Tax Progressivity and Mobility Costs

Daniele Coen-Pirani*

March 22, 2025

Abstract

This paper examines how mobility costs influence the effectiveness and desirability of tax progressivity using a general equilibrium spatial model. A key feature of the model is that workers' idiosyncratic productivity depends on location. The interaction of amenities, idiosyncratic shocks and moving costs implies that progressive taxation distorts location choices by reducing incentives for agents to relocate to their most productive areas. Using a quantitative framework, I find that the negative effect of tax progressivity on output is weakest when mobility costs are either relatively low or high. The optimal degree of tax progressivity balances the costs of spatial tax distortions against the benefits of enhanced insurance, leading to relatively high optimal progressivity at both extremes of mobility costs.

Keywords: Spatial economics, Tax progressivity, Spatial tax distortions, Mobility costs, Welfare.

*Department of Economics, University of Pittsburgh. E-mail: coen@pitt.edu. Thanks to Mark Colas, Hal Cole, Ed Glaeser, Andreas Ferrara, Wenli Li, Fernando Parro, Marla Ripoll, Juan Carlos Suarez-Serrato, Pin Sun, and conference and seminar participants for useful feedback and comments on earlier drafts of this paper. An earlier draft of this paper was titled "Tax Progressivity and Barriers to Internal Labor Migration: An Equivalence Result." All remaining errors are my own.

1 Introduction

The analysis of optimal degree of progressivity in the tax and transfer system aims to theoretically isolate and empirically quantify the benefits and costs of redistribution. On the benefit side, greater progressivity reallocates resources toward individuals with lower initial human capital and provides insurance against uninsured idiosyncratic shocks (see e.g., Heathcote, Storesletten and Violante (2017)). On the cost side, higher progressivity may generate deadweight losses by discouraging labor supply and investment in human capital (see e.g. Krueger and Ludwig (2013) and Guvenen, Kuruscu and Ozkan (2014)). Importantly, much of the debate on optimal tax progressivity has been framed in models that abstract from spatial considerations, assuming a single, economy-wide labor market where workers and firms interact.

A parallel literature (Albouy (2009), Eeckhout and Guner (2017)) argues that in spatial models, progressive taxation introduces additional distortions by influencing not only labor supply but also location choices. In such models, workers select their location based on a utility index that includes consumption of non-tradeable goods and amenities, in addition to consumption of tradeable goods. A progressive tax on nominal earnings distorts location choices toward areas where utility is derived from untaxed factors, such as high amenities or low housing costs.

While this literature has provided valuable insights into the interaction between taxation and location choices, the baseline Roback (1982)-style model abstracts from moving costs. However, both casual observation and rigorous estimation suggest that moving costs are a significant spatial friction shaping mobility decisions. In the U.S. about 60% of U.S. born male workers aged 15-65 works in the same state in which they were born and Zabek (2024) reports that half of U.S.-born adults live within 50 miles of their birthplace. Kennan and Walker (2011) and Giannone et al. (2023) estimate structural models of migration and find moving costs in the order of hundreds of thousands of dollars.[1]

This paper examines how mobility costs costs influence the effectiveness and desirability of tax progressivity using a general equilibrium spatial model. The framework builds on Bryan and Morten (2019), where agents choose their work location based on pairwise moving costs relative to their birth location, location-specific idiosyncratic productivity shocks, and differences in wages, rents, and amenities. I extend this model by incorporating a progressive tax-transfer system, endogenous labor supply, and a rental housing market. When interacted

[1] Moving costs are often mentioned as a source of misallocation and reduced output, especially in developing countries (e.g. Bryan, Chowdhury and Mobarak (2014), Bryan and Morten (2019) and Lagakos, Mobarak and Waugh (2023)). Heise and Porzio (2022) emphasize the importance of moving costs and home bias for location choices in Germany.

with moving costs, idiosyncratic productivity shocks introduce a novel source of spatial tax distortions, beyond the one due to the heterogeneity in amenities and rents across locations emphasized by Albouy (2009), Eeckhout and Guner (2017), and Colas and Hutchinson (2021). Specifically, higher tax progressivity weakens incentives for individuals to locate in places where they are idiosyncratically more productive, regardless of location-wide productivity and amenities.[2] I refer to both of these distortions - those arising from amenities and rents, as well as those stemming from idiosyncratic productivity shocks and mobility costs - as the "spatial tax distortion channel."

The model is calibrated to match key U.S. data on bilateral migration rates across states, migrant earnings, and heterogeneous housing supply elasticities. In the quantitative model, higher tax progressivity reduces aggregate output through the spatial tax distortion channel and by discouraging labor supply. A reform replacing the benchmark progressive tax with a linear tax increases aggregate output by about 10%, with the spatial tax distortion channel accounting for roughly one-third of this increase and labor supply responses making up the rest. Despite these distortions, the optimal degree of tax progressivity is positive and increases with workers' risk aversion, as it provides insurance against idiosyncratic productivity shocks and differences in their initial conditions.

With this benchmark model in place, I address two key questions. First, how do moving costs shape the response of aggregate output to changes in tax progressivity? Second, how does optimal tax progressivity depend on moving costs?

For the first question, I analyze the same linear tax reform under different moving cost scenarios. Under benchmark moving costs, the reform reduces spatial tax distortions, leading to a 3.5% increase in output. With zero moving costs, this effect diminishes to 1.8%, about half the benchmark magnitude. Interestingly, high moving costs also weaken the spatial tax distortion channel, resulting in a similar output response when moving costs are twice the benchmark level.

This non-monotonic effect arises due to two opposing forces. When spatial frictions are high, only workers with substantial pre-tax earnings gains choose to move, so the pool of marginal agents who are potentially affected by higher tax progressivity is small. However, because their gains are large, higher tax progressivity strongly discourages their migration. When spatial frictions are low, many workers are on the margin between moving and staying, but their individual gains from migration are smaller, reducing the tax distortion effect. At both extremes, the spatial tax distortion is relatively modest, peaking at intermediate levels of moving costs.

[2] As such this mechanism is active even if locations are ex-ante symmetric in terms of amenities, productivity and moving costs, as in the dynamic migration model in Coen-Pirani (2021).

For the second question, I find that optimal tax progressivity follows a similar non-monotonic pattern. It is relatively high when moving costs are either low or high and lowest at intermediate moving costs. This mirrors the response of aggregate output: when spatial distortions are small, the planner sets higher progressivity, whereas at intermediate levels, the planner optimally reduces it. This result might help explain why European countries exhibit lower internal labor mobility than the U.S. and more generous social insurance systems (Hassler et al. (2005)).

This paper contributes to the literature on spatial economics and taxation, particularly the effects of national tax policies on the geographic distribution of economic activity.[3] Compared to prior work, it offers several key contributions.

First, much of the existing literature on spatial tax distortions has primarily examined the welfare benefits of policies designed to mitigate or neutralize these distortions. For example, Albouy (2009) and Colas and Hutchinson (2021) estimate the welfare gains from replacing the U.S. tax system with lump-sum taxation. In contrast, this paper goes beyond evaluating policies - such as a linear tax - that minimize spatial tax distortions. Instead, it provides a broader analysis of optimal tax progressivity, explicitly accounting for both non-spatial labor supply distortions and the insurance benefits of progressive taxation.

While labor supply distortions have been a key focus in research on optimal tax progressivity (Diamond and Saez (2011)), standard spatial models typically assume inelastic labor supply. By incorporating labor supply distortions, this paper allows for a direct quantification of their relative importance within a unified framework - a comparison that has not yet been explored in the literature.

Among other roles, progressive taxation serves as a mechanism for insuring workers against idiosyncratic earnings risk, a novel feature of the paper that is central to research on non-spatial optimal taxation (Heathcote, Storesletten and Violante (2017)). As a result of these shocks, the earnings distribution in each location has full support, and progressive taxation redistributes income within locations, not just across them, as in common in much of the spatial literature.[4] The analysis of insurance also highlights the role of risk aversion in shaping optimal tax policy. In spatial models agents make decisions after observing idiosyncratic shocks, so risk aversion plays no role in location choices. However, risk aversion

[3] My focus is on national tax progressivity, rather than state-level taxation, which is absorbed into local amenities in the model. While I do not explicitly examine state and local tax heterogeneity, additional distortions from such variation have been studied by Desmet and Rossi-Hansberg (2013) and Fajgelbaum et al. (2019).

[4] A Theil index decomposition of earnings inequality into within and between U.S. states components reveals that the within-state component accounts for over 96% of earnings inequality. In the spatial literature, only Colas and Hutchinson (2021)'s model considers heterogeneous agents (of two ex-ante different types, skilled and unskilled workers), but they abstract from uninsured earnings risk.

is fundamental in determining optimal tax progressivity because the social planner evaluates welfare ex-ante, before individual shocks are realized.

A second key distinction from the existing literature is the explicit focus on moving costs. Colas and Hutchinson (2021) explore spatial tax distortions in a model with college- and non-college-educated workers but emphasize distributional effects rather than the role of moving costs in tax reforms and welfare. My model directly assesses the implications of moving costs for optimal tax progressivity.[5] Related work by Zabek (2024) and Zerecero (2021) examines how home bias - the preference for residing in one's birthplace - affects the effectiveness of place-based policies, concluding that home bias reduces the distortionary impact of such policies. This finding parallels my result that tax progressivity is less distortionary when moving costs are high. A novel contribution of this paper is showing that the same holds when moving costs are relatively low.

A feature of my approach is to study optimal tax progressivity policies using a specific tax-transfer function.[6] This functional form closely approximates the existing U.S. tax and transfer system (Heathcote, Storesletten and Violante (2017)), but constrains the set of alternative policies that a planner might implement. An alternative approach, as pursued by Fajgelbaum and Gaubert (2020) in a general spatial model, is to derive first-best allocations and the implied optimal spatial policies. The main advantage of my approach relative to theirs is its relative simplicity, which allows for a tractable analysis of optimal tax progressivity.

The remainder of the paper is structured as follows. Section 2 introduces the model. Section 3 describes the calibration of the model. Section 4 discusses the impact of varying tax progressivity on the model's equilibrium and welfare. Section 5 analyzes the link between moving costs, tax progressivity, output and welfare. Section 6 reports sensitivity analysis results. Section 7 concludes. The paper's online appendix provides additional details on the model, its results, and the data used to calibrate it.

2 Theoretical Framework

In this section I introduce the model economy, define the competitive equilibrium, and discuss the welfare criterion.

[5] I abstract from idiosyncratic preference shocks - a labor mobility friction sometimes included in spatial models (e.g. Colas and Hutchinson (2021)) - as they alone cannot explain why most workers remain in their birth state.

[6] This tax function was first introduced by Feldstein (1969) and Persson (1983), and more recently employed by Benabou (2002) and Heathcote, Storesletten and Violante (2017) in their analysis of optimal tax progressivity in non-spatial environments.

2.1 Workers' Decision Problem

The economy consists of J locations and a continuum of measure one of workers who are initially exogenously distributed across origin locations, indexed by o. Each worker makes a one-time decision about where to live and work and how much to work.

Each worker's human capital h_{ido} has a destination-dependent idiosyncratic component, denoted by s_{id} and an origin-dependent component, denoted by q_o:

$$h_{ido} = q_o s_{id}. \tag{1}$$

The destination-specific components $\{s_{id}\}_{d=1}^{J}$ are drawn from a Fréchet distribution with cumulative distribution function:

$$F(s_{i1}, s_{i2}, ..., s_{iJ}) = \exp\left(-\left(\sum_d s_{id}^{-\frac{\tilde{\theta}}{1-\rho}}\right)^{1-\rho}\right), \tag{2}$$

where $\tilde{\theta}$ measures the extent of skill dispersion and ρ the correlation in skills across locations.

The pretax earnings of an individual originating from o is equal to the product of their human capital, their labor effort l, and the wage per efficiency units of skill supplied in d, denoted by w_d:

$$y_{ido} = w_d h_{ido} l. \tag{3}$$

The national government's tax and transfer system maps an individual's market earnings into their post-tax and transfer earnings $\tilde{y}_{ido}$:

$$\tilde{y}_{ido} = \chi y_{ido}^{1-\mu}. \tag{4}$$

The net taxes paid by an agent with market income y are therefore $T(y) = y - \chi y^{1-\mu}$. The parameter $\mu < 1$ captures the degree of tax progressivity. In particular, $(1-\mu)$ is equal to elasticity of after-tax income to before-tax income:

$$1 - \mu = \frac{1 - T'(y)}{1 - T(y)/y} \text{ for each } y.$$

If marginal and average tax rates are equal - a linear tax system - the parameter $\mu = 0$. If the marginal tax rate is higher than the average tax rate, instead, the tax system is progressive and $\mu > 0$. The parameter χ, instead, determines the aggregate net taxes paid by agents in the economy, with higher values of χ corresponding to lower aggregate net taxes. Using data on pre-government (y) and post-government $(\tilde{y})$ household income for the early 2000s, Heathcote, Storesletten and Violante (2017) show that the functional form in equation (4) provides a very good approximation of the U.S. tax and transfer system.[7]

An agent's utility function takes the following form:

$$u_{do}\left(c,z,l\right)=\alpha_{d}\left(1-\tau_{do}\right)\left(\frac{c}{1-\beta}\right)^{1-\beta}\left(\frac{z}{\beta}\right)^{\beta}\exp\left(-\frac{\varphi}{\xi}l^{\xi}\right), \tag{5}$$

where c is consumption of the final good, z is housing consumption, and l denotes time spent working. The term in l in equation (5) represents the disutility of labor effort. Utility increases in the parameter α_d which denotes location d's exogenous amenities, inclusive of state-specific linear taxes and public goods. Higher mobility costs between o and d - denoted by τ_{do} - reduce utility.

The budget constraint of an agent with human capital h_{ido} at a destination d is:

$$c+r_{d}z=\chi\left(w_{d}h_{ido}l\right)^{1-\mu}, \tag{6}$$

where r_d denotes rents per unit of housing in d.

Maximizing equation (5) subject to (6) yields the following optimal consumption and labor supply choices:

$$c_{ido}=\left(1-\beta\right)\chi\left(w_{d}h_{ido}l^{*}\right)^{1-\mu}, \tag{7}$$

$$z_{ido}=\beta\chi\left(w_{d}h_{ido}l^{*}\right)^{1-\mu}r_{d}^{-1}, \tag{8}$$

$$l^{*}=\left(\frac{1-\mu}{\varphi}\right)^{\frac{1}{\xi}}. \tag{9}$$

Equations (7) and (8) show that agents spend constant fractions of their after-tax earnings on consumption and housing. Equation (9) represents the optimal supply of labor. The latter is independent of the level of wages and human capital as the Marshallian elasticity of labor supply is equal to zero.[8] Notice, however, that higher tax progressivity discourages labor supply. The extent to which this happens depends on the parameter ξ, with higher values

[7] Specifically, they regress $\ln\tilde{y}$ on $\ln y$ and estimate $\mu=0.18$, with a regression R^2 equal to 0.91.

[8] A zero Marshallian elasticity is consistent with the data I use to calibrate the model. In the 2000 Census data, the elasticity of annual weeks of work to weekly wages for U.S. born males aged 15-65 is very close to zero. The estimated elasticity is 0.019 with a standard error 0.0008. A similar result holds if instead of annual weeks worked and weekly wages, the elasticity is computed using annual hours of work and hourly wages.

of ξ being associated with a smaller negative effect of tax progressivity on labor supply.[9]

Replacing equations (7)-(9) into (5) I obtain the agent's indirect utility function:

$$U_{ido} = \alpha_d r_d^{-\beta} (1 - \tau_{do}) \chi (w_d h_{ido} l^*)^{1-\mu} \exp\left(-\frac{\varphi}{\xi} (l^*)^{\xi}\right), \tag{10}$$

Each agent i chooses a destination location d to maximize their indirect utility function:

$$\max_j U_{ijo}.$$

With Fréchet shocks, the fraction of the population initially in o that moves to d is given by:

$$\pi_{do} = \frac{\tilde{w}_{do}^{\theta}}{\sum_{j=1}^{J} \tilde{w}_{jo}^{\theta}}, \tag{11}$$

where $\theta \equiv \tilde{\theta}/(1-\rho)$. The term $\tilde{w}_{do}$ represents the composite fundamental that matters for migration decisions:

$$\tilde{w}_{do} = \left[\alpha_d r_d^{-\beta} (1 - \tau_{do})\right]^{\frac{1}{1-\mu}} w_d. \tag{12}$$

The latter depends on the wage per unit of human capital, amenities, rents, moving costs, and tax progressivity. Notice, instead, that the level of taxes (χ), labor supply (l^*), and the human capital parameter q_o don't matter for migration choices because they do not vary by destination.

2.2 Technology, Housing Markets, and the Government Budget Constraint

The final (numeraire) consumption good is produced using a variety of intermediate inputs (Armington (1969)) according to the production function:

$$Y = \left(\sum_{d=1}^{N} y_d^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}}, \tag{13}$$

where $\sigma > 1$ denotes the elasticity of substitution among inputs. Each variety is produced in a location d according to the linear production function:

$$y_d = A_d L_d, \tag{14}$$

[9]The parameter φ determines the overall level of labor supply.

where L_d denotes aggregate effective hours of work in d. Due to agglomeration effects, labor productivity may depend positively on L_d:

$$A_d = \bar{A}_d L_d^{\gamma}, \tag{15}$$

with $\gamma \geq 0$. The location-specific supply of aggregate effective hours is given by:

$$L_d = \sum_o v_o \pi_{do} \bar{h}_{do} l^*, \tag{16}$$

where v_o denotes the (exogenous) mass of agents initially located in o. The average human capital of workers who migrate from o to d is defined as:

$$\bar{h}_{do} = q_o E_{do}\left[s_{id}\right], \tag{17}$$

where $E_{do}\left[.\right]$ denotes the expectation over the distribution of idiosyncratic shocks conditional on an agent choosing to move from o to d. The properties of the Fréchet distribution imply that the conditional expectation in equation (17) takes the following form:

$$E_{do}\left[s_{id}\right] = \pi_{do}^{-\frac{1}{\tilde{\theta}}} \Gamma\left(1 - \frac{1}{\tilde{\theta}}\right), \tag{18}$$

where $\Gamma(.)$ denotes the gamma function.[10]

Finally, assume that the inverse supply function for housing in location d is given by:[11]

$$r_d = B_d T_d^{\psi_d}, \tag{19}$$

where T_d denotes the quantity of housing, B_d is a location-specific parameter, and ψ_d the inverse elasticity of housing supply in d. Land is assumed to be owned by absentee landlords who receive all rent payments.

The aggregate demand for housing in a location reflects the fact that each agent spends a portion β of their disposable income on housing:

$$r_d T_d = \beta \sum_o v_o \pi_{do} E_{do}\left[\tilde{y}_{ido}\right]. \tag{20}$$

[10] See the online appendix. The term $\pi_{do}^{-\frac{1}{\tilde{\theta}}}$ in equation (17) captures a selection effect: as more individuals move from o to d, the idiosyncratic location shock of the marginal mover declines.

[11] The latter can be derived from a constant returns to scale production function for housing whose inputs are land and units of the final good (see the online appendix for details). Differently from Bryan and Morten (2019), I explicitly model the housing market to capture the direct effect of income redistribution on equilibrium rents.

where the average after-tax income of those who move from o to d is:

$$E_{do}\left[\tilde{y}_{ido}\right]=\chi\left(w_d q_o l^*\right)^{1-\mu}\pi_{do}^{-\frac{1-\mu}{\theta}}\Gamma\left(1-\frac{1-\mu}{\tilde{\theta}}\right).$$

The model is closed by the national government's budget constraint:

$$\sum_{d,o} v_o \pi_{do} E_{do}\left[T\left(y_{ido}\right)\right]=g, \tag{21}$$

where the left-hand side of this equation denotes aggregate net taxes and the right-hand side represents government consumption, denoted by g.

Government policy is represented by the triple (μ, χ, g). Given any two policy parameters in this triple, the third one is determined by the government's budget constraint, equation (21). In what follows, I treat μ parametrically and assume that χ adjusts to clear the government's budget constraint, while government purchases g remain constant in all counterfactual experiments. Therefore, I abstract from fiscal externalities arising from changes in public goods provision when tax progressivity varies.[12]

I conclude the model description by defining a competitive equilibrium.

Definition 1. *Given a tax progressivity policy μ and an amount of public good consumption g, a competitive equilibrium for this economy consists of location-specific wages, rents, and prices of intermediate goods $\{w_d, r_d, p_d\}$, an allocation $\{Y, L_d, H_d, l^*, y_d, \pi_{do}, y_{ido}, \tilde{y}_{ido}\}$ and a government policy χ such that:*

1. Given prices, intermediate producers in d choose effective hours L_d to maximize profits:

$$\max_{L_d}\left(p_d A_d-w_d\right)L_d. \tag{22}$$

where p_d denotes the relative price of good d in units of numeraire.

2. Given prices, the final goods sector chooses $\{y_d\}$ to maximize profits:

$$\max_{\{y_d\}}\left(\sum_{d=1}^{J} y_d^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}}-\sum_{d=1}^{J} p_d y_d. \tag{23}$$

3. Given prices, workers choose location and labor supply optimally, so that migration probabilities are given by π_{do}, defined in equation (11) and labor supply is given by equation (9).

[12] I assume that g affects the utility function (10) multiplicatively or additively, so it doesn't influence agents' location or labor supply behavior. Since g is kept constant in all counterfactual experiments, I have omitted including it explicitly in the utility function (10) to avoid burdening the notation.

4. Each location's supply of aggregate effective hours, output, and productivity are given, respectively, by equations (16), (14), and (15).

5. Equilibrium rents satisfy equations (19) and (20).

6. *Each agent's market income y_{ido} is given by equation (3) and their post-tax income $\tilde{y}_{ido}$ by equation (4).*

7. The policy parameters χ is such that the government's budget constraint in equation (21) is balanced.

2.3 Welfare

Welfare is defined as the expected utility of being born in this economy:

$$W = \sum_o v_o E\left[\max_j u\left(U_{ijo}\right)\right], \tag{24}$$

where the expectation $E[.]$ is taken with respect to the distribution of idiosyncratic location-specific $\{\varepsilon_d\}$ shocks and v_o is the probability of being born in location o.

The welfare function depends on the increasing and weakly concave function $u(.)$, which captures agents' risk attitudes. This is an important point that warrants discussion. In this economy, individuals make decisions after observing the realization of idiosyncratic shocks, meaning they do not face risk in their decision-making. As a result, a monotonic transformation of their original utility function via $u(.)$ does not influence migration or labor supply choices.

However, agents are subject to ex-ante risk due to the randomness of both their idiosyncratic shock vector and their birthplace o. Since welfare is evaluated ex-ante - before this uncertainty is resolved - the function $u(.)$ plays a crucial role in shaping welfare outcomes. In what follows, I assume that the function $u(.)$ takes the following iso-elastic form:

$$u\left(x\right) = \begin{cases} \frac{x^{1-\phi}}{1-\phi} & \text{if } \phi \\ \ln x & \text{if } \phi = 1 \end{cases}, \tag{25}$$

where $\phi \geq 0$ is the coefficient of relative risk-aversion. A higher degree of risk-aversion strengthens the case for higher tax progressivity, as it enhances the benefits of insurance against low realizations of idiosyncratic shocks.

As shown in the appendix, when $\phi \neq 1$ welfare is given by:[13]

$$W = (1-\phi)^{-1}\left((l^*)^{1-\mu}\exp\left(-\frac{\varphi}{\xi}(l^*)^{\xi}\right)\right)^{1-\phi}\mathit{\Gamma}\left(1-\frac{(1-\phi)(1-\mu)}{\tilde{\theta}}\right)\sum_o v_o m_o^{1-\phi}, \quad (26)$$

where $m_o = \chi q_o^{1-\mu}\left(\sum_d \tilde{w}_{do}^{\theta}\right)^{\frac{1-\mu}{\theta}}$. I summarize welfare effects using an equivalent variation measure: the proportion of current consumption an agent would need to receive in the benchmark economy to be indifferent between being born in that economy with welfare $W$and being born in an economy where a counterfactual policy is implemented, achieving welfare W'.[14]

Tax progressivity influences welfare through a variety of channels reviewed below.[15]

Non-spatial labor supply distortions. Tax progressivity introduces a wedge between the private and social returns to exerting labor effort, a non-spatial distortion. In particular, with a progressive tax system the social return is higher than the private return because an individual's labor effort also benefits the recipients of the transfers that are funded by a progressive tax. Private and social returns to labor effort are equalized under a linear tax system. Thus, maximizing this portion of the welfare function - the first term in parenthesis in equation 26 - would call for setting $\mu = 0$.

Spatial distortions. Higher tax progressivity reduces aggregate output by decreasing workers' incentives to locate in areas where they are individually or collectively more productive. Moreover, the planner, differently from individual workers, internalizes local agglomeration effects, and so might have an incentive to increase the concentration of workers in the most productive locations where the benefits of agglomeration are highest. For these reasons, output maximization calls for a regressive tax system ($\mu < 0$). The potential of this policy to increase output rises with the elasticity of housing supply in the most productive states.[16]

Insurance and redistribution among workers. In this model, no market mechanism allows agents with different birth locations and idiosyncratic productivity shocks to

[13] See the online appendix for the expression corresponding to the case $\phi = 1$.

[14] Formally, given the transformation (25), the equivalent variation is $ev = \left(\frac{W'}{W}\right)^{\frac{1}{1-\phi}} - 1$ if $\phi \neq 1$ and $ev = \exp\left(W(\mu') - W(\mu)\right) - 1$ if $\phi = 1$.

[15] The online appendix illustrates the various components of the welfare function analytically using a simplified version of the model.

[16] As pointed out by Hsieh and Moretti (2019) and others, some of the most productive areas of the U.S. tend to be characterized by relatively low housing elasticities. Among the 48 continental states, the five states with the highest productivity $\bar{A}_d$ (California, Connecticut, New Jersey, New York, and Illinois) are characterized by an average housing supply elasticity of 0.51, while for the rest of the states the average housing supply elasticity is 0.67. See Section 3 for more details on these calculations.

equalize their marginal utilities of consumption. The planner can in principle increase welfare through this channel by selecting a progressive tax system. The incentive to do so increases with agent's coefficient of relative risk aversion, ϕ. The benefits of insurance provision by a progressive tax system have been highlighted in the context of incomplete markets models (e.g. Heathcote, Storesletten and Violante (2017)). In a spatial setting such as this one, there are additional reasons to select a relatively high degree of tax progressivity because agents respond to geographic shocks by selecting whether and where to migrate. Specifically, the indirect utility function in equation (10) indicates that, due to moving costs, migrating agents experience a lower marginal utility of consumption compared to those who remain in place. Furthermore, it suggests that workers relocating to high-amenity and low-rent locations have a higher marginal utility of consumption than those settling in low-amenity areas. Both of these spatial considerations increase the planner's incentives to select a relatively high degree of tax progressivity.

Redistribution between workers and landlords. The welfare function in equation (26) reflects only the utility of workers and not that of landlords. Since higher tax progressivity induces workers to move to lower rent areas, it redistributes resources away from landowners and towards workers. Therefore, this effect calls for a higher degree of tax progressivity.

3 Quantitative Implementation

To implement the model quantitatively, I set some parameters a-priori while measuring others using the structural equations of the model.

The parameters set a-priori are $\{\theta, \rho, \sigma, \gamma, \mu, \xi, \varphi, \beta, g, (\psi_d)\}$. I set some of the elasticity parameters to the values estimated by Bryan and Morten (2019) using U.S. data, in order to increase comparability with their results: $\theta = 27.6$, $\rho = 0.9$, $\sigma = 8$, $\gamma = 0.05$. The benchmark tax progressivity parameter μ and the labor supply elasticity parameter ξ are set to 0.18 and 3 respectively, following Heathcote, Storesletten and Violante (2017). These authors set $\xi = 3$ to match a Frisch elasticity of labor supply - which corresponds to $1/(\xi - 1)$ in this model - equal to $1/2$ (see the review paper by Keane (2011)). The parameter φ is set so that labor supply l^* in the benchmark calibration is normalized to one.

The housing share parameter is set to $\beta = 0.3$, which is consistent with the estimates by Davis and Ortalo-Magne' (2011) for the 2000 Census. Government purchases are assumed to be equal to 9% of aggregate output in the benchmark economy. This corresponds to the U.S. federal government's average consumption share of the sum of private and federal government's consumption for the period 2000-2006.

To set the state-level housing supply elasticity parameters $\{\psi_d\}$ I follow the aggregation methodology proposed by Baum-Snow and Han (2023).[17] The latter show how, under some assumptions, it is possible to aggregate their Census tract-level estimates of housing supply elasticities to larger spatial units. The average inverse housing supply elasticity I obtain is 1.58 with a cross-state standard deviation of 0.31.[18] For comparison, the average inverse supply elasticity estimated by Colas and Hutchinson (2021) is 0.57, which is also close to Saiz (2010)'s estimate. Thus, the quantitative model features a relatively large response of rents to exogenous changes in housing demand. The estimated inverse elasticities are positively correlated (correlation coefficient, 0.41) with productivity $\bar{A}_d$, so housing supply is steeper in more productive states.

Given these parameters, I compute the remaining parameter values $\left\{\bar{A}_d, B_d, \alpha_d, \tau_{do}, q_o\right\}$ by requiring the model to be consistent with the observed bilateral migrations flows, bilateral average weekly wages, and each state's average rents. In doing so I take as given the initial distribution of population by state of birth $\{v_o\}$.

The data inputs I use are the empirical counterparts of bilateral migration flows, the initial distribution of population by state of birth v_o, and the average weekly wages of individuals born in a state o that reside in state d. This data is drawn from the 2000 U.S. Census of Population (Ruggles et al. (2021)), with a location o representing an individual's state of birth, and a location d representing the state where the agent is observed at the time of the 2000 Census. The sample consists of male workers ages 15-65, born in the U.S., who are head of household, with positive and non-missing data on earned income and weeks worked. I drop from the analysis individuals born or residing in Alaska, Hawaii and the District of Columbia.

When bringing the model to the data, I assume that migration flows and bilateral average wages are observed with (classical) measurement error. In particular, I assume that we observe:

$$f^*_{do} = f_{do} \exp u^f_{do},$$

$$\overline{\text{wage}}^*_{do} = \overline{\text{wage}}_{do} \exp u^w_{do}, \tag{27}$$

where f_{do} are the migration flows (in levels) from o to d, $\overline{\text{wage}}_{do}$ represents the average

[17] Thanks to an anonymous referee for suggesting this approach.

[18] The specific aggregation assumption I make, following Baum-Snow and Han (2023) is that, in response to shifts to housing demand at the state-level, "agents get redistributed across neighborhoods in a way that maintains the same relative home prices across neighborhoods." The online appendix presents additional descriptive statistics on the distribution of housing supply elasticities across states.

wages in d of agents who move from o to d, and a star denotes an observed variable. The error terms u's are independently and identically distributed random variables such that $E[u^f_{do}] = E[u^w_{do}] = 0$. By definition, the observed migration rates are related to the actual (unobserved) ones by the following relationship:

$$\ln \pi^*_{do} = \ln \pi_{do} + \ln \frac{v_o}{v^*_o} + u^f_{do}, \tag{28}$$

where $v_o = \sum_d \exp\left(-u^f_{do}\right) f^*_{do}$ and $v^*_o = \sum_d f^*_{do}$ are, respectively, the original population of o and its observed counterpart in the data.

The data and the parameters set a-priori are used to estimate the moving cost parameters $\{\tau_{do}\}$, the amenity parameters $\{\alpha_d\}$, and the measurement errors $\left\{u^f_{do}, u^w_{do}\right\}$. The moving cost parameters $\{\tau_{do}\}$ are estimated from the migration probability data, under the assumption that they are symmetric ($\tau_{do} = \tau_{od}$) and the normalization $\tau_{oo} = 0$. To do so, replace equation (12) into equation (11), use equation (28) and re-arrange to obtain the estimating equation:

$$\ln \pi^*_{do} = \ln (1 - \tau_{do})^{\frac{\theta}{1-\mu}} + \ln \left(\left(\alpha_d r_d^{-\beta}\right)^{\frac{\theta}{1-\mu}} w^\theta_d \right) - \ln \sum_{j=1}^{J} \tilde{w}^\theta_{jo} + \ln \frac{v_o}{v^*_o} + u^f_{do}. \tag{29}$$

This is a regression equation with origin-by-destination symmetric fixed effects to capture the moving cost term, destination fixed effects that reflect the wage-amenity-rent composite term $\left(\alpha_d r_d^{-\beta}\right)^{\frac{1}{1-\mu}} w_d$ at destination, and origin fixed effects that capture the denominator of the location choice probabilities and the ratio v_o/v^*_o. The estimates of moving costs can be obtained from the origin-by-destination symmetric fixed effects, given values of θ and μ. For simplicity, I estimate equation (29) by ordinary least squares because the bilateral 48×48 matrix of migration flows has only five zeroes.[19]

The average estimated moving cost is 0.14, which is close to the value of 0.15 reported by Bryan and Morten (2019) in Table 4 of their paper.[20]

The composite amenity-rent term $\left\{\alpha_d r_d^{-\beta}\right\}$ can be disentangled from the unit wage w_d using the model's bilateral wage, $\overline{\text{wage}}_{do} = w_d \bar{h}_{do}$, together with bilateral wage data $\{\overline{\text{wage}}^*_{do}\}$, and the measurement error relationship (27). Simple algebra then yields the following esti-

[19]Estimating equation (29) by Poisson Pseudo Maximum Likelihood to account for the possibility of zero flows (Silva and Tenreyro (2006)) gives rise to very similar estimates of migration costs.

[20]The slope of the relationship between the estimated moving costs and the log of the Euclidean distance between states' population centroids is 0.021 (p-value 0.00), which is also close to the figure 0.023 reported by Bryan and Morten (2019) in their Figure 3.

mating equation:

$$\ln\left(\overline{\text{wage}}^{*}_{do}\left(\pi^{*}_{do}\right)^{\frac{1}{\theta}}\left(\Gamma\left(1-1/\tilde{\theta}\right)\right)^{-1}\right)=\ln w_d+\ln q_o+\ln\left(\frac{v_o}{v^{*}_o}\right)^{\frac{1}{\theta}}+u^{w}_{do}+\frac{u^{f}_{do}}{\theta}. \tag{30}$$

Intuitively, this equation states that observed bilateral wages and/or migration probabilities from o to d are relatively high if unit wages in d are high, or average human capital of agents from o is high, or because of measurement error in wages or migration flows. The regression equation (30) yields estimates of $\{w_d\}$. The estimate of the unit wage allows me to back out the composite amenity-rent terms $\left\{\alpha_d r_d^{-\beta}\right\}$ using the estimated destination fixed effect in equation (29). From the composite term $\alpha_d r_d^{-\beta}$, I back out the amenity parameter α_d using measured rents. This procedure also yields the values of measurement errors $\left\{u^{w}_{do}, u^{f}_{do}\right\}$ and the set of origin parameters $\{q_o\}$, after the normalization $q_1 = 1$. Finally, I back out B_d from equations (19) and (20), after plugging in them measured rents r_d and disposable earnings at destination.

Summary statistics on these parameters are reported in the online appendix. The estimated amenities $\{\alpha_d\}$ are relatively low in the Northeastern states and California and they are highly negatively correlated (-0.91) with the productivity parameters $\left\{\bar{A}_d\right\}$ across states.

4 Tax Progressivity in the Benchmark Model

In this section I use the quantitative model to investigate the impact of tax progressivity on output and compute optimal tax progressivity. Throughout this section I keep moving costs and all other parameters - except for μ - constant at their benchmark value. Section 5 addresses the main questions of the paper by investigating the importance of mobility costs for optimal tax progressivity.

Figure 1 shows the effect of varying tax progressivity on aggregate output and its components. With a linear tax system $(\mu = 0)$, aggregate output Y increases by 10.1%.[21] This increase reflects an increase in labor supply l^* by 6.6% and an increase in output per unit of effort, defined as $\bar{y} = Y/l^*$, by 3.5%. The latter reflects the effect of tax progressivity on the spatial allocation of workers in the economy and is therefore a measure of spatial tax distortions associated with the progressive tax. Endogenizing labor supply makes it possible to compare spatial tax distortions with those related to labor supply within the same envi-

[21] The results are quantitatively very similar if considering the economy's gross domestic product, which consists of Y plus total land rents collected by landlords, instead of Y only.

ronment. Notice that, while labor supply distortions are larger, spatial tax distortions are quantitatively significant.

It is useful to compare spatial tax distortions in this setting with those measured by Eeckhout and Guner (2017) in a related, but distinct, model.[22] In their model, the tax function takes the same form as in equation (4), while moving costs are zero, labor supply is exogenous, and there are no idiosyncratic productivity shocks. They find (see their Table 2) that reducing tax progressivity from their benchmark value $\mu = 0.12$ to $\mu = 0.02$ increases aggregate output by 1.4%. In my model, this shift increases aggregate output per unit of labor effort by 1.9%, a comparable amount.[23]

In this model, lower tax progressivity increases output through two distinct channels. First, it incentivizes agents to relocate to states with relatively high productivity, as shown by Albouy (2009), Eeckhout and Guner (2017), and Colas and Hutchinson (2021). Second, it encourages agents to move to locations where their idiosyncratic productivity is higher, regardless of the state's overall productivity. To assess the impact of these two forces on spatial tax distortions, I decompose the change in output per unit of labor effort in Figure 1 into two components: one driven by the redistribution of population across states and another reflecting the change in average output per unit of effort while holding the population distribution across states fixed at its $\mu = 0.18$ level.[24] I find that more than 60% of the change in output per unit of effort resulting from varying tax progressivity is driven by the "idiosyncratic productivity shocks" channel, while the remaining portion is accounted by changes in the distribution of population across states.

Figure 2 plots the relationship between tax progressivity and the equivalent variation, the welfare measure defined in Section 2.3, for different values of the risk-aversion parameter ϕ. The equivalent variation displays an inverse U-shaped relationship with respect with μ and its maximum increases with risk-aversion. The optimal degree of tax progressivity is approximately equal to the benchmark ($\mu = 0.18$) if risk-aversion is $\phi = 0.2$ and increases to $\mu = 0.33$ when agents' utility is logarithmic. This means that, while in the economy with $\mu = 0.18$, a 1% increase in market income for an agent leads to an 82% increase in consumption, when $\mu = 0.33$, it only leads to a 67% increase in consumption. To put these

[22]While Colas and Hutchinson (2021) and Albouy (2009) discuss spatial tax distortions, they do not report counterfactual results for aggregate output.

[23]Notice that output per unit of labor effort is the analog of output in Eeckhout and Guner (2017) because labor supply is exogenous in their model.

[24]The online appendix shows three additional figures that are useful to interpret the model's mechanisms and are consistent with this decomposition. One shows the impact of moving from the current tax system to a linear tax ($\mu = 0$) on state-level population and output per capita. The second illustrates the effect of a linear tax reform on state-level rents. The third one shows the impact of tax progressivity on gross and excess migration flows.

numbers in perspective, I estimate values of μ equal to 0.32 for Germany and 0.33 for Sweden and Norway using tax data for OECD countries for the years after 2000.

Table 1 provides more details and summary statistics for the linear and optimal tax progressivity reforms. The linear tax reform increases the concentration of workers in the most productive states, leading to an increase in average rents. While output increases, welfare - computed assuming that agents have a unit degree of risk-aversion - falls because a linear tax system does not offer insurance to agents who draw low productivity shocks. The results for a policy reform that increases tax progressivity to $\mu = 0.33$ are almost mirror images of those generated by a linear tax reform, except for welfare. The optimal tax progressivity reform increases welfare by about 1.4%.

5 Tax Progressivity and Moving Costs

In this section I use the quantitative model to address the main questions raised in this paper: to what extent do the effects of tax progressivity on aggregate output depend on migration costs? How do migration costs affect the optimal degree of tax progressivity?

Specifically, I repeat the tax progressivity analysis of Section 4 for different moving costs, focusing on output per unit of effort and welfare as the main variables of interest.[25]

Since there are $J(J-1)$ bilateral moving costs ($J = 48$), to simplify the analysis, I follow Bryan and Morten (2019) and assume that the counterfactual moving costs $\{\tau'_{do}\}$ are obtained from the benchmark ones $\{\tau_{do}\}$ using the follow one-parameter transformation:

$$\tau'_{do} = 1 - (1 - \tau_{do})^{1+\kappa}, \tag{31}$$

with $\kappa \in [-1, 1]$. The parameter κ can be interpreted as the percentage increase (decrease if $\kappa < 0$) in moving costs relative to the benchmark model.[26] When $\kappa = -1$ all moving costs are eliminated ($\tau'_{do} = 0$), while when $\kappa = 1$ moving costs are approximately doubled.

Ouput per unit of labor effort. Figure 3 represents the percentage change of aggregate output per unit of labor effort, $\bar{y}$, when tax progressivity goes from its benchmark $\mu = 0.18$ to $\mu = 0$ (a linear tax) for moving costs in the range $\kappa \in [-1, 1]$. The figure shows that moving costs have a large effect on the response of output to changes in tax progressivity. Switching to a linear tax in the benchmark model ($\kappa = 0$) increases output per unit of

[25] Notice that, since labor supply does not vary with moving costs, the response of aggregate output to tax progressivity depends on moving costs in the same way as the response of output per unit of labor effort does.

[26] To see this, rewrite equation (31) as $\ln(1 - \tau'_{do}) = (1 + \kappa)\ln(1 - \tau_{do})$ and use the approximation $\ln(1 - \tau_{do}) \approx -\tau_{do}$. Therefore, we have $\kappa \approx (\tau'_{do} - \tau_{do})/\tau_{do}$.

labor effort by 3.5%. In an economy with costless mobility ($\kappa = -1$), the same policy yields a 1.8% increase in output per unit of labor effort, almost half the original impact.

Figure 3 also illustrates the non-monotonic impact of moving costs on the output response to tax progressivity: doubling moving costs relative to the benchmark ($\kappa = 1$) is associated with a 1.9% increase in output per unit of effort following a linear tax reform. Interestingly, the current level of moving costs ($\kappa = 0$) is close to the point where the impact of tax progressivity reform is quantitatively largest.

Intuition. To better grasp the intuition behind these results, it is convenient to consider a simplified version of the model with two symmetric locations that produce the same good, abstracting from housing, production externalities, and labor supply. Upon observing the realization of productivity shocks in the home (s_{is}) and away (s_{im}) locations, an agent chooses whether to "stay" or "move", in which case they suffer a moving cost τ. Agent i chooses to move if and only after-tax earnings associated with moving and net of mobility costs are larger than those associated with staying:

$$s_{im}^{1-\mu}(1-\tau) > s_{is}^{1-\mu}.$$

Taking the logarithm of both sides of this equation and rearranging, this condition can be re-written as:

$$\ln s_{im} - \ln s_{is} > \frac{\ln(1-\tau)^{-1}}{1-\mu}, \tag{32}$$

where the left-hand side of this equation represents the proportional gross earning gains from migration and the right-hand side is the migration threshold. The latter increases with mobility costs and tax progressivity.

Panel (a) of Figure 4 plots the population density of ($\ln s_{im} - \ln s_{is}$) together with the migration threshold. The dark blue shaded area shows the mass of agents that choose to migrate when tax progressivity is $\mu = 0.18$. The light blue shaded area corresponds to the mass of agents who choose to move when taxes are linear but don't move when tax progressivity is $\mu = 0.18$.

Panel (b) of Figure 4 considers the same tax reform from $\mu = 0.18$ to $\mu = 0$ for two values of the moving cost parameter τ, one lower and one higher than in panel (a). Both shaded areas in panel (b) represent the mass of agents that choose to move when taxes are linear but don't move when tax progressivity is $\mu = 0.18$. The red area corresponds to the low moving costs case, while the green area corresponds to the high moving costs one. The moving cost parameters are such that these two areas are equal, so that the linear tax reform induces the same mass of agents to migrate under low and high moving costs.

The intuition for this result has to do with how the linear tax reform affects incentives to move and the mass of agents at the margin. The base of each shaded area in panel (b) of Figure 4 represents the impact of the reform on migration incentives at a point in the distribution of gross earnings gains. If moving costs are small, the marginal mover is characterized by smaller net and gross earnings gains than if moving costs are large. Therefore, the linear tax reform has a smaller impact on marginal individuals in the former than in the latter case. However, when moving costs are low, more agents are affected by the tax reform than when moving costs are high - the density is higher at lower moving costs. In panel (b) of Figure 4, these two effects have the same magnitude, so the linear tax reform produces the same impact on the incentives to move - and on output per unit of effort - for low and high moving costs.

Welfare. Figure 5 represents the effect of moving costs on optimal tax progressivity for different coefficients of relative risk aversion ϕ. Notice that optimal tax progressivity displays a U-shaped relationship with respect to moving costs and that this relationship becomes more pronounced for higher degrees of risk aversion. As discussed in the previous section, the U-shape reflects the effect of spatial tax distortions on aggregate output per unit of labor effort. These are relatively small for both high and low values of moving costs and therefore a utilitarian planner chooses to redistribute more in these circumstances.[27]

Interpretation. The upshot of this analysis and discussion is that moving costs are an important determinant of the effectiveness and desirability of tax policies. In the benchmark economy, a linear tax is associated with a 10% increase in aggregate output relative to the current degree of tax progressivity and the reduction in spatial tax distortions accounts for 3.5 percentage points of this increase. In a version of the same economy with no moving costs ($\kappa = -1$), aggregate output would increase by 8.4%, of which 6.6 percentage points are due to higher labor supply and 1.8 to reduced spatial tax distortions. In other words, with zero moving costs the response of aggregate output to a linear tax reform would be 16% smaller than in the benchmark.

An alternative way to interpret this result is by asking: what value of the labor supply elasticity parameter ξ would be required - under benchmark moving costs ($\kappa = 0$) - for a linear tax reform to deliver an 8.4% increase in aggregate output? The answer is $\xi = 4$ (see Table 2, column (5)), which implies a Frisch elasticity of labor supply, given by $1/(\xi - 1)$, equal to $1/3$ - compared to the benchmark value of $1/2$. This represents a substantial

[27] The online appendix reports the equivalent variation associated with the optimal degree of tax progressivity as a function of moving costs. The welfare gain (in consumption equivalent terms) is 1.4% for benchmark moving costs, 1.6% if moving costs are zero, and 1.8% when they are twice as large as in the benchmark. Thus, also this welfare measure inherits the U-shaped pattern that characterizes optimal tax progressivity.

decline. In their review of the literature, Whalen and Reichling (2017) report that "the Frisch elasticity most relevant for fiscal policy analysis range from 0.27 to 0.53," placing the benchmark estimate (1/2) and the alternative estimate (1/3) at nearly opposite ends of the empirical range.

Quantitatively, the optimal degree of tax progressivity - and its sensitivity to moving costs - depends on the coefficient of relative risk aversion. For a unit risk aversion coefficient ($\phi = 1$) and benchmark moving costs ($\kappa = 0$), the optimal degree of tax progressivity μ is 0.33 - comparable to levels observed in Scandinavian countries (see Section 4).

The optimal degree of tax progressivity rises to 0.35 when moving costs are eliminated ($\kappa = -1$), and to 0.37 when moving costs are doubled relative to the U.S. benchmark ($\kappa = 1$). These changes can again be interpreted through the lens of an equivalent shift in the Frisch elasticity of labor supply. Specifically, when the labor supply elasticity parameter is $\xi = 4$ - implying a Frisch elasticity of 1/3 - and all other parameters, including moving costs, remain at their benchmark values, the optimal tax progressivity is approximately 0.36. Thus, eliminating or doubling moving costs produces changes in optimal tax progressivity that are roughly equivalent to those resulting from a decline in the Frisch elasticity from 1/2 to 1/3, holding moving costs constant.

6 Sensitivity Analysis

In this section I briefly comment on the robustness of the results of Sections 4 and 5 to variations in key model parameters.

Table 2 reports the effect of linear (Panel A) and optimal (Panel B) tax progressivity reforms on the main variables of interest under different parameter configurations. The optimal degree of tax progressivity is calculated assuming that agents' coefficient of relative risk aversion equals one ($\phi = 1$). Column (1) considers the case in which housing supply elasticities are doubled in each location (ψ_d is halved in each d). This amplifies the response of output, output per unit of labor effort, and population to a given change in tax progressivity, while attenuating the increase in average rents relative to the benchmark. Since spatial tax distortions are larger in this case, optimal tax progressivity declines relative to the benchmark (see Table 1).

A larger budget share of housing (column 2) dampens population reallocation effects after a tax progressivity reform, as it implies a larger adjustment in housing rents. As a result, optimal tax progressivity increases. A lower elasticity of substitution across inputs in equation (13) (column 3) increases the price response of state-level output to population shocks, limiting reallocation. Nonetheless, the aggregate effects of a linear tax reform remain similar to the benchmark.

Column (4) reports results for the case in which individual shocks are less correlated across locations.[28] This intensifies spatial tax distortions, making tax progressivity more distortionary and amplifying the impact of reducing it on aggregate output. This results in lower optimal tax progressivity. Finally, a lower labor supply elasticity (column 5) reduces the output response, but spatial tax distortions remain largely unchanged, as they depend on labor supply only indirectly.[29] A less elastic labor supply results in a higher optimal degree of tax progressivity.

Figure 6 shows the impact of a linear tax reform on output per unit of labor effort as a function of moving costs for the parameter values considered in Table 2. The non-monotonic relationship between moving costs and the output effects of a linear tax reform is evident for all parameter configurations. Figure 7 plots the optimal degree of tax progressivity as a function of moving costs for the same parameters. Optimal tax progressivity follows a U shape with respect to moving costs in all of these cases.

7 Conclusions

This paper examines how moving costs influence the effectiveness and desirability of tax progressivity using a general equilibrium spatial model. The model highlights that progressive taxation distorts location choices by discouraging workers from relocating to areas where they are most productive, leading to a reduction in aggregate output. Relative to the existing spatial literature, this paper identifies and quantifies a novel source of spatial tax distortions arising from the interaction between location-specific idiosyncratic productivity shocks and mobility costs.

I find that the negative impact of tax progressivity on output is most pronounced when moving costs are either very low or very high. When spatial frictions are high, only workers with substantial pre-tax earnings gains move, so the number of affected individuals is limited. However, because their gains are large, higher tax progressivity significantly discourages their migration. Conversely, when spatial frictions are low, many workers are at the margin between moving and staying, but their individual migration gains are smaller, reducing the distortionary effect of taxation. As a result, spatial tax distortions exhibit a non-monotonic

[28] In this experiment, as I lower the correlation parameter ρ I keep the variance of shocks $\tilde{\theta}$ unchanged. This increases the dispersion of idiosyncratic shocks across locations - so agents are more likely to move - without increasing the overall variance of idiosyncratic shocks.

[29] Although labor supply does not vary across space, it influences output per unit of labor effort by affecting agents' income. In turn, income shapes the demand for housing, which impacts equilibrium rents. The resulting rent adjustments vary across locations due to differences in housing supply elasticity.

pattern, peaking at intermediate levels of moving costs.

From a policy perspective, the paper finds that the optimal degree of tax progressivity follows a U-shaped pattern with respect to moving costs. At both low and high levels of moving costs, optimal tax progressivity is relatively high because spatial tax distortions are smaller. The quantitative model shows that the impact of a linear tax reform on output per unit of effort - an indicator of spatial tax distortions - is roughly twice as large in the benchmark economy as in economies with either zero or double the benchmark moving costs. In both of these alternative cases, the optimal degree of tax progressivity is similar to that obtained in an economy with benchmark moving costs and a Frisch labor supply elasticity that is one-third lower.

Beyond the specific quantitative findings, this paper raises broader questions about the interaction between tax policy and spatial frictions. One possible avenue for future research is to examine how different types of moving costs - such as financial constraints, informational frictions, and social networks - affect migration decisions and tax policy outcomes. Additionally, while this paper focuses on national tax progressivity, future work could explore how regional and local tax policies interact with mobility frictions and shape the spatial distribution of economic activity.

References

Albouy, David. 2009. "The Unequal Geographic Burden of Federal Taxation." *Journal Political Economy*, 117(4): 635–667.

Armington, Paul. 1969. "A Theory of Demand for Products Distinguished by Place of Production." *IMF Staff Papers*, 16(1): 159–178.

Baum-Snow, Nathaniel, and Lu Han. 2023. "The Microgeography of Housing Supply." *Journal of Political Economy.*

Benabou, Roland. 2002. "Tax and Education Policy in a Heterogeneous-Agent Economy: What Levels of Redistribution Maximize Growth and Efficiency?" *Econometrica*, 70(2): 481–517.

Bryan, Gautam, and Melanie Morten. 2019. "The Aggregate Productivity Effects of Internal Migration: Evidence from Indonesia." *Journal of Political Economy*, 127(5): 2229–2268.

Bryan, Gharad, Shyamal Chowdhury, and Ahmed Mushfiq Mobarak. 2014. "Underinvestment in a Profitable Technology: The Case of Seasonal Migration in Bangladesh." *Econometrica*, 82: 1671–1748.

Coen-Pirani, Daniele. 2021. "Geographic Mobility and Redistribution." *International Economic Review*, 62(3): 921–952.

Colas, Mark, and Kevin Hutchinson. 2021. "Heterogeneous Workers and Federal Income Taxes in a Spatial Equilibrium." *American Economic Journal: Economic Policy*, 13(2): 100–134.

Davis, Morris A., and Francois Ortalo-Magne'. 2011. "Household expenditures,wages,rents." *Review of Economic Dynamics*, 14: 248–261.

Desmet, Klaus, and Esteban Rossi-Hansberg. 2013. "Urban Accounting and Welfare." *American Economic Review*, 103(6): 2296–2327.

Diamond, Peter, and Emmanuel Saez. 2011. "The Case for a Progressive Tax: From Basic Research to Policy Recommendation." *Journal of Economic Perspectives*, 25(4): 165–190.

Eeckhout, Jan, and Nezih Guner. 2017. "Optimal Spatial Taxation: Are Big Cities Too Small?" *CEMFI Working Paper 1705.*

Fajgelbaum, Pablo, and Cecile Gaubert. 2020. "Optimal Spatial Policies, Geography, and Sorting." *The Quarterly Journal of Economics*, 125: 959–1036.

Fajgelbaum, Pablo, Eduardo Morales, Juan Carlos Suarez Serrato, and Owen Zidar. 2019. "State Taxes and Spatial Misallocation." *Review of Economic Studies*, 86: 333–376.

Feldstein, Martin S. 1969. "The Effects of Taxation on Risk Taking." *Journal of Political Economy*, 77(5): 755–764.

Giannone, Elisa, Nuno Paixao, Xinle Pang, and Qi Li. 2023. "Unpacking Moving: A Spatial Equilibrium Model with Wealth." *CEPR Discussion Paper No. 18072.*

Guvenen, Fatih, Burhanettin Kuruscu, and Serdar Ozkan. 2014. "Taxation of Human Capital and Wage Inequality: A Cross-Country Analysis." *Review of Economic Studies*, 81: 818–850.

Hassler, John, Jose V. Rodriguez-Mora, Kjetil Storesletten, and Fabrizio Zilibotti. 2005. "A Positive Theory of Geographic Mobility and Social Insurance." *International Economic Review*, 46: 263–303.

Heathcote, Jonathan, Kjetil Storesletten, and Gianluca Violante. 2017. "Optimal Tax Progressivity: An Analytical Framework." *Quarterly Journal of Economics*, 132: 1693–1754.

Heise, Sebastian, and Tommaso Porzio. 2022. "Labor Misallocation Across Firms and Regions." *NBER working paper 30298.*

Hsieh, Chang-Tai, and Enrico Moretti. 2019. "Housing Constraints and Spatial Misallocation." *American Economic Journal: Macroeconomics*, 11(2): 1–39.

Keane, Michael P. 2011. "Labor Supply and Taxes: A Survey." *Journal of Economic Literature*, 49(4): 961–1075.

Kennan, John, and Robert Walker. 2011. "The Effect of Expected Income on Individual Migration Decisions." *Econometrica*, 79: 211–251.

Krueger, Dirk, and Alexander Ludwig. 2013. "Optimal Progressive Labor Income Taxation and Education Subsidies When Education Decisions and Intergenerational Transfers are Endogenous." *American Economic Review*, 103(3): 496–501.

Lagakos, David, Ahmed Mushfiq Mobarak, and Michael E. Waugh. 2023. "The Welfare Effects of Encouraging Rural-Urban Migration." *Econometrica*, 91(3): 803–837.

Persson, Mats. 1983. "The distribution of abilities and the progressive income tax." *Journal of Public Economics*, 22(1): 73–88.

Roback, Jennifer. 1982. "Wages, Rents, and the Quality of Life." *Journal of Political Economy*, 90(6): 1257–1278.

Ruggles, Steven, Sarah Flood, Sophia Foster, Ronald Goeken, Jose Pacas, Megan Schouweiler, and Matthew Sobek. 2021. "IPUMS USA: Version 11.0 [dataset]. Minneapolis, MN: IPUMS." *https://doi.org/10.18128/D010.V11.0.*

Saiz, Albert. 2010. "The Geographic Determinants of Housing Supply." *Quarterly Journal of Economics*, 125(3): 1253–96.

Silva, J.M.C. Santos, and Silvana Tenreyro. 2006. "The Log of Gravity." *The Review of Economics and Statistics*, 88: 641–658.

Whalen, Charles, and Felix Reichling. 2017. "Estimates of the Frisch Elasticity of Labor Supply: A Review." *Eastern Economic Journal*, 43(1): 37–42.

Zabek, Mike. 2024. "Local Ties in Spatial Equilibrium." *American Economic Journal: Macroeconomics*, 16(2): 287–317.

Zerecero, Miguel. 2021. "The Birthplace Premium." *Working Paper.*

	Linear tax	Optimal tax progressivity
	(1)	(2)
	$\mu = 0$	$\mu = 0.33$
Aggregate output	10.1	-10.3
Labor supply	6.6	-6.8
Output per unit of labor effort	3.5	-3.5
Average rental price of housing	8.0	-7.7
Aggregate land rents of landlords	11.5	-11.7
Population in 5 most productive states	10.9	-9.1
Equivalent variation ($\phi = 1$)	-4.7	1.4

Table 1: Effect of tax progressivity reform. Column (1) corresponds to a linear tax reform from $\mu = 0.18$ to $\mu = 0$. Column (2) represents the results of the optimal tax progressivity reform from $\mu = 0.18$ to $\mu = 0.33$ (assuming unit risk aversion, $\phi = 1$). Each entry, except for the equivalent variation, represents the percentage change in the row variable in the counterfactual economy relative to the benchmark one. The equivalent variation is a percentage amount.

Panel A	Linear tax reform				
	(1)	(2)	(3)	(4)	(5)
	$\{\psi_d/2\}$	$\beta = 0.4$	$\sigma = 4$	$\rho = 0.86$, $\theta = 20$	$\xi = 4$
Aggregate output	10.4	9.8	10.1	10.8	8.4
Labor supply	6.6	6.6	6.6	6.6	5.0
Output per unit of labor effort	3.8	3.2	3.4	4.2	3.4
Average rental price of housing	6.8	7.5	8.0	8.4	6.9
Aggregate land rents of landlords	12.0	11.1	11.4	12.2	9.6
Population in 5 most productive states	12.5	8.3	7.9	10.2	11.0
Equivalent variation ($\phi = 1$)	−4.2	−5.4	−4.8	−4.8	−4.8

Panel B	Optimal tax progressivity				
	(1)	(2)	(3)	(4)	(5)
	$\{\psi_d/2\}$	$\beta = 0.4$	$\sigma = 4$	$\rho = 0.86$, $\theta = 20$	$\xi = 4$
Optimal tax progressivity	0.30	0.37	0.33	0.32	0.36
Aggregate output	−8.0	−13.05	−10.3	−10.5	−10.5
Labor supply	−5.1	−9.0	−6.8	−6.5	−6.3
Output per unit of labor effort	−2.9	−4.1	−3.5	−4.0	−4.1
Average rental price of housing	−4.9	−9.6	−7.9	−7.8	−8.1
Aggregate land rents of landlords	−9.3	−14.8	−11.7	−11.9	−12.0
Population in 5 most productive states	−8.2	−8.3	−7.4	−7.7	−11.2
Equivalent variation ($\phi = 1$)	0.9	2.3	1.4	1.3	1.8

Table 2: Effect of a linear tax progressivity reform (Panel A) and of an optimal tax progressivity reform (Panel B). Each column represents a different parameter configuration relative to the benchmark calibration of the model. Each row entry, except for the equivalent variation, represents the percentage change in the relevant variable in the counterfactual economy relative to the benchmark one. The equivalent variation is a percentage.

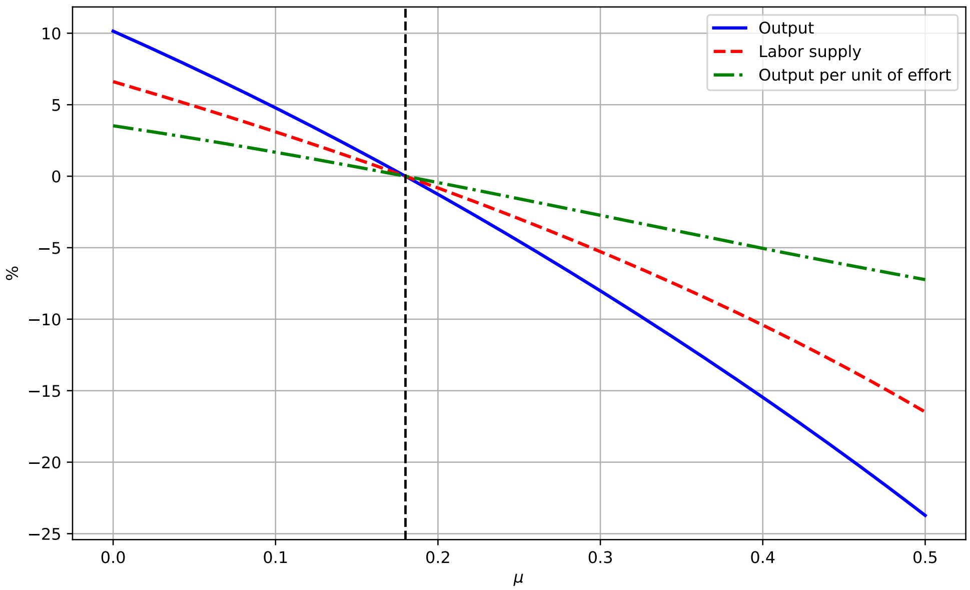

Figure 1: The percent impact of varying tax progressivity on aggregate output (Y), labor supply (l^*) and output per unit of labor effort ($\bar{y} = Y/l^*$) relative to the benchmark degree of tax progressivity, $\mu = 0.18$ (denoted by the vertical black dash line). All parameters other than μ are at their benchmark values.

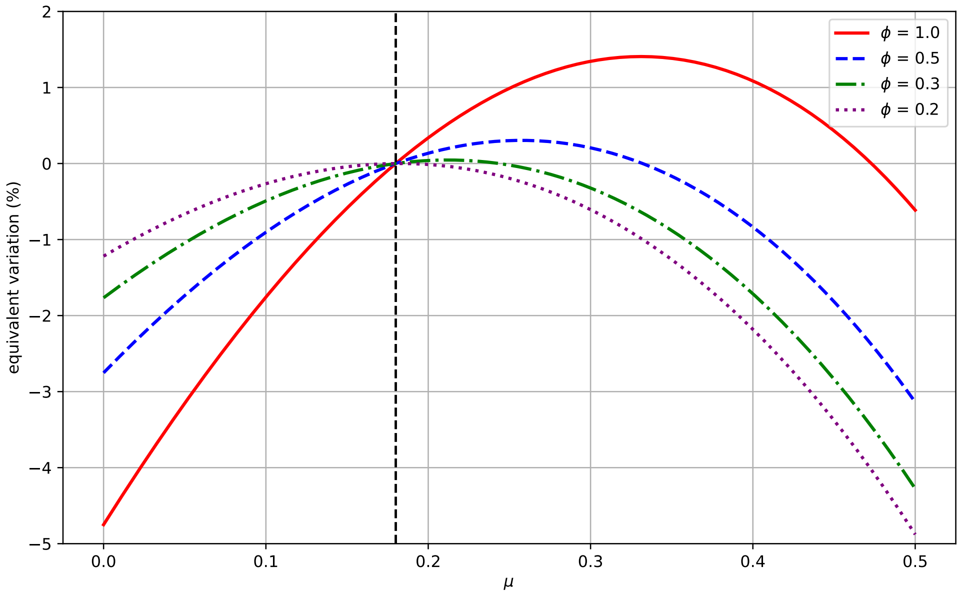

Figure 2: The effect of tax progressivity on welfare. The figure plots the relationship between the equivalent variation, defined in Section 2.3, and tax progressivity μ for different values of the risk-aversion parameter ϕ. The vertical black dashed line corresponds to the benchmark degree of tax progressivity $\mu = 0.18$. All parameters other than μ and ϕ are at their benchmark values.

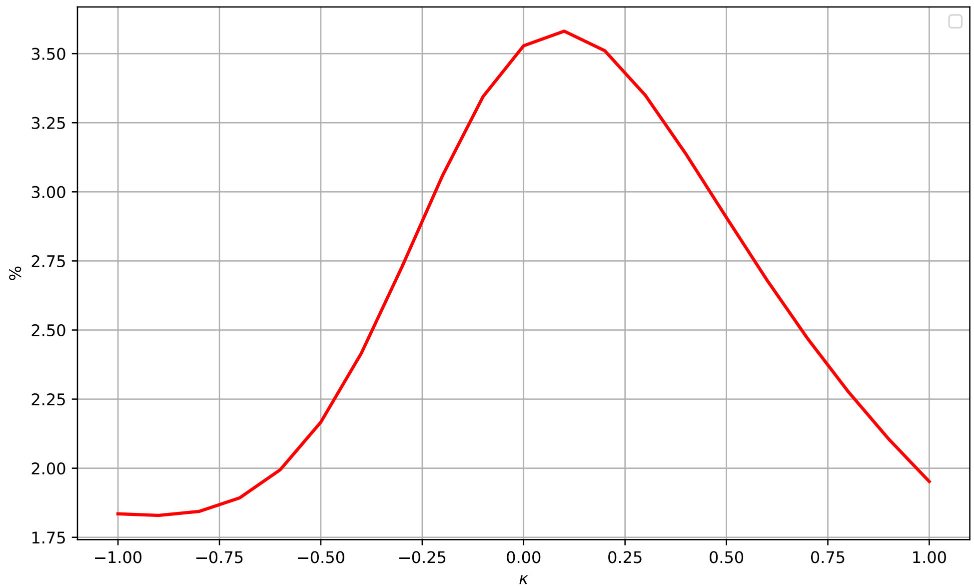

Figure 3: This figure plots the percent gain in aggregate output per unit of labor effort ($\bar{y} = Y/l^*$) associated with moving from the benchmark tax system ($\mu = 0.18$) to a linear tax system ($\mu = 0$). An increase in κ over 0 corresponds to an increase in moving cost, while a decline relative to 0 corresponds to smaller moving costs. If $\kappa = -1$, all moving costs are zero ($\tau_{do} = 0$ for all (d, o) combinations).

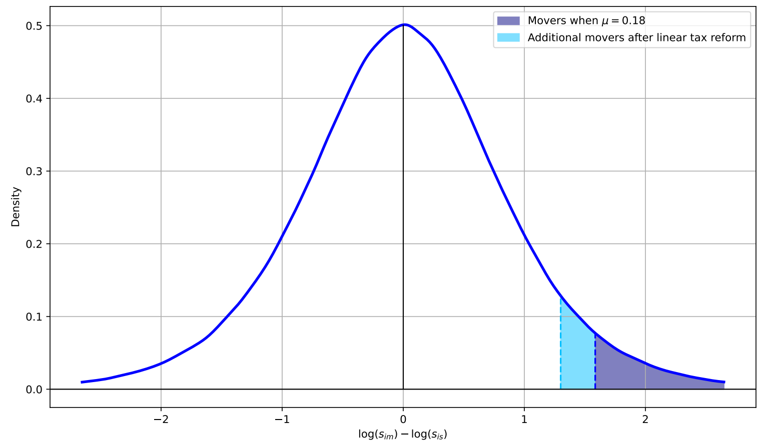

(a) The dark blue shaded area represents the mass of workers who chooses to move when tax progressivity is $\mu = 0.18$. The light blue shaded area represent the additional mass of workers of do not move when $\mu = 0.18$ but choose to do so after a linear tax reform. The blue density represents the probability distribution of $(\ln s_{im} - \ln s_{is})$. Parameters: $\theta = 2$, $\rho = 0, \tau = 0.73$.

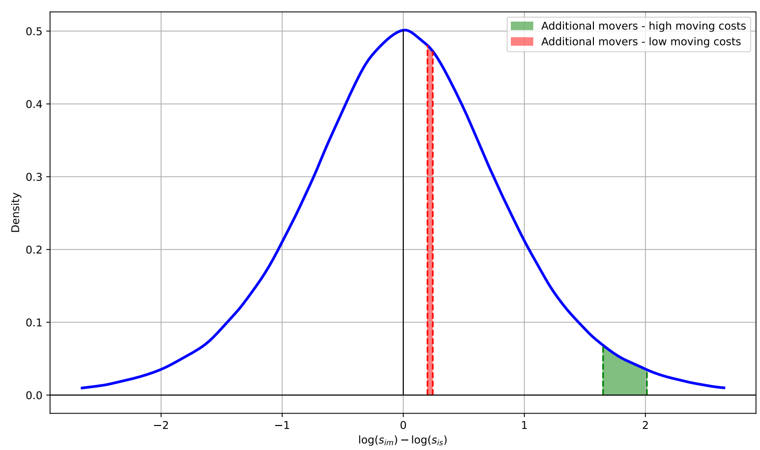

(b) Impact of a linear tax reform on the mass of movers at low (τ_L) and high (τ_H) moving costs. The red (green) shaded area represents the mass of agents that choose to move after the linear tax reform in the low (high) moving cost case. The blue density represents the probability distribution of $(\ln s_{im} - \ln s_{is})$. The green and red areas are approximately equal by construction. Parameters: $\theta = 2$, $\rho = 0, \tau_L = 0.18, \tau_H = 0.81$.

Figure 4: Intuition for the effect of a linear tax reform on migration rates at different moving costs.

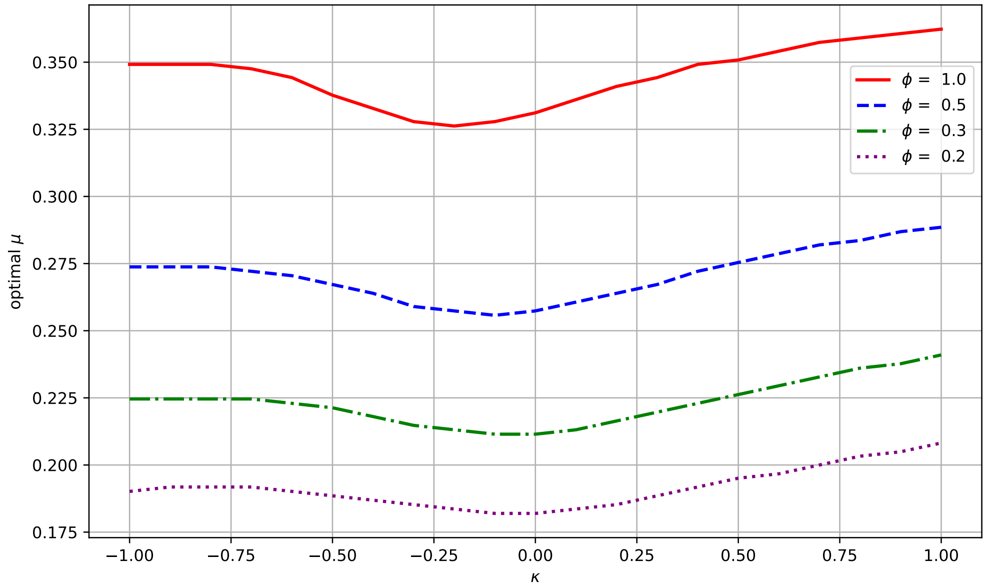

Figure 5: Optimal tax progressivity and moving costs. Each line corresponds to a different coefficient of relative risk aversion ϕ. Benchmark moving costs correspond to $\kappa = 0$. An increase in κ over 0 corresponds to an increase in moving cost, while a decline relative to 0 corresponds to smaller moving costs. If $\kappa = -1$, all moving costs are zero ($\tau_{do} = 0$ for all (d, o) combinations).

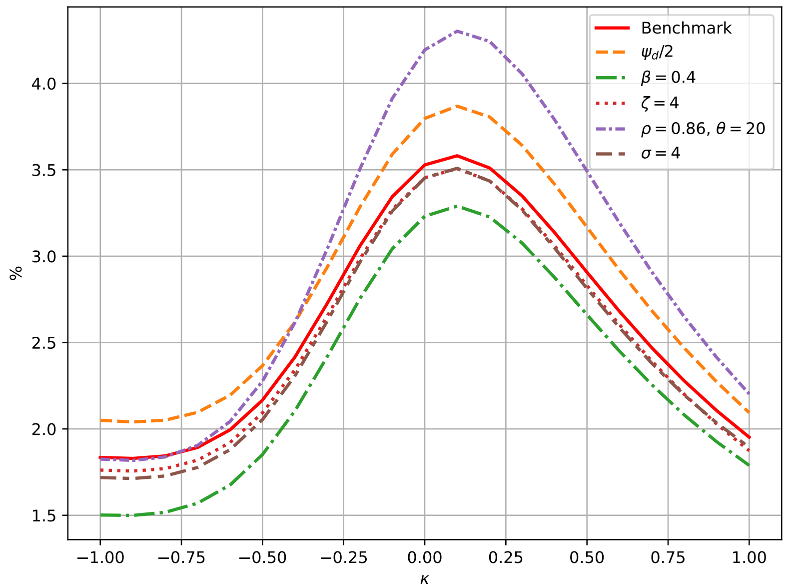

Figure 6: This figure plots the percent gain in aggregate output per unit of labor effort ($\bar{y}$) associated with moving from the current tax system (tax progressivity $\mu = 0.18$) to a linear tax system ($\mu = 0$) for each of the parameter configurations in Table 2. Benchmark moving costs correspond to $\kappa = 0$. An increase in κ over 0 corresponds to an increase in moving cost, while a decline relative to 0 corresponds to smaller moving costs. If $\kappa = -1$, all moving costs are zero ($\tau_{do} = 0$ for all (d, o) combinations).

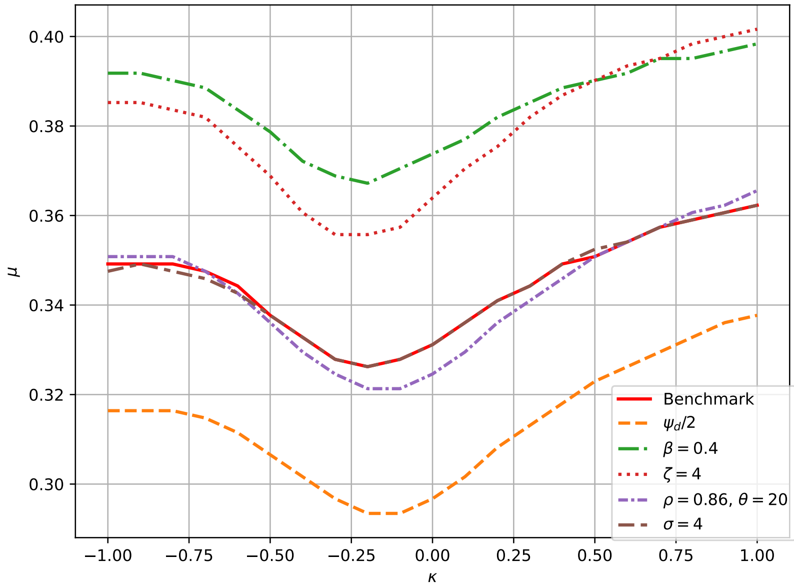

Figure 7: Optimal tax progressivity and moving costs. Each line corresponds to one of the parameter configurations in Table 2. Benchmark moving costs correspond to $\kappa = 0$. An increase in κ over 0 corresponds to an increase in moving cost, while a decline relative to 0 corresponds to smaller moving costs. If $\kappa = -1$, all moving costs are zero ($\tau_{do} = 0$ for all (d, o) combinations). The degree of risk aversion underlying all plots is $\phi = 1$.