 1 through 3, respectively). Optimal taxes is a non-linear function of a bunch of things, so to get a sense of the relative importance of the channels through which redistributive preferences operate, we regress optimal taxes on the MU metro effect, college share, and the average tax of substitutes.[15]

Several patterns emerge as the weight on non-college workers increases. First, and most directly connected to the preceding discussion, the coefficient on college share rises monotonically—from 16.43 at weight 2 to 24.36 at weight 4 to 29.12 at weight 8. This steepening confirms that the positive relationship between college share and optimal taxes documented in Figure 15 is amplified by redistributive preferences favoring non-college workers. As the planner places greater weight on the welfare of less-educated individuals, the penalty on metros with college-educated populations increases. Second, observe that this increase in no way removes the importance of location-specific MU. The coefficient on the MU metro effect attenuates in absolute value, falling from −34.08 to −24.91 to −22.00 as the Pareto weight increases. This pattern reflects a shift in the composition of the planner's objective. Under stronger redistributive preferences, the planner prioritizes transfers toward metros with high concentrations of non-college workers, even if measured location-specific marginal utility (the AKM metro effect) is not especially high in those locations. In other words, cross-type redistribution partially crowds out the within-type, cross-location redistribution that the MU metro effect captures. Third, the coefficient on average substitute tax remains stable, rising only slightly from 0.80 to 0.87.

9.6 Decomposition by Tax Formula Components

Table 5 presents regression decompositions that formalize the relationships documented in the preceding subsections. By regressing optimal taxes on the key components of our formula, we quantify the relative contributions of metro-level marginal utility, demographic composition, and substitute metro taxes to the cross-sectional variation in optimal place-based transfers.

Column (1) regresses optimal taxes on three variables: the AKM metro effect for marginal utility, metro college share, and the average tax faced by substitute metros. All three regressors are highly significant and carry the expected signs. A one-unit increase in the log marginal utility metro effect reduces optimal taxes by approximately $34,000. This large magnitude reflects the central role of marginal utility in determining optimal spatial transfers. Metros with a 10 percentage point higher college share face roughly $1,600 higher optimal taxes, reflecting the lower Pareto weight assigned to college workers in our baseline specification. Own taxes move approximately 80 cents for every dollar increase in the migration-weighted average tax of substitutes.

Column (2) replaces the average substitute tax with its underlying determinants—the average MU metro effect and average college share of the substitute metros. This specification examines whether the substitute tax effect operates through the same channels that determine own taxes.

[15] If we are fully flexible, then these three components will fully explain the variation in optimal taxes.

Table 5: Optimal tax decomposition

	(1) Optimal Tax	(2) Optimal Tax
MU "Metro Effect" (AKM)	-34.08***	-43.91***
	(3.75)	(4.51)
College Share	16.43***	17.37***
	(3.36)	(4.44)
Avg. Tax of Substitutes	0.80***	
	(0.07)	
Avg. MU "Metro Effect" of Substitutes		-19.97*
		(9.01)
Avg. College Share of Substitutes		36.44***
		(10.49)
Constant	-5.57***	-21.66***
	(1.51)	(3.48)
Observations	301	301
R^2	0.599	0.474

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Metros whose substitutes have higher marginal utility (and thus receive subsidies) themselves face lower optimal taxes. Metros whose substitutes have more college-educated populations face higher optimal taxes. Both align with the logic of the optimal tax formula where a metro's tax is pulled toward the taxes of its substitutes, and those substitute taxes are themselves determined by marginal utility and demographic composition.

The high R^2 indicates that these three components, own MU, own college share, and substitute taxes, explain the majority of cross-metro variation in optimal taxes.[16] Together, these decompositions confirm that the optimal taxes generated by our algorithm adhere closely to the theoretical structure of the optimal tax formula.

9.7 Sensitivity to Migration Elasticities

In this counterfactual exercise, we examine what happens to the distribution of optimal taxes when we make individuals more elastic, holding location-specific marginal utility of incomes fixed.

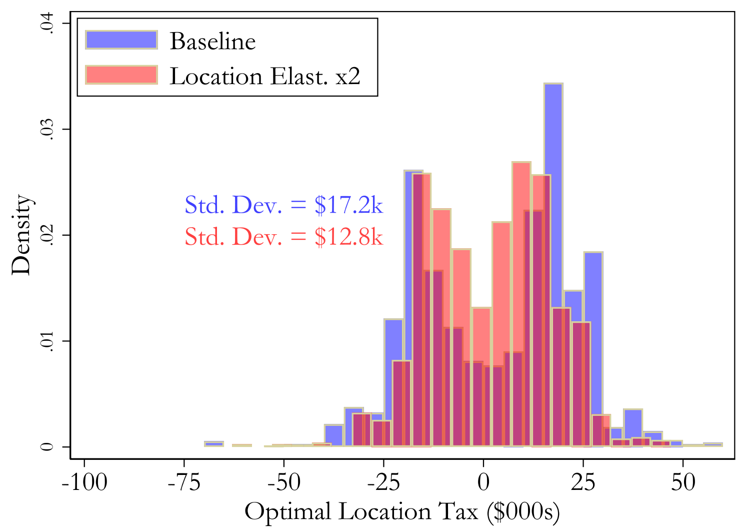

Figure 16: Sensitivity to Migration Elasticities

Figure 16 compares the distribution of optimal metro taxes under our baseline migration elasticities (blue) with the distribution obtained when all location choice elasticities are doubled (red). We find that this leads to a compression of the optimal tax distribution where the standard deviation falls by nearly 25.5% from $17,200 under baseline elasticities to $12,800 when elasticities are doubled.

This compression reflects the trade-off in optimal spatial taxation that emerges directly from our formula. Recall that the MU deviation term is scaled by the inverse of the own-tax migration

[16] The somewhat lower R^2 of 0.47 in column (2) relative to 0.599 in column (1) reflects the fact that average substitute tax is a sufficient statistic that aggregates information beyond just the mean MU and college share of substitutes, since it also depends on the average tax of the substitutes of these substitute metros.

response:

$$\tau_j = \underbrace{\frac{\sum_{j'\neq j} \frac{\partial \mathcal{P}_{j'}}{\partial \tau_j} \tau_{j'}}{\sum_{j'\neq j} \frac{\partial \mathcal{P}_{j'}}{\partial \tau_j}}}_{\text{Avg. tax of substitutes}} + \underbrace{\frac{1}{\frac{\partial \mathcal{P}_j}{\partial \tau_j}} \frac{\mu_j - \bar{\mu}}{\bar{\mu}}}_{\text{MU deviation term}}$$

Since $\frac{\partial \mathcal{P}_j}{\partial \tau_j} < 0$, higher migration elasticities reduce the magnitude of the second term. Intuitively, when individuals are more responsive to tax differentials, a larger share of the population leaves taxed metros, shrinking the local tax base. Raising an additional dollar for redistribution therefore requires imposing a greater burden on the remaining residents, increasing the marginal cost of redistribution. Facing these higher costs, the planner optimally compresses the tax distribution, accepting smaller redistributive gains to avoid the fiscal burden that arises when taxes fall on an increasingly narrow base.

In this counterfactual exercise, since elasticities only double but do not change in sign, the compression is approximately symmetric around the center of the distribution. This symmetry follows from the structure of our formula which indicates that higher elasticities scale down the MU deviation term proportionally for all metros, regardless of whether that term is positive or negative. Metros that warranted large subsidies under baseline elasticities still receive subsidies, but of smaller magnitude; likewise for metros facing large taxes. Consequently, the bimodal structure of the distribution persists under higher elasticities, with the two modes shifting closer together. This persistence indicates that the qualitative pattern of optimal policy where high-MU metros are subsidized and low-MU metros are taxed remains even when behavioral responses are substantially larger.

The resulting substantial optimal tax differentials represent a middle ground between the extremes of no spatial redistribution which one would obtain under perfect mobility and unconstrained transfers which one would obtain if workers were immobile.

9.8 Decomposition of the Across-Metro Redistribution Motive

The optimal tax formula highlights that spatial redistribution is driven by deviations of metro-level marginal utility from the population-weighted mean, $\mu_j - \bar{\mu}$. This term captures the welfare gains from tilting transfers toward locations where marginal utility is high. In this section, we discuss what accounts for the variation in this motive across locations.

Let $MU_{ij} = \Pi_{g(i)} \times MU_{ij}^{within} \times MU_i^{loc}$ denote the Pareto-weighted marginal utility of individual i in location j. Denote by $\overline{MU}_i = \sum_j p_{ij} MU_{ij}$ the expected marginal utility for individual i across their candidate locations, and by $\overline{MU}_{g(i)}$ the average expected marginal utility among individuals of type $g(i)$. We decompose the redistribution motive as:

$$\mu_j - \bar{\mu} = \underbrace{\sum_i p_{ij} \left(MU_{ij} - \overline{MU}_i \right)}_{\text{Within-person}} + \underbrace{\sum_i p_{ij} \left(\overline{MU}_i - \overline{MU}_{g(i)} \right)}_{\text{Across-person, within-type}} + \underbrace{\sum_i p_{ij} \left(\overline{MU}_{g(i)} - \bar{\mu} \right)}_{\text{Across-type}} \tag{21}$$

The first term captures within-person variation: holding the individual fixed, for individuals who choose metro j with positive probability, how much does marginal utility provided by metro j compare to the marginal utility provided by other metros that individual might choose? The second term captures across-person, within-type variation: among individuals of the same type (e.g., non-college workers), compared to average households of their type, do households who choose metro j consider menus of locations with higher or lower than average marginal utility? The third term captures across-type variation: how much does the type composition of metro j differ from the national average, and how does this compositional difference translate into marginal utility differences given the Pareto weights?

Table 6 presents a decomposition of the cross-metro variance in optimal tax rates into four components: (A) the within-person redistribution motive (within-respondent differences in marginal utility across locations), (B) the within-type redistribution motive (sorting across metros within a skill group), (C) the across-type redistribution motive (differences in the skill composition of metros, interacted with Pareto weights), and (D) the "neighbors" term that captures the fiscal-externality channel through the average tax of substitute metros.

For our purposes, the key takeaway is that the redistributive motive in the tax formula is not driven solely by cross-type differences. Treating location solely as a tag for type would obtain any redistributive motive entirely from component (C). In contrast, we find that components (A) and (B) together account for a non-negligible share of the redistribution motive, and these are precisely the components that require our survey to identify. In other words, within-person and within-type channels provide additional justification for place-based redistribution beyond cross-type redistribution alone. Even holding fixed the planner's Pareto weights across skill groups, the survey uncovers systematic within-type differences in where marginal utility is high, and these differences imply a distinct geographic pattern of desired transfers. Absent the survey, the planner could at best implement redistribution that follows type composition, or else rely on assumptions about how within-type marginal-utility differences relate to observable location attributes.

A final takeaway concerns the role of substitution patterns in shaping the fiscal-externality term in the optimal-tax formula. Under a simple non-mixed logit benchmark, the marginal outflow from any metro would reallocate across all other metros in proportion to their aggregate population shares. In that case, the "average tax of substitutes" faced by movers leaving a given metro would be approximately the same for every origin metro–essentially a common, market-share-weighted average of taxes elsewhere–and would therefore explain little (or none) of the cross-metro variation in optimal taxes. In our data, however, substitution patterns are far from proportional to national market shares. Because respondents consider and substitute disproportionately among a limited set of geographically and economically related metros, the migration-weighted average tax faced by the marginal outflow varies substantially across origin metros. This dispersion in the average tax of substitutes is a key contributor to the cross-metro variation in optimal tax rates.

Table 6: Variance decomposition

Component	Share of total	Share within A+B+C
A. Within-person	0.032	0.088
B. Within-type (across people)	0.154	0.425
C. Across-type	0.176	0.487
D. Neighbors / substitute-tax term	0.638	—

Notes: Column 2 reports each component's share of total cross-metro variation. Column 3 re-normalizes A–C so that A+B+C=1 and decomposes the redistribution motive only.

10 Panel Analysis of Marginal Utility

The cross-sectional results above describe how average marginal utility varies across locations at baseline expectations. In this section we use the panel structure of the survey to study how marginal utility within a given location changes when that location becomes more or less attractive due to experimentally induced changes in net income offers (and family/friends), holding all other location attributes fixed. These experimental shocks shift both (i) the level of resources available in a location and (ii) the set of individuals who would choose that location, allowing us to separate within-location changes in marginal utility from composition effects due to sorting.

To organize the interpretation of these panel responses, we begin with a decomposition of how the average marginal utility in a location changes with a location-specific income shock into a direct effect and a selection (sorting) effect.

Denote by MUC_j the marginal utility of income in location j, and consider the average marginal utility of income conditional on choosing j. The derivative with respect to the wage in location j can be decomposed into a direct and a selection effect:

$$\frac{d}{dw_j}E[MUC_j|j \text{ chosen}] = \underbrace{\int \frac{\partial MUC_j(w_j, e_i)}{\partial w_j} f(e_i|j \text{ chosen}) de_i}_{\text{Direct effect}} + \underbrace{\int MUC_j(w_j, e_i) \frac{\partial f(e_i|j \text{ chosen})}{\partial w_j} de_i}_{\text{Selection effect}} \tag{22}$$

Furthermore, the selection effect can be re-written as:

$$\frac{1}{P_j(w_j)} \int \frac{\partial}{\partial w_j} \Pr\big(U_j \geq U_k \mid e_i\big) \Big[MUC_j(w_j, e_i) - E\big(MUC_j \mid j \text{ chosen}\big)\Big] f(e_j)\, de_i$$

The direct effect captures how marginal utility changes for a fixed individual in a fixed location when income in that location changes. The selection effect captures how the composition of residents changes as the location becomes more attractive and marginal movers enter or exit, which in turn changes the average marginal utility among those who choose the location. This decomposition clarifies why, in our setting, within-location average marginal utility can increase with income even though standard separable models imply only a direct (typically negative) effect. Note that this form of selection effect is not present in a model with separable shocks.

Our sufficient-statistics policy calculations use baseline estimates of $\{\mu_j\}$ as inputs, abstracting from how policy changes might alter these marginal utilities through sorting. The structural model in the next section relaxes this abstraction by allowing μ_j to change endogenously with policy-induced changes in net incomes and the composition of residents. We estimate the structural model to match these panel responses, which trace how respondents' marginal-utility measures move as we experimentally vary location-specific net incomes within their menus. This allows the data to discipline how marginal utility changes as policy changes and induces sorting.

10.1 Experimental Design and Estimation

To identify causal relationships between income and marginal utility of incom, we leverage our experimental design where respondents evaluated multiple hypothetical scenarios with varying income levels and presence of family/friends across locations. The exogenous variation we induce in our survey removes selection based concerns. Additionally, the panel structure enables us to include multiple fixed effects that additionally account for unobserved heterogeneity.

As described earlier, for each respondent i in each scenario k, we construct $\log(MUC_{ijk}) \equiv \log(1/b_{ijk})$ for each location j, where b_{ijk} represents the income boost required in location j to match the utility gain from a \$5,000 income boost in the highest-MUC location. Our estimating equation is:

$$\log(MUC_{ijk}) = \beta_1 \log w_{ijk} + \beta_2 f_{ijk} + \theta_{ik} + \theta_{ij} + e_{ijk} \tag{23}$$

where θ_{ik} and θ_{ij} represent individual-by-scenario and individual-by-location fixed effects, respectively. Our experimental variation allows us to use this specification to isolate the causal effect of income changes (and changes in presence of family and friends) on marginal utility of income conditional on individual-location and individual-scenario fixed effects.

Next, we analyze the underlying distribution of the elasticity of MU with respect to expected income in two steps. First, we exploit the within respondent design of our survey to estimate respondent-specific elasticities and estimate:

$$\log(MUC_{ijk}) = \beta_{i1} \log w_{ijk} + \beta_{i2} f_{ijk} + \theta_{ik} + \theta_{ij} + e_{ijk} \tag{24}$$

Second, we implement an empirical Bayes shrinkage procedure to adjust for sampling variance, shrinking individual-specific estimates toward the sample mean based on the estimated signal-to-noise ratio.

10.2 Panel Results

Our panel analysis yields an estimated elasticity of MUC with respect to expected income (β_1) of 0.45 (SE = 0.01), larger than our cross-sectional estimates. This difference likely reflects our experimental design's ability to isolate causal effects from confounding factors. The presence of family and friends (β_2) increases MUC by approximately 4% (SE = 0.00), indicating a complementarity between social connections and the marginal utility of income.

The distribution of these estimates from the empirical Bayes procedure reveals substantial heterogeneity in how respondents' marginal utility responds to income changes across locations, as shown in Figure 17. The mean elasticity is 0.43 with a standard deviation of 0.16 yielding a standard error of 0.003, indicating considerable variation in the strength of this relationship.

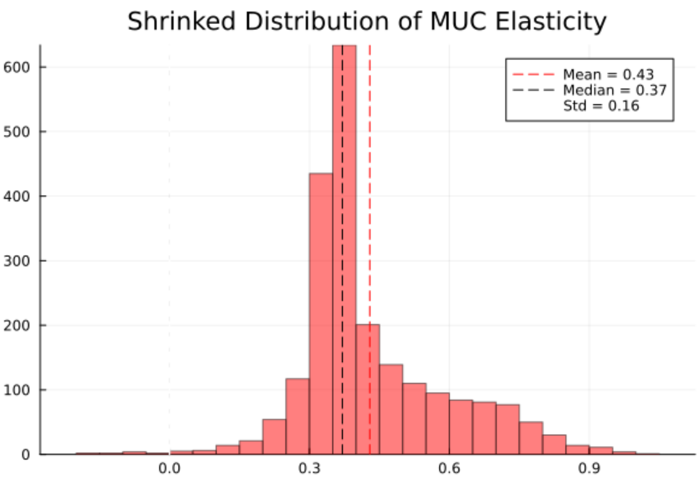

Figure 17: Distribution of Estimates of MUC Elasticity

The panel estimates inform the theoretical relationship between income and MU conditional on location choice described above. The positive overall elasticity of 0.45 indicates that the selection effects dominate the direct effects, with the latter typically negative in standard models. This pattern suggests that marginal entrants to higher-income locations have systematically higher MU than inframarginal residents. In standard linearly separable models the selection effect by construction is zero thereby not allowing any heterogeneity between inframarginal and marginal individuals. These findings have significant implications for optimal redistribution policies across locations, as they demonstrate that the standard approach of simply redistributing from high-income to low-income locations may not align with welfare maximization when accounting for heterogeneity in marginal utility of income.

10.3 Rejection of separable model

Several of our main empirical results findings are difficult or impossible to rationalize with discrete location choice model that specifies consumption utility as additively separable from location-specific shocks or attachment draws, such as $\log(w_{jk}) + \epsilon_{ij}$ where w_{jk} represents income in city j for skill group k, and ϵ_{ijk} are location specific attachment draws that captures individual-specific location preferences. This model places strong restrictions on how MUC varies across locations (on average) and across individuals within a location. First, MUC is a (decreasing) function of the consumption level (wage). As a result, high wage cities must have lower average MUCs, and within a city individuals with the same wage will have the same MUC regardless of the strength of their attachment to the city. These model predictions do not change when the utility specification is augmented to allow non-wage factors to affect location choice by including additive linear terms in observable location characteristics or including an additive composite location amenity parameter.

Our empirical findings reject these restrictions implied by the additively separable model. Specifically, the standard additively separable model cannot rationalize the findings that the average MUC of each location conditional on the location being chosen is positively correlated with the location's wage level, but variation in wage levels across locations explains a small fraction of the overall variation in average MUC across locations. Within individual, the expected MUC in a location conditional on that location being chosen is *increasing* the location's wage.

The selection effect ($\int MUC_j(w_j, e_i)\frac{\partial f(e_i|j \text{ chosen})}{\partial w_j}de_i$) in the separable model is zero by construction. We introduce an alternative utility function that allows for substitutability between income and person-specific amenities, which generates these patterns in equilibrium allowing for selection effects.

11 Model

Next, to investigate what underlying preferences and distributions of location attachment are consistent with our main empirical results, and to study their implications for optimal redistributive policy, we develop and estimate a simple structural model. In this preliminary draft, we study a stylized model of a homogeneous population (and thus a choice problem represented by a single random utility specification) that we estimate by targeting key average elasticities and sample moments.

11.1 Non-separable Structural Model

Each individual i has a feasible menu of locations $\mathcal{M}_i \subseteq \{1, \ldots, J\}$. For each location $j \in \mathcal{M}_i$, individual i earns wage w_{ij} and draws an idiosyncratic shock ε_{ij}, where $\varepsilon_{ij} \sim \text{pareto}(1, 1)$. Individual i chooses a location $j \in \mathcal{M}_i$ to maximize expected utility.

Locations are grouped into a finite number of types. Location j belongs to a discrete type $\tau(j) \in \{1, \ldots, T\}$. For each individual–location pair, $rank(i, j)$ denotes the rank that individual i assigns to location $j \in \mathcal{M}_i$. The model includes amenity parameters $\{A_{\tau(j)}, a_{rank(i,j)}, B_{\tau(j)}, b_{rank(i,j)}\}$ capturing type-specific and rank-specific components. Preferences over consumption and amenities are represented by

$$U_{ij} = \frac{\left[\left(B_{\tau(j)}\, b_{rank(i,j)}\, w_{ij}\right)^{\rho} \; + \; \left(\lambda\, A_{\tau(j)}\, a_{rank(i,j)}\, \varepsilon_{ij}\right)^{\rho} \right]^{\frac{1-\sigma}{\rho}}}{1 - \sigma}.$$

The parameter set is

$$\Theta = \left\{\sigma, \lambda, \rho, \{A_t, B_t\}_{t=1}^{T}, \{a_r, b_r\}_{r=1}^{R} \right\},$$

where $R = \|\mathcal{M}_i\| = 4$ and $T = 4$.

11.2 Identification

We normalize $A_1 = B_1 = a_1 = b_1 = 1$. Table tab:identification describes the mapping of specific moments in the data to identify the parameters of the model.

Table 7: Mapping of Moments to Model Parameters

Parameter	**Targeted Moment**
λ	Elasticity of location choice with respect to wage
ρ	Elasticity of EMUC with respect to wage
σ	Within-location MU curvature (CRRA = 1; imposed)
$A_j + B_j$	Choice probability for city type $j \ \forall \ j \neq 1$
$a_j + b_j$	Impact of rank of j on choice probability relative to rank of 1 (holding types fixed) $\forall \ j \neq 1$
$\frac{A_j}{A_j+B_j}$	Average MUC in city type j relative to MUC in city 1 (holding ranks fixed) $\forall \ j \neq 1$
$\frac{a_j}{a_j+b_j}$	Impact of rank of j on EMUC relative to rank of 1 $\forall \ j \neq 1$

11.3 Estimation

Estimation of our non-separable model follows in two-steps. First, to reduce dimensionality, we conduct a k-means clustering to classify cities into 4 discrete types. Second, after discretizations we estimate the model using Indirect Inference (Gourieroux et al., 1993).

$$\widehat{\Theta} = \arg\min_{\Theta} (\tilde{\gamma}(\Theta) - \widehat{\gamma})'(\tilde{\gamma}(\Theta) - \widehat{\gamma}),$$

where $\widehat{\gamma}$ denotes the vector of auxiliary estimates in the data and $\tilde{\gamma}(\Theta)$ the same vector computed from simulated data given parameter vector Θ. The auxilliary models that we estimate using baseline data:

$$\ln(P_{ij}) = \theta_i + \sum_{j\neq 1} \beta_j \text{Type}_j + \sum_{j\neq 1} \alpha_j \text{Rank}_j + \psi w_{ij} + u_{ij}$$

$$\ln(MUC_{ij}) = \theta_i + \sum_{j\neq 1} \gamma_j \text{Type}_j + \sum_{j\neq 1} \delta_j \text{Rank}_j + \psi w_{ij} + e_{ij}$$

For $j = 2, 3, 4$, β_j identifies $(A_j + B_j)$, α_j identifies $(a_j + b_j)$, γ_j identifies $A_j/(A_j + B_j)$, δ_j identifies $a_j/(a_j + b_j)$.

11.4 Parameter estimates

We report the estimates of our structural model in Table 8

Table 8: Estimated Model Parameters

Parameter	Type 1	Type 2	Type 3	Type 4
A	1.000	0.952	0.942	0.835
B	1.000	1.351	1.174	1.832
a	1.000	0.294	0.160	0.069
b	1.000	1.245	1.208	1.201
ρ		0.892		
σ		2.956		
λ		185.738		

11.5 Fit of non-targeted moments

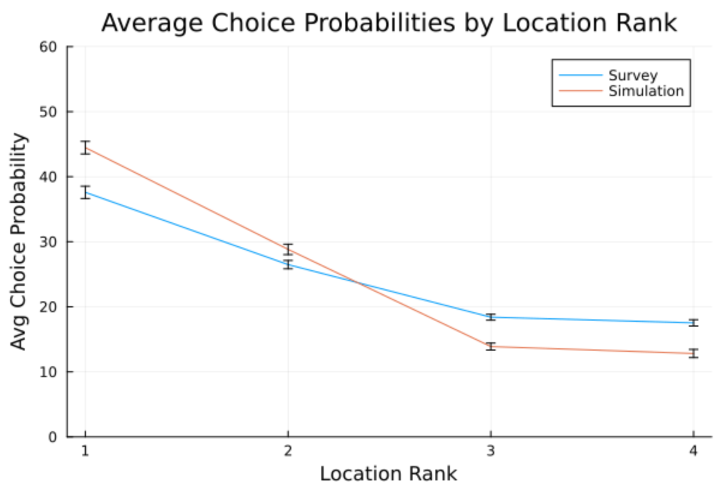

12 Welfare comparison of other policies with optimal taxes obtained for sufficient statistics

In progress.

13 Conclusion

This paper develops and implements a new approach to measuring a central input to optimal place-based policy: how the marginal utility of income varies across locations. We show that this object is not identified from location choice data alone without strong assumptions on preference separability, and the implied methods are practically infeasible to implement. We find that these assumptions are not consistent with our data. Our solution is to measure marginal utilities directly through a survey that elicits within-person comparisons of the value of additional income across locations where each respondent could plausibly live. Combined with experimentally varied scenarios that shift location-specific incomes and contextual factors, our survey also recovers the full matrix of substitution elasticities across locations needed to characterize the efficiency costs of spatial redistribution. Our survey thus measures the two key inputs required to fully characterize optimal spatial transfers in absence of agglomeration externalities.

Our main empirical findings can be summarized as follows. First, the marginal utility of income exhibits substantial geographic variation that cannot be explained by income differences alone. Controlling for expected income, marginal utility is lower in high housing cost locations and higher in locations with better schools, stronger labor markets, and lower crime. Second, substitution patterns across locations display strong regional clustering: when a metro becomes less attractive, nearby and culturally similar metros absorb the bulk of outflows. Third, optimal metro-specific taxes vary widely, and exhibits a bimodal distribution. The bimodality and geographic clustering arise because metros that serve as close substitutes tend to share similar characteristics, generating groups of metros that are taxed or subsidized together. Each metro's optimal tax equals the migration-weighted average tax of its substitutes, adjusted for the metro's deviation of marginal utility from the national mean. The metro-level marginal utility effect, local college share, and average substitute tax together explain sixty percent of the variation. Notably, we document that optimal policy is not the same as a transfer from high wage locations to low wage ones, since marginal utility differences are not solely explained by income differences.

A key finding concerns the sources of the redistribution motive. We decompose the variation in metro-level marginal utility into three components: within-person variation (how marginal utility differs across locations for the same individual), across-person within-type variation (how individuals of the same type who sort into different metros differ in their expected marginal utility), and across-type variation (how the type composition of metros differs). Within-person variation and across-person within-type variation accounts for half of the variation in marginal utility differences.

Our results also highlight the sensitivity of optimal policy to key parameters. Increasing the

Pareto weight on non-college workers steepens the gradient between optimal taxes and college share, while attenuating the role of location-specific marginal utility effects. Doubling migration elasticities compresses the optimal tax distribution by approximately 25 percent, as higher mobility raises the fiscal cost of tax differentials. The qualitative patterns of optimal taxation persist across specifications, but magnitudes depend on both normative choices about redistribution and empirical estimates of behavioral responses.

Our analysis abstracts from agglomeration externalities and productivity spillovers, focusing exclusively on the redistribution motive for place-based transfers. Future research could extend our work by integrating both efficiency and equity considerations.

References

Abowd, J.M., Kramarz, F., Margolis, D.N., 1999. High wage workers and high wage firms. Econometrica 67, 251–333.

Alam, M.M.U., Mookerjee, M., Roy, S., 2021. Worker-side Discrimination: Beliefs and Preferences—Evidence from an Information Experiment on Jobseekers. Working Paper.

Andrew, A., Adams, A., 2025. Revealed beliefs and the marriage market return to education. The Quarterly Journal of Economics , qjaf020.

Blass, A.A., Lach, S., Manski, C.F., 2010. Using elicited choice probabilities to estimate random utility models: Preferences for electricity reliability. International Economic Review 51, 421–440.

Blundell, R., Pistaferri, L., Preston, I., 2008. Consumption inequality and partial insurance. American Economic Review 98, 1887–1921.

Bound, J., Holzer, H.J., 2000. Demand shifts, population adjustments, and labor market outcomes during the 1980s. Journal of labor Economics 18, 20–54.

Canay, I.A., 2011. A simple approach to quantile regression for panel data. The econometrics journal 14, 368–386.

Diamond, R., 2016. The determinants and welfare implications of u.s. workers' diverging location choices by skill: 1980-2000. The American Economic Review 106, 479–524.

Donald, E., Fukui, M., Miyauchi, Y., 2025. Optimal dynamic spatial policy. Technical Report. National Bureau of Economic Research.

Efron, B., 2016. Empirical bayes deconvolution estimates. Biometrika 103, 1–20.

Fajgelbaum, P.D., Gaubert, C., 2020. Optimal spatial policies, geography and sorting. Quarterly Journal of Economics 135, 959–1036.

Fajgelbaum, P.D., Gaubert, C., 2025. Optimal spatial policies. Technical Report. National Bureau of Economic Research.

Fajgelbaum, P.D., Morales, E., Suárez Serrato, J.C., Zidar, O., 2019. State taxes and spatial misallocation. The Review of Economic Studies 86, 333–376.

Gaubert, C., Kline, P., Yagan, D., 2021. Placed-based redistribution. NBER Working Paper w28337.

Gourieroux, C., Monfort, A., Renault, E., 1993. Indirect inference. Journal of applied econometrics 8, S85–S118.

Kuziemko, I., Norton, M., Stantcheva, S., Saez, E., 2015. How elastic are preferences for redistribution: Evidence from randomized survey experiments. American Economic Review 105, 1478–1508.

Notowidigdo, M.J., 2011. The incidence of local labor demand shocks. Technical Report. National Bureau of Economic Research.

Rossi-Hansberg, E., Sarte, P.D., Schwartzman, F., 2022. Cognitive hubs and spatial redistribution. Working Paper.

Serrato, J.C.S., Wingender, P., 2014. Estimating the Incidence of Government Spending. Technical Report. Duke University and the International Monetary Fund.

Serrato, J.C.S., Zidar, O.M., 2015. Who Benefits from State Corporate Tax Cuts? A Local Labor Markets Approach with Heterogeneous Firms. Technical Report. Duke University and The University of Chicago Booth School of Business.

Walters, C., 2024. Empirical bayes methods in labor economics, in: Handbook of Labor Economics. Elsevier. volume 5, pp. 183–260.

Wiswall, M., Zafar, B., 2018. Preference for the workplace, investment in human capital, and gender. The Quarterly Journal of Economics 133, 457–507.

A Appendix: Derivation of Proposition 2.1

This appendix derives the sufficient–statistics condition in Proposition 2.1.

A.1 First–order conditions

The planner chooses τ to maximize

$$W(\tau) = \sum_{t=1}^{T} \Pi_t N_t U_t(\tau)$$

subject to $\sum_{j=1}^{J} \tau_j L_j(\tau) = G$, where $L_j(\tau) = \sum_t N_t P_j^t(\tau)$. Let Λ denote the multiplier on the budget constraint. The Lagrangian is

$$\mathcal{L}(\tau, \Lambda) = \sum_t \Pi_t N_t U_t(\tau) + \Lambda \left(\sum_{j=1}^{J} \tau_j L_j(\tau) - G \right).$$

A marginal increase in τ_j reduces consumption one–for–one for individuals who choose metro j, so

$$\frac{\partial U_t}{\partial \tau_j} = -P_j^t(\tau)\, \mu_j^t,$$

where μ_j^t is the expected marginal utility of income for type t among those who choose j. The FOC for τ_j is therefore

$$0 = -\sum_t \Pi_t N_t P_j^t(\tau)\, \mu_j^t + \Lambda L_j(\tau) + \Lambda \sum_{j'=1}^{J} \tau_{j'} \frac{\partial L_{j'}(\tau)}{\partial \tau_j}.$$

Define the Pareto–weighted average marginal utility among residents of metro j as

$$\mu_j \equiv \frac{\sum_t \Pi_t N_t P_j^t(\tau)\, \mu_j^t}{L_j(\tau)}.$$

Dividing the FOC by $L_j(\tau)$ yields

$$\frac{\mu_j - \Lambda}{\Lambda} = \frac{1}{L_j(\tau)} \sum_{j'=1}^{J} \tau_{j'} \frac{\partial L_{j'}(\tau)}{\partial \tau_j}. \tag{25}$$

A.2 Identifying Λ

Let δ denote a common marginal increase in all metro head taxes, so τ_j becomes $\tau_j + \delta$ for all j. Under the assumption that a common tax change shifts all location utilities by the same amount, location choice probabilities and hence the vector $L(\tau)$ are unchanged locally. The budget constraint implies

$$\sum_j (\tau_j + \delta) L_j(\tau) = G + \delta,$$

so $dG/d\delta = 1$. Welfare changes by

$$\frac{dW}{d\delta} = \sum_t \Pi_t N_t \frac{dU_t}{d\delta} = -\sum_t \Pi_t N_t \mu_t, \qquad \mu_t \equiv \sum_j P_j^t(\tau) \mu_j^t.$$

Define $\bar{\mu} \equiv \sum_t \Pi_t N_t \mu_t$. Since $dG/d\delta = 1$, this implies $dW/dG = -\bar{\mu}$.

On the other hand, $\partial \mathcal{L}/\partial G = -\Lambda$, so at the optimum $dW/dG = -\Lambda$. Hence $\Lambda = \bar{\mu}$.

Substituting $\Lambda = \bar{\mu}$ into (25) delivers Proposition 2.1.

Algebraic manipulation provides optimal taxes for city j is a sum of average tax of substitute metros and deviation of MU of metro j from the average as:

$$\tau_j = \underbrace{\frac{\sum_{j' \neq j} \frac{\partial \mathcal{P}_{j'}}{\partial \tau_j} \tau_{j'}}{\sum_{j' \neq j} \frac{\partial \mathcal{P}_{j'}}{\partial \tau_j}}}_{\text{Average tax of substitute metros}} + \underbrace{\frac{1}{\frac{\partial \mathcal{P}_j}{\partial \tau_j}} \cdot \frac{\mu_j - \bar{\mu}}{\bar{\mu}}}_{\text{MU-based adjustment}}$$

A.3 Two–city derivation

This appendix derives equation (4) in the two–metro special case.

Consider two locations, indexed by $j = 1, 2$, and a single individual type. The government provides a head transfer b to everyone in metro 1 and finances it with a head tax t on everyone in metro 2, subject to a balanced budget.

Let W_1 and W_2 denote pre–tax incomes. Under the policy, consumption is $W_1 + b$ in metro 1 and $W_2 - t$ in metro 2. Individuals draw idiosyncratic shocks (ϵ_1, ϵ_2) and choose the metro that maximizes utility, so expected indirect utility is

$$V(b,t) = \mathbb{E}_{\epsilon_1,\epsilon_2} \left[\max \left\{ u(W_1 + b, \epsilon_1),\, u(W_2 - t, \epsilon_2) \right\} \right].$$

Let $L_1(b)$ denote the equilibrium share of the population in metro 1 as a function of the transfer, with $1 - L_1(b)$ living in metro 2. The balanced budget constraint implies

$$L_1(b)\, b = [1 - L_1(b)]\, t(b), \qquad \text{so} \qquad t(b) = \frac{L_1(b)}{1 - L_1(b)}\, b.$$

We can write welfare as a function of the single instrument b:

$$\widetilde{V}(b) \equiv V\big(b, t(b)\big).$$

Differentiating,

$$\frac{d\widetilde{V}}{db} = \frac{\partial V}{\partial b} + \frac{dt}{db}\frac{\partial V}{\partial t} + \frac{dL_1}{db}\frac{\partial V}{\partial L_1}.$$

The final term is zero by the envelope theorem: households who are marginal and switch metros in response to a small change in (b, t) are indifferent between the two metros at the margin, so changes in the equilibrium cutoff (and hence L_1) do not contribute to first order. Therefore the optimal transfer satisfies

$$\frac{\partial V}{\partial b} + \frac{dt}{db}\frac{\partial V}{\partial t} = 0. \tag{26}$$

Let μ_1 and μ_2 denote the expected marginal utilities of income for households who choose metros 1 and 2, respectively, under the policy $(b, t(b))$. Then

$$\frac{\partial V}{\partial b} = L_1(b)\, \mu_1, \qquad \frac{\partial V}{\partial t} = -[1 - L_1(b)]\, \mu_2.$$

Differentiating the balanced–budget expression $t(b) = \frac{L_1(b)}{1-L_1(b)} b$ yields

$$\frac{dt}{db} = \frac{L_1}{1-L_1} + \frac{b}{(1-L_1)^2} \frac{dL_1}{db}.$$

Substituting these expressions into (26) gives

$$\begin{aligned} 0 &= L_1\mu_1 + \left[\frac{L_1}{1-L_1} + \frac{b}{(1-L_1)^2}\frac{dL_1}{db}\right]\big(-[1-L_1]\mu_2\big) \\ &= L_1\mu_1 - L_1\mu_2 - \mu_2\,\frac{b}{1-L_1}\frac{dL_1}{db}. \end{aligned}$$

Rearranging,

$$L_1(\mu_1 - \mu_2) = \mu_2\,\frac{b}{1-L_1}\frac{dL_1}{db},$$

and dividing both sides by $L_1\mu_2$ yields equation (4):

$$\frac{\mu_1 - \mu_2}{\mu_2} = \frac{1}{1-L_1}\left(\frac{dL_1}{db}\frac{b}{L_1}\right).$$

A.4 Computation details

For completeness, we briefly summarize the iterative algorithm that implements equation (17):

1. **Initialization.** Start from an initial guess $\tau^{(0)}$ for metro–level taxes.
2. **Compute after–tax incomes and choice probabilities.** Using $\tau^{(k)}$, compute after–tax incomes in each metro for each individual, and obtain choice probabilities $P_j^t(\tau^{(k)})$ and aggregate shares $L_j(\tau^{(k)})$.
3. **Construct the derivative matrix.** Using the constant elasticity specification and the observed probabilities, convert to derivatives $\partial L_{j'}(\tau^{(k)})/\partial \tau_j$ and form the matrix $A(\tau^{(k)})$.
4. **Update taxes.** Construct the right–hand side vector $b^{(k)}$ from the current metro–level marginal utilities $\{\mu_j(\tau^{(k)})\}$, and solve

   $$A(\tau^{(k)})\,\tau^{(k+1)} = b^{(k)}$$

   using a Gauss–Seidel iterative linear solver.
5. **Iterate.** Repeat steps 2–4 until the tax vector converges.

This algorithm implements the sufficient–statistics characterization using only the estimated marginal utilities by metro and the substitution matrix across metros that we recover from the survey, and it delivers the optimal place–based head tax (or subsidy) for each metro.

B Appendix: Canonical models with Gumbel or Fréchet shocks

A key question for optimal spatial policy is whether marginal utilities of consumption can be recovered from observed location choices. In this section, we examine two canonical models—an additive separable model with Type 1 Extreme Value (logit) shocks and a multiplicative model with Fréchet shocks—and show that both yield the same identification formula linking ratios of cross-location elasticities to marginal utility ratios.

B.1 Model 1: Additive Separable with Logit Shocks

Consider a standard discrete choice model where utility for individual i in location j is additively separable:

$$U_{ij} = u(W_j) + a_j + \sigma\epsilon_{ij} \tag{27}$$

where $u(W_j)$ is utility from wages, a_j captures location amenities common to all individuals, and ϵ_{ij} is drawn iid Type 1 Extreme Value (Gumbel) and σ is scale parameter. The marginal utility of income in location j is $MU_j = u'(W_j)$. With T1EV shocks, the choice probability takes the logit form:

$$P_j = \frac{\exp\left(\frac{1}{\sigma}[u(W_j) + a_j]\right)}{\sum_k \exp\left(\frac{1}{\sigma}[u(C_k) + a_k]\right)}$$

The cross-price effect is:

$$\frac{\partial P_j}{\partial w_k} = -\frac{1}{\sigma} \cdot P_j \cdot P_k \cdot u'(W_k)$$

The negative sign indicates substitution: raising wages in k draws people away from j. The magnitude depends on the baseline shares P_j and P_k, and critically, on the marginal utility of income in k.

B.2 Model 2: Multiplicative with Fréchet Shocks

An alternative specification, common in the quantitative spatial economics literature, specifies utility as multiplicative:

$$U_{ij} = w_j^{\alpha} \cdot A_j \cdot \epsilon_{ij} \tag{28}$$

where w_j is the wage in location j, A_j captures location amenities, and ϵ_{ij} is drawn iid Fréchet with shape parameter θ, so that $F(\epsilon) = \exp(-\epsilon^{-\theta})$.

With Fréchet shocks, the choice probability is:

$$P_j = \frac{(w_j^{\alpha} A_j)^{\theta}}{\sum_k (w_k^{\alpha} A_k)^{\theta}} \tag{29}$$

Taking logs, this can be rewritten as:

$$P_j = \frac{\exp\left(\theta[\alpha \ln w_j + \ln A_j]\right)}{\sum_k \exp\left(\theta[\alpha \ln w_k + \ln A_k]\right)} \tag{30}$$

which similar to Model 1 has a logit functional form.

The cross-price effect in the Fréchet model is:

$$\frac{\partial P_j}{\partial w_k} = -\theta\alpha \cdot P_j \cdot P_k \cdot \frac{1}{w_k} \tag{31}$$

C Appendix: Survey design

C.1 Details on design choices

Definition of a location: A location in our survey is defined as a metropolitan statistical area (MSA). We obtain the list of all states and MSAs within each state from the U.S. Census Bureau. We provide a drop-down menu of states followed by a drop-down menu of MSAs within the chosen state. At the beginning of the survey, we define a location as a metropolitan area: *Note: A metropolitan area is a large region comprising one or more counties, which includes a city with a significant population and its surrounding areas that are economically interconnected. It is defined by factors such as population size and the extent of linkage among nearby communities in terms of jobs, commuting, and transportation.*

Additionally, the three locations other than the respondent's hometown are required to be unique and different from the current location of the respondent and the location where they grew up.

Choice of language: We administered the survey in English and in Spanish. The Spanish translation was done by experts at the Rutgers Survey Center. We also conducted a pilot survey and a focus group discussion with a small sample of Spanish-speaking individuals to ensure that the translation was accurate and that the survey was understandable.

Displaying the exogenous location-specific attributes: We color coded the direction of the shocks to highlight the incomes to reduce cognitive load on the respondents. This decision was based on tests of whether the color coding had any effect on the responses in two simultaneous rounds of pilot data collection where we found no difference in the responses between the color-coded and non-color-coded versions of the pilot survey, except for shorter times to completion in the color-coded version. Thus, we decided to use the color-coded version in the final survey to reduce the cognitive load on the respondents. In order to avoid the possibility of any parallel being drawn with political affiliation, we specifically avoided using the colors blue and red. To represent increment (reduction) in income we used green (yellow). We did not color code the presence or absence of family and friends. We could have used the same color coding for the presence of family and friends, but we decided against it to avoid assuming whether the presence of family and friends increases or decreases utility, which in the case of income is more straightforward. In the same spirit of minimizing cognitive load, we showed respondent the shocked income instead of showing both the percentage change and the new income.

Ordering: Since we are identified within individuals, we are not concerned about the order of the locations. However, we randomized the order of the scenarios to avoid any ordering effects.

Facilitating ease in responses: We made several choices in order to facilitate ease in responding to the survey. Here we some of the most important ones:

- We provided a slider for the respondents to report their expected income and the compensating differential. This also avoids the possibility of different respondents using different units of income.
- In each exogenously varied scenario, respondents were shown their baseline probabilities and compensating differentials to facilitate comparison.
- In order to emphasize that their response in the compensating differential questions is only conditional on that location being their choice in 5 years, each option of location was followed with the phrase "*had I chosen to live there in 5 years*"

C.2 Background

The Survey Research Center at Rutgers University has helped us design and implement the survey. Our survey was administered in the July and November of 2024 to a nationally representative sample of adults. The survey was available in English and Spanish, and participants that completed the survey were given small payments in line with typical compensation given to online survey participants. The survey respondents themselves were anonymous, and no personally identifiable data was collected.

The survey started with a small pilot in April of 2024. Based on responses and feedback from the pilot, we adjusted survey questions and then administered the survey to a larger sample as a second pilot (July). Once we reviewed the results from that pilot and decided not to implement any changes, we then launched the full survey in November.[17]

C.3 Locations and Income

We start the survey by asking basic questions about current and previous location and current and expected future incomes.

- We first ask the state, metropolitan area or nearest city, and zip code that the respondent considers his or her home town, and then the same data for where the respondent currently lives.
- We next ask the respondent their total pre-tax income over the previous 12 months. If the respondent lives with one or more adults, the respondent is asked to only report income from his or her job.
- Then we ask the respondent to report his or her expected individual annual income 5 years from now.

[17] The second pilot tested whether responses varied depending on whether certain phrases were highlighted or not. The responses did not seem to depend on highlighting. For the full survey, we used the highlighted version.

- Then we ask the respondent their current total annual *household* income. Respondents are instructed to report income from all sources and include incomes from all adults in the household.

C.4 Baseline Data

The survey next asks respondents to choose the three most likely state and metropolitan areas where they may live 5 years from now. The first question is the most likely, the second is the second-most likely, and the third is the least likely.

To help fix ideas, going forward we will consider an example where the most likely location is Atlanta, GA, the second most likely location is Austin, TX, the third most likely location is Seattle, WA. In this example, the respondent's hometown is Madison, WI and current location is Washington, DC.

After the respondent reports the three places he or she is most likely to live, the respondent is asked to report the probabilities he or she will be "living in these locations 5 years from now." The set of locations includes the respondent's hometown, and the respondent is asked to make sure the probabilities sum to 100 percent. A sample response is below.

Madison, WI	10	%
Atlanta-Sandy Springs-Alpharetta, GA	20	%
Austin-Round Rock-Georgetown, TX	40	%
Seattle-Tacoma-Bellevue, WA	30	%
Total	100	%

We will call the collection of the three most likely locations respondents expect to live 5 years from now and the home town as the respondent's "likely future locations."

The respondent is then asked to report his or her expected annual total income in 5 years in each of the likely future locations. The respondent answers these questions by moving a slider, shown below.

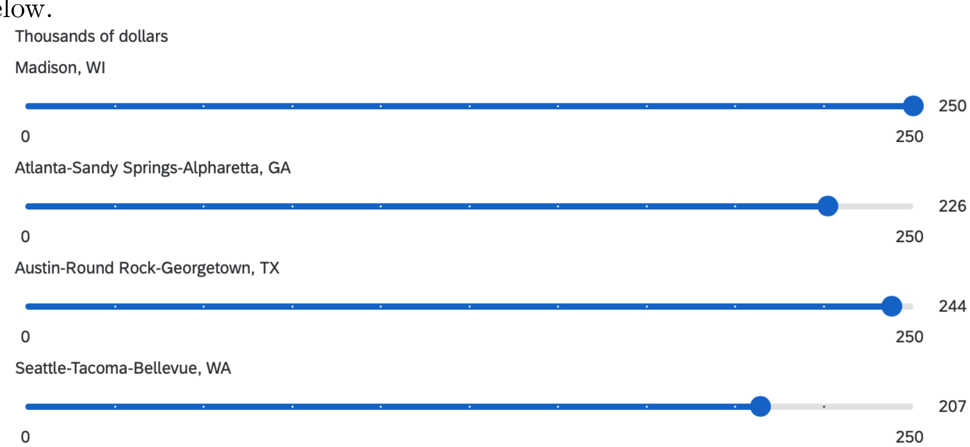

The next questions ask about the presence of savings and friends and family in each of the likely future locations as well as expected future costs of living.

- We first ask respondents if they have enough cash on hand to "cover expenses for a few months in the event of a job loss or a relatively small event causing some unanticipated expenses." A sample response is shown below.

- We next ask if respondents have a social network of family or friends in each of the likely future locations. In each location, the answer is restricted to be "yes" or "no."

- The respondent is then asked, *In the future, if your life circumstances become very difficult, in which location(s) do you expect to get the most support from your social network of family and friends?* The respondent is asked to rank the likely future locations from 1 (most likely to get support) to 4 (least likely).

- Next, the respondent is asked, *Thinking about your hobbies and things you like to do for entertainment, where do you think you will have the most fun?* The respondent is asked to rank the likely future locations from 1 (most fun) to 4 (least fun).

- Finally, the respondent is asked, *Given the sorts of goods and services you like to buy and given the kind of housing unit you wish to live in, which locations do you think would be most expensive for you?* The respondent is asked to rank the likely future locations from 1 (most expensive) to 4 (least expensive).

C.5 Experimental Variation in Hypothetical Scenarios

The next two blocks of questions are key for estimating the marginal utility of consumption. First, we ask how probabilities over location chioces change as assumptions about income and the presence of family and friends vary. Instructions leading up to this block of questions are as follows:

> *Earlier you told us four places you could imagine living 5 years from now. You also gave us your best guess about what your annual income would be in each location, whether family would be present, and your probability of living in each place.*
>
> *Now we are interested in how the probability of living in these locations would be affected if your anticipated income or family ties were changed, holding all other aspects of the location the same.*

For your convenience, we will also show your reported expectation of income and presence or not of family/friends in parenthesis below the changed income and family ties.

Respondents are shown new values of income and family and friends in each location, with the baseline values shown in parentheses. We select the new values quasi-randomly to span a large support of possible variation. The variation in incomes and presence of family and friends for each MSA, by scenario, is shown below

Table 9: Survey-induced exogenous variation across scenarios

	Option shown to respondent			
Scenario	Hometown	Alt. 1	Alt. 2	Alt. 3
Panel A: Income growth factor (multiplicative)				
1	1.2520	0.9162	1.3590	0.8913
2	0.7849	1.1480	0.8662	1.2660
3	1.2090	0.8617	0.6198	1.1050
4	0.8848	1.4440	1.2390	0.7793
5	0.7254	0.7256	1.3950	1.1870
6	1.0650	1.2130	0.9453	0.7797
7	0.8673	0.8899	0.7091	1.4020
8	1.4440	0.9491	1.1100	1.2090
Panel B: Family and friends indicator				
1	0	1	1	0
2	0	1	1	0
3	1	0	0	1
4	1	0	0	1
5	1	0	0	1
6	1	1	0	0
7	0	0	1	0
8	0	1	1	1

Notes: Each row is a scenario used in the survey. Columns correspond to the respondent's hometown option and three alternative locations. Panel A reports the randomized income growth multipliers applied to each option. Panel B reports the randomized presence of family and friends in each option (binary indicator).

Continuing with our example, in a first scenario respondents are shown the following:

Please report your revised chances of living in each location if, in 5 years, your expected household income in these locations and presence or not of family/friends are different than what you expected as shown below, while holding all other aspects of the locations the same.

Location	**Future Income**	**Family/ Friends**
Madison, WI	313k (was 250k)	No (was Yes)
Atlanta-Sandy Springs-Alpharetta, GA	206k (was 226k)	Yes (was No)
Austin-Round Rock-Georgetown, TX	332k (was 244k)	Yes (was Yes)
Seattle-Tacoma-Bellevue, WA	184k (was 207k)	No (was No)

Given the new assumptions, respondents are asked to recompute the probabilities they live in each of the likely future locations in 5 years:[18]

*Your **reported % chances of living in each location is indicated in parentheses.***
*Please make sure that **your new % chances sum to 100**.*

Madison, WI (10%)	New%	15
Atlanta-Sandy Springs-Alpharetta, GA (20%)	New%	15
Austin-Round Rock-Georgetown, TX (40%)	New%	45
Seattle-Tacoma-Bellevue, WA (30%)	New%	25
Total	New%	100

Respondents are shown 9 scenarios in total: baseline plus 8. Call this entire set of 9 scenarios as the "probability scenarios" which help us identify the elasticity of location choice with respect to expected income and the presence of friends and family.

The next block of questions allows for estimation of the marginal utility of consumption in each of the likely future locations, as we show later. Instructions for this block of questions are as follows:

> *The following questions explore a circumstance in which your income unexpectedly and permanently jumps **after** you have chosen your future location. We will then ask the same two questions under a number of different scenarios where your expected household income and presence or not of family/friends might differ from your stated expectations.*

[18]Note the probabilities at the baseline assumptions are shown in parentheses.

1. The first question will ask you to choose the location in which you would benefit most from the permanent jump in income, had you chosen to live in that location. Call the benefit arising from this permanent jump in income in the location you choose your "change in happiness." Notes that this location may not necessarily be the one where you are most likely to move. Simply think of where the permanent jump in income would be most helpful had you chosen to live there.

2. Next, we will ask you to imagine living in the other locations 5 years from now, and how much of a jump in income would be needed for your change in happiness at each place to be equal to the change in happiness at the location in the first question.

Respondents will be asked the same two questions nine times. The first time, they will be shown a scenario corresponding to their baseline assumptions. The next eight times, they will be shown scenarios that are the same as the scenarios shown in the probability questions. In each of these nine scenarios, respondents will be instructed to hold all factors other than income and the presence of family and friends the same.

- The first question is as follows: *Fast forward 5 years. Out of many possible future versions of yourself, think about 4 versions A, B, C, and D each of whom has chosen to live in one of the locations you have specified and with the incomes and presence of family/friends as given below.*

  In the first scenario, respondents are shown baseline assumptions

Location	Future Income	Family/ Friends
Madison, WI	250k	Yes
Atlanta-Sandy Springs-Alpharetta, GA	226k	No
Austin-Round Rock-Georgetown, TX	244k	Yes
Seattle-Tacoma-Bellevue, WA	207k	No

  The question continues: *Fast forward 5 years and imagine yourself in each of the locations below. In which location would you benefit the most from an unexpected and permanent increase of $5,000 in household income, had you chosen to live there? Note that this location may not necessarily be the one where you are most likely to move. Call the benefit arising from this permanent jump in income in the location you choose your "change in happiness."* Respondents choose a location from the list of likely future locations:

  - ○ Madison, WI , had I chosen to live there in 5 years
  - ◉ Atlanta-Sandy Springs-Alpharetta, GA , had I chosen to live there in 5 years
  - ○ Austin-Round Rock-Georgetown, TX , had I chosen to live there in 5 years
  - ○ Seattle-Tacoma-Bellevue, WA , had I chosen to live there in 5 years

- The second question is: *Now, for each of the other versions living in the other locations, indicate the size of the jump in household income that they would need living in those locations 5 years from now such that their change in happiness after receiving this jump in household income is the same as the change in happiness from the $5,000 jump in household income given to the version living in [Atlanta-Sandy Springs-Alpharetta, GA].*[19] Respondents then adjust a slider for each of the other three likely future locations to answer that question:

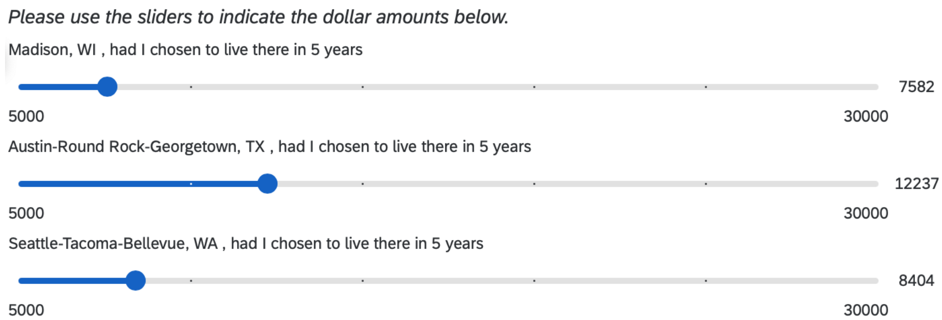

The respondents then answer the same 2 questions 8 more times, but with the assumptions about income and friends and family in each location changed, for example:

Location	Future Income	Family/ Friends
Madison, WI	313k (was 250k)	No (was Yes)
Atlanta-Sandy Springs-Alpharetta, GA	206k (was 226k)	Yes (was No)
Austin-Round Rock-Georgetown, TX	332k (was 244k)	Yes (was Yes)
Seattle-Tacoma-Bellevue, WA	184k (was 207k)	No (was No)

C.5.1 Economic Policy and Political Leanings

After we ask these 18 questions (2 questions under 9 sets of assumptions), we then ask a block of 7 questions about economic policy and political leanings that we borrow from Kuziemko et al. (2015). We are interested in whether political preferences on views on redistribution across agents are correlated with expected variation in their own marginal utility of income across potential likely future places. We list these questions below in italics:

- *On economic policy matters, where do you see yourself on the liberal-conservative spectrum?*. Respondents choose from "Very Liberal," "Liberal," "Moderate," "Conservative," and "Very Conservative."

- *In general, how important do you think it is to stay informed about economic policy?"* Respondents choose from "Very important," "Somewhat important," "Not very important," and "Not Important at all."

[19] Note that the location in brackets will be the location chosen as the answer to the previous question.

- *What would you say are the main reasons why you wish to be well informed about economic policy?* Respondents choose one of the following responses "Affects personal finances," "Affects business or profession," "Relevant to stock market and investments," "Economic issues are important politically and might affect my vote," and "To be a responsible citizen, I like to keep informed."

- *How knowledgeable do you consider yourself on economic policies and issues?* Respondents choose one of the following responses "Highly knowledgeable," "Somewhat knowledgeable," "Not very knowledgeable," and "Not knowledgable at all"

- *Do you agree that the government should assist people in response to temporary crises, for example cash transfers to people that lost their job during a recession or households that lost their house and possessions due to hurricane, fire, tornado, or flood?* Respondents choose one of the following response, "Definitely," Probably," "Probably not," and "Definitely not."

- *Do you agree that the government should implement welfare programs using tax dollars?* Respondents choose one of the following response, "Definitely," Probably," "Probably not," and "Definitely not."

- *Your responses show that a little extra money in certain scenarios is useful. Do you believe that the government should help individuals with some cash assistance in such scenarios.* Respondents choose one of the following responses, "Definitely," Probably," "Probably not," and "Definitely not."

C.5.2 Demographic Questions

In the final 9 questions of the survey, we ask each respondent to list (1) his or her age (restricted to one of 7 intervals), (2) gender (male, female, or other), (3) if the respondent is of Hispanic, Latino, or Spanish origin (yes/no), (4) ethnicity or race (one of six labels including "Other,"), (5) current marital status, (6) number of adults in the household, (7) children under 18 in the household, (8) highest level of education (one of nine categories), and (6) employment status over the past 12 months (one of seven categories).

D Empirical Bayes Algorithm

Consider a collection of n parameter estimates $\{\hat{\theta}_i\}_{i=1}^n$ with associated standard errors $\{\hat{\sigma_i}\}_{i=1}^n$. The hierarchical model is specified as:

$$\text{Truth:} \quad \theta_i \sim N(\mu, \tau^2) \qquad \text{(Prior)}$$
$$\hat{\theta}_i|\theta_i \sim N(\theta_i, \hat{\sigma}^2) \qquad \text{(Sampling)}$$

where:

- μ is the population mean, τ^2 is the prior variance
- θ_i represents the true parameter value for observation i
- $\hat{\sigma}^2$ is the estimated sampling variance. To estimate this, we:
    1. Estimate θ_i for each i separately using Canay (2011)
    2. Obtain residual vector for each i
    3. Stack residuals and estimate its variance to obtain $\hat{\sigma}^2$

Given the specified model, we employ an empirical Bayes approach where $\hat{\theta}_i \sim N(\mu, \hat{\sigma}^2 + \tau^2)$. This framework allows us to optimally combine individual estimates with information from the entire sample to reduce estimation error.

The log likelihood function for this specification can be expressed as:

$$\ell(\mu, \tau^2; \{\hat{\theta}_i, \hat{\sigma}^2\}) = -\frac{n}{2}\log(2\pi) - \frac{n}{2}\log(\hat{\sigma}^2 + \tau^2) - \frac{1}{2(\hat{\sigma}^2 + \tau^2)}\sum_{i=1}^{n}(\hat{\theta}_i - \mu)^2$$

Since the weights are all equal in our homoskedastic case, the estimate of the prior mean reduces to a simple average: $\hat{\mu} = \frac{1}{n}\sum_{i=1}^{n}\hat{\theta}_i$. Conditioning on this estimated mean, the likelihood function for τ^2 becomes:

$$\ell(\tau^2; \{\hat{\mu}, \hat{\theta}_i, \hat{\sigma}^2\}) = -\frac{n}{2}\log(2\pi) - \frac{n}{2}\log(\hat{\sigma}^2 + \tau^2) - \frac{1}{2(\hat{\sigma}^2 + \tau^2)}\sum_{i=1}^{n}(\hat{\theta}_i - \hat{\mu})^2$$

To estimate the variance component τ^2, we derive the first-order condition:

$$\frac{n}{2(\hat{\sigma}^2 + \hat{\tau}^2)} - \frac{1}{2(\hat{\sigma}^2 + \hat{\tau}^2)^2}\sum_{i=1}^{n}(\hat{\theta}_i - \hat{\mu})^2 = 0$$

Solving this equation yields the maximum likelihood estimate:

$$\hat{\tau}^2 = \frac{1}{n}\sum_{i=1}^{n}(\hat{\theta}_i - \hat{\mu})^2 - \hat{\sigma}^2 \tag{32}$$

This approach produces identical shrinkage factors for all observations, where $\lambda_i = \lambda = \frac{\hat{\tau}^2}{\hat{\sigma}^2+\hat{\tau}^2}$.

The resulting empirical Bayes estimate for each observation is a weighted average: $\tilde{\theta}_i = (1-\lambda)\hat{\mu} + \lambda\hat{\theta}_i$, reflecting the balance between individual estimates and the sample-wide mean in proportion to the estimated signal-to-noise ratio. The overall signal-to-noise ratio, identical for each observation i, can be expressed as $\frac{\hat{\tau}^2}{\hat{\sigma}^2} = \frac{1}{n}\sum_{i=1}^{n} \frac{(\hat{\theta}_i-\hat{\mu})^2}{\hat{\sigma}^2} - 1$.