Redistributing the Gains From Trade through Progressive Taxation

PRELIMINARY AND WORK IN PROGRESS.

Spencer G. Lyon
New York University

Michael E. Waugh
New York University and NBER

September 2017

ABSTRACT

This paper studies the optimal degree of tax progressivity as an economy opens to trade. We answer this question within a standard incomplete markets model with frictional labor markets and Ricardian trade. Labor market frictions lead to losses in labor income and reductions in labor force participation for import-competition-exposed workers. In an incomplete markets setting, adverse shocks to comparative advantage are imperfectly insured and import-competition-exposed workers experience welfare losses. A progressive tax system transfers resources to the losers from trade and substitutes for imperfect insurance. However, a progressive tax system reduces incentives to work and for labor to reallocate from comparatively disadvantaged locations to locations with a comparative advantage. We calibrate the model to match the observed labor market response to trade exposure and solve for the optimal degree of progressivity of the tax system with varying levels of exposure to trade.

Email: spencer.lyon@stern.nyu.edu, mwaugh@stern.nyu.edu.

There are many concerns about the forces of globalization—that the losses from trade are large; that there are insufficient mechanisms to insure against these losses; that globalization simply propagates existing inequality. The standard answer to these concerns is helpful in theory: that there exists a Pareto improving transfer scheme that can compensate the losers from trade, yet still preserve the gains for the winners. In practice, this answer is less helpful given the limited mechanisms and incentive problems policy makers face in implementing any kind of transfer scheme.

Evidence suggests that that these concerns are warranted. Autor, Dorn, and Hanson (2013) (henceforth ADH) show that exposure to Chinese import competition led to losses in labor income and reductions in labor force participation for import-competition-exposed workers. Krishna and Senses (2014)) show that import penetration has had a significant association with labor income risk in the United States. Moreover, given that risk sharing is often found to be incomplete (see, e.g., Cochrane (1991), Attanasio and Davis (1996)), this suggests that the labor market consequences of trade led to welfare losses.

Building on the insights of Varian (1980) and Eaton and Rosen (1980), one way to mitigate and insure against these losses is via the tax system. That is, use a progressive tax system to provide social insurance and transfer resources from the winners from trade to the losers. This paper evaluates this possibility by measuring the optimal degree of tax progressivity as an economy becomes exposed to trade.

We implement these ideas by building off of our parallel work in Lyon and Waugh (2017). In this work, we develop an open-economy, standard incomplete markets model in which households face uninsurable income shocks (some of which are trade related), and yet have several margins to mitigate labor income risk. Specifically, households can self-insure, opt-out of the labor force, and/or move when local labor market conditions are not favorable. These mechanisms lead to a tension on the optimal degree of tax progressivity. On the one hand, a progressive tax system provides a mechanism to substitute for imperfect insurance against labor income risk and to transfer resources to the losers from trade. However, a progressive tax system reduces incentives to work and for labor to migrate.

In our model, the government uses a parametric, log-linear labor-income tax and transfer scheme to redistribute resources. In particular, we closely follow approach of Benabou (2002), Conesa and Krueger (2006), and Heathcote, Storesletten, and Violante (2014) where the progressivity of net-taxes is parameterized directly.[1] As in this work, we take a stand on the social welfare function. Specifically, we focus on a utilitarian planner which places equal weight on households within the domestic economy. The optimal degree of progressivity is the measured

[1] Heathcote, Storesletten, and Violante (2014) show that this functional form provides a good approximation of the actually tax and transfer scheme in the US data. Guner, Kaygusuz, and Ventura (2014) provide an exploration of this and alternative tax functions.

as the progressivity parameter that maximizes social welfare.

The optimal degree of progressivity and how it varies with trade exposure is ultimately a quantitative question. We essentially take our parameter estimates from Lyon and Waugh (2017). In the model, there are three key elasticities: how wages respond to trade shocks, how labor supply responds to trade-induced changes in wages, and how migration responds to changes in wages. The former elasticity essentially determines the amount of trade-induced income risk there is the model; the latter two elasticities speak to the ability of households to adjust in response to trade-induced income shocks. These three elasticities map into parameters regarding the elasticity of substitution across goods, the disutility of working, and migration costs. We then discipline these three parameters by asking the model to match the ADH facts regarding (i) decline in labor income in import-competition-exposed areas (ii) decline in labor force participation and (iii) migration facts in ADH and updated data.

Given the calibrated model, we ask several policy questions of the model. First, what is the optimal policy and what trade-offs does the policy maker face? We find that the US tax system should be more progressive than it currently is, but that the gains from any change are small. Unlike the closed economy studies of Conesa and Krueger (2006) and Heathcote, Storesletten, and Violante (2014), we show that the quantitatively important tradeoff the policy maker faces is not about labor supply, but the reduction in migration (and hence allocative efficiency) versus providing social insurance. The migration margin is new and unexplored—providing social insurance to households reduces their incentive to move from low productivity to high productivity places, thus, reducing allocative efficiency.

Second, how does optimal policy change with increased openness to trade? We find that the tax system should generally become more progressive as the economy becomes more open. However, only at relatively large levels of openness (at least for the US, but comparable to economies such as Canada and Mexico) do we see quantitatively large, prescribed changes in the optimal policy.

Finally, we ask how does openness change the costs from moving away from optimal (or current) policy? While related to the previous question, this question is about concavity of the "Laffer-like" curve in social welfare and less about how optimal policy changed. We fine that the costs associated with moving a way from the optimal or current policy become larger as we become more open. For example, a move to a flat tax system when the economy is very open would lead to a two and half percent decrease in welfare where as at the baseline the move to a flat tax system would lead to a half a percent decrease in welfare.

Related Literature. This paper is work in progress and we are currently surveying the literature to understand our place in it. This paper is closely related to the quantitative studies of Conesa and Krueger (2006) and Heathcote, Storesletten, and Violante (2014) who study the optimal

progressivity of the US tax scheme in heterogenous agent, incomplete market models. This paper is distinguished by its focus on trade and open economy issues. Moreover, our question is less on what the optimal system is vis-a-vi the current system but about how it should change in response to a specific counterfactual.

With regards to open economy issues, the most closely related is Antràs, De Gortari, and Itskhoki (2016). In a static, open economy Mirrlees (1971) framework, they study the welfare consequences of trade-induced inequality its interaction with redistributive polices. Our work complements their study by exploring an economy with alternative motives for redistributive taxation—that is motives arising from risk and incomplete insurance. In our model, the key motive for progressive taxation is to redistribute resources towards the "unlucky" from the "lucky." While related, this insurance motive is distinct from inequality aversion per se, as in Mirrlees (1971). This is the sense in which our paper builds most closely on the earlier work from Varian (1980), Eaton and Rosen (1980), and Mirrlees (1974).

Our modeling framework is related, but distinct from an exciting and growing body work on trade and labor market dynamics (see, e.g., Kambourov (2009), Artuç, Chaudhuri, and McLaren (2010), Dix-Carneiro (2014), Caliendo, Dvorkin, and Parro (2015)). Our key departure is by studying an economy in which households face income shocks and incomplete markets, but with ways to self insure. We view this as an important departure for two reasons. As discussed above, this opens up the door to a motive for government policy to provide social insurance and increased social insurance as the economy opens to trade. Second, in our setting, the motive for households to migrate is for insurance.[2] This motive is important as it creates a tension between losses in allocative efficiency that come with less migration versus the gains from social insurance.

1. Model

Here we describe a model of international trade with households facing incomplete markets and frictions to move across labor markets. The first section discusses the production structure, the second section discusses the government and the tax function, the third section discusses the households, and finally an equilibrium is discussed and defined.

1.1. Production

The model has an intermediate goods sector and a final good sector that aggregates the intermediate goods. Within a country, there is a continuum of intermediate goods indexed by $\omega \in [0, 1]$. As in the Ricardian model of Dornbusch, Fischer, and Samuelson (1977) and Eaton

[2]Insurance motivated moves are distinct from migration in Artuç, Chaudhuri, and McLaren (2010) and Caliendo, Dvorkin, and Parro (2015) which arise from shocks to preferences across locations/sectors.

and Kortum (2002), intermediate goods are not nationally differentiated and, thus, intermediate ω produced in one country is a perfect substitute for the same intermediate produced by another country.

Intermediate goods are produced by competitive firms with linear production technologies,

$$q(\omega) = z(\omega)\ell, \tag{1}$$

where z is the productivity level of firms and ℓ is the number of efficiency units of labor. Intermediate goods productivity z evolves stochastically according to an AR(1) process in logs

$$\log z_{t+1} = \phi \log z_t + \epsilon_t \tag{2}$$

where ϵ_t is distributed normally with mean zero and standard deviation σ_ϵ. The innovation ϵ_t is independent across time, across goods, and across countries.

Firms producing variety ω will face competitive product and labor markets with households that can potentially supply labor. Competition implies that a household choosing to work in market ω earn the value of their marginal product of labor, which is the price of the good times the firm's productivity z.

Finally, transporting goods across countries is costly. Specifically, variety producing firms face iceberg trade costs $\tau \geq 1$ when exporting their product. This means that for a firm to deliver one unit of the intermediate good abroad, it must produce $\tau > 1$ units for shipment.

Intermediate goods are aggregated by a competitive final goods producer. This final goods producer has a standard CES production function:

$$Q = \left[\int_0^1 q(\omega)^\rho d\omega \right]^{\frac{1}{\rho}}, \tag{3}$$

where $q(\omega)$ is the quantity of individual intermediate goods ω demanded by the final goods firm and ρ controls the elasticity of substitution across variety, which is $\theta = \frac{1}{1-\rho}$.

1.2. Government

The government both consumes resources G and levies taxes and transfers to finance these expenditures. Define net tax revenues at labor income level w as $T(w)$. Following Benabou (2002) and Heathcote, Storesletten, and Violante (2014), we assume that net tax revenues is of the following parametric class:

$$T(w) = w - \delta w^{1-\tau_p} \tag{4}$$

There are two parameters in (4). The δ parameter determines the average rate. The parameter τ_p directly controls the progressivity of the tax scheme. Heathcote, Storesletten, and Violante (2014) describe several ways to see how τ_p determines the progressivity. The most straightforward way is to note that $1-\tau_p$ is equal to one minus the marginal tax rate relative to one minus the average tax rate. Thus, if when the progressivity parameter is greater than zero, marginal rates exceed average rates and the tax system is progressive. If the progressivity parameter equals zero, than the marginal rates equal average rates and the tax system is a flat tax with rate $1-\delta$. Heathcote, Storesletten, and Violante (2014) and Guner, Kaygusuz, and Ventura (2014) show that this functional form provides a good approximation of the actual tax and transfer scheme in the US data.

1.3. Households

In country i there is a continuum of infinitesimally small households of mass L_i. Each household is infinitely lived and seeks to maximize expected discounted utility

$$E\sum_{t=0}^{\infty}\beta^t\left\{\log(c_t)-B\frac{h_t^{1-\gamma}}{1-\gamma}\right\} \tag{5}$$

where E is the expectation operator and β is the subjective discount factor. Period utility depends on both consumption of the final good and the disutility of labor. As discussed below, we model labor supply as being purely on the extensive margin, thus the parameter γ is irrelevant and only B matters.

Households live and work along the same dimension as the intermediate goods. That is a household's location is given by ω, the intermediate goods sector it can work in. Given their current location, households can chose if to work, to move and work someplace else in the future, and accumulate a non-state contingent asset. Below, we describe each of these choices in detail.

Working is a discrete choice between zero hours and $\bar{h}$. Thus, the labor supply is purely on the extensive margin. If a household works, it receives income from employment in the intermediate goods sector that the household resides. If a household does not work, if receives some (un-taxed) home production w_n.

The wage per efficiency unit that a household receives depends on the value of the marginal product of labor in that market. This depends directly on a country's productivity level in that market and indirectly through all countries productivity levels in the production of that same good through trade. Thus, the relevant state variable determining the wage per efficiency unit are all countries' productivities. The wage per efficiency unit a household earns is $w(\mathbf{s})$ where

$\mathbf{s}$ denotes the vector of productivity levels across countries.[3] Given the tax function in 4, the posttax earnings of the household is

$$\tilde{w}(\mathbf{s}) = \delta \left(w(\mathbf{s}) \bar{h} \right)^{1-\tau_p} . \tag{6}$$

Households can move to an alternative intermediate goods sector ω' at some cost. Paying $m > 0$ in units of the final good to allows the household to change where they can work in next period. If a household chooses to move in a given period, they remain in their current location in that period and can choose supply labor in the original intermediate goods sector. They begin working in the sector ω' starting in the period after they move. The value of the new location can take several forms. One is the best labor market as in Lucas and Prescott (1974); an alternative is a random labor market. We explore both specifications.

Per the discussion above, households residing in a intermediate goods location face idiosyncratic labor income risk associated with fluctuations in productivity and (as discussed below) fluctuations in world prices. We do not allow for any insurance markets against this risk, but let households accumulate a non-state contingent asset a that pays gross return R. For now, we treat R as exogenous and not solved for in equilibrium. An interpretation is that in the economy is a large supply of assets at this rate. Households face a lower bound on asset holding $-\phi$, so agents can acquire debt up to the value ϕ.

Given the description of the environment, the household's period t budget constraint is

$$a_{t+1} + P_h c_t + \iota_{m,t} m_t \leq R a_t + \iota_{n,t} \tilde{w}(\mathbf{s}) + (1 - \iota_{n,t}) w_n \tag{7}$$

where R is the gross interest rate on assets, c_t is consumption, m_t is the moving cost a household faces multiplied by ι_t which is an indicator function equalling one if a household moves in period t and zero otherwise. Finally, $\iota_{n,t}$ is an indicator function equalling one if a household works.

The individual state variable of a household is its asset holdings a. The sector level state variable relevant to the household is the wage $w(\mathbf{s})$ where the household resides. This depends on domestic productivity and wold prices which are summarized in the vector $\mathbf{s}$. Thus, we denote the state variable of a household as $(a, \mathbf{s})$.

Given these states, the recursive formulation of the household's problem is a discrete choice between four options: the value of staying and working, the value of staying and not working, the value of moving and working, the value of moving and not working. The value function of

[3] While standard, it is worth discussing how the location ω is not a relevant state variable—because intermediate goods enter the final goods production function symmetrically, outcomes in an intermediate goods only depend upon productivity; not any ω specific variable.

the household is the max over these four options. The value of staying and working is

$$V^{s,w}(a,\ \mathbf{s}) = \max_{a' \geq -\phi} \left[u(Ra + \tilde{w}(\mathbf{s}) - a') - \bar{h}B + \beta EV(a', \mathbf{s'}) \right]. \tag{8}$$

where $u(\bullet) = \frac{c^{1-\sigma}}{1-\sigma}$. The value of staying and not working is

$$V^{s,nw}(a,\ \mathbf{s}) = \max_{a' \geq -\phi} \left[u(Ra + w_n - a') + \beta EV(a', \mathbf{s'}) \right] \tag{9}$$

The value of moving and working is

$$V^{m,w}(a,\ \mathbf{s}) = \max_{a' \geq -\phi} \left[u(Ra + \tilde{w}(\mathbf{s}) - a' - m) - \bar{h}B + \beta V^{m}(a') \right]. \tag{10}$$

where there are two key distinctions relative to (8). First, the moving cost, m is paid. Second, the continuation value is $V^m(a')$. As discussed above, the value of moving can take several forms: the best labor market, a random labor market, etc. Finally, the value of moving and not working is

$$V^{m,nw}(a,\ \mathbf{s}) = \max_{a' \geq -\phi} \left[u(Ra + w_n - a' - m) + \beta EV^{m}(a') \right] \tag{11}$$

Putting these together describes the value function of a household with asset level a and labor market states of $\mathbf{s}$

$$V(a,\ \mathbf{s}) = \max \left[V^{s,w},\ V^{s,nw},\ V^{m,w},\ V^{m,nw} \right]. \tag{12}$$

2. Equilibrium

In this section, we close the model by focusing on a small open economy equilibrium. Specifically, we solve for a home country equilibrium given world prices. The small open economy assumption is that there is no feedback from home country actions into world prices. Relative to the literature, this is similar to the second specification solved in Artuç, Chaudhuri, and McLaren (2010). Moreover, it has the advantage (say relative to Caliendo, Dvorkin, and Parro (2015)) to be relatively simple, yet allows us to specific about the interaction between trade flows and capital flows. And, thus, we are able to talk meaningfully about the role of trade imbalances.

World Prices. World prices for commodity ω evolve according to an independent AR(1) process in logs as well:

$$\log p_w(\omega)_{t+1} = \phi \log p_w(\omega) + \epsilon_{w,t} \tag{13}$$

where $\epsilon_{w,t}$ is distributed normally with mean zero and standard deviation σ_w and independent of the innovation to the home countries productivity ϵ_t. We express these prices in units of the numeraire, which we take to be the final good in the home country.

As we discuss below, domestic prices will depend both on domestic productivity and the world price. Thus, we carry around the the vector $\mathbf{s} = \{ z_h, p_w \}$ and the relevant state variable determining the price $p(\mathbf{s})$.

Finally, a note on notation. We denote $\pi(\mathbf{s})$ as the stationary distribution of productivity states and world prices induced by (2) and (13). And denote $\pi(\mathbf{s'}, \mathbf{s})$ as the probability of transiting from state $\mathbf{s}$ to $\mathbf{s'}$.

2.1. Production Side of the Economy

Below we describe the equilibrium conditions associated with the production side of the economy. These take as given the choices of the household. In particular, define $\mu_h(\mathbf{s})$ as the mass of households supplying labor in market of type $\mathbf{s}$.

Final Goods Production. Final goods producer's problem is:

$$\max_{q(\mathbf{s})} P_h Q_h - \int_0^1 p_h(\mathbf{s}) q(\mathbf{s}) \pi(\mathbf{s}) dz \tag{14}$$

which gives rise to the following the demand curve for an individual variety

$$q_h(\mathbf{s}) = \left(\frac{p_h(\mathbf{s})}{P_h} \right)^{-\theta} Q_h. \tag{15}$$

where Q_h is aggregate demand of the final good; P_h is the price associated with the final good.

Intermediate Goods Production. The intermediate goods producers problem is

$$\max_{q(\mathbf{s})} p_h(\mathbf{s}) q(\mathbf{s}) - P_h w(\mathbf{s}) \frac{q(\mathbf{s})}{z_h}, \tag{16}$$

or to chose the quantity produced to maximize profits. Competition implies that the wage per efficiency unit (in units of the final good) the firm hires labor at is:

$$w_h(\mathbf{s}) = \frac{p_h(\mathbf{s}) z_h}{P_h}, \tag{17}$$

or the value of the marginal product of labor. Only at the wage in (17) are intermediate goods producers willing to produce.

Intermediate Goods, International Trade, and Market Clearing. To formulate the pattern of

trade, we denote the set of prevailing prices that the final goods producer in the home country faces as $p_h(\mathbf{s})$, τp_w. The final goods producer in each country purchases intermediate goods from the low cost supplier. This decision gives rise to three cases with three different market clearing conditions.[4] Specifically, if the good is non-traded, if the good is imported, if the good is exported.

Below, we describe demand and production in each of these cases .

- **Non-traded.** If the good is non-traded, then the domestic price for the home country must satisfy the following inequality: $\frac{p_w}{\tau} < p_h(\mathbf{s}) < \tau p_w$. That is from home country's perspective it is optimal to source the good domestically and it is not optimal for the home country to export the product.

  In this case, the market clearing condition (for the home country) is:

$$\pi(\mathbf{s}) \left(\frac{p_h(\mathbf{s})}{P_h} \right)^{-\theta} Q_h = z_h \mu_h(\mathbf{s}) \bar{h}, \tag{18}$$

  or that domestic demand equals production. The left-hand-side part is the measure of intermediate goods markets multiplied by the demand in those markets which must equal supply or the productivity of domestic suppliers multiplied by the supply of labor units in that market.

- **Imported.** If the good is imported, then the domestic price for the home country must be $p_h(\mathbf{s}) = \tau p_w$. Why? The home price can not be any larger as arbitrage implies that the domestic price must equal the price of the same imported good from the world. And if the price was lower, then it would not be imported. In this situation with frictional labor markets, we need to take into account that there may be some domestic production. So the quantity of imports is

$$\underbrace{\pi(\mathbf{s}) \left(\left(\frac{\tau p_w}{P_h} \right)^{-\theta} Q_h \right) - z_h \mu_h(\mathbf{s}) \bar{h}}_{\text{imports}} > 0 \tag{19}$$

  The key distinction here relative to the two country model is that we are not specifying world production of the commodity. That is at the given world price p_w, home demand is able to be met. All of this must be larger than zero. Below, it will be useful to define

[4] This is more nuanced than the standard formulation in Eaton and Kortum (2002) due to the frictional labor market. That is there are situations in which an intermediate good is both imported and produced domestically, which is not the case in the Eaton and Kortum (2002) model.

domestic demand as

$$\pi(\mathbf{s})\left(\left(\frac{\tau p_w}{P_h}\right)^{-\theta} Q_h\right) \;=\; z_h \mu_h(\mathbf{s})\bar{h} \;+\; \text{imports}(\mathbf{s}), \tag{20}$$

or domestic consumption must equal domestic production plus imports of that commodity.

- **Exported.** If the good is exported, then it must be that $p_h(\mathbf{s})\tau = p_w$. Why? If the home price was larger, then it would not be purchased on the world market. And the price can not be lower as arbitrage implies that the price of the exported good sold in the world market must equal the prevailing price in that market. At this price, the quantity of exports is

$$\underbrace{\pi(\mathbf{s})\left(\frac{p_w/\tau}{P_h}\right)^{-\theta} Q_h - z_h \mu_h(\mathbf{s})\bar{h}}_{-\text{exports}} \;<\; 0 \tag{21}$$

  The left-hand-side is the measure of intermediate goods markets multiplied by domestic demand net of production. All of this should be negative implying that the country is an exporter. Below, it will be useful to define this in the following way

$$\pi(\mathbf{s})\left(\frac{p_w/\tau}{P_h}\right)^{-\theta} Q_h \;=\; z_h \mu_h(\mathbf{s})\bar{h} \;-\; \text{exports}(\mathbf{s}) \tag{22}$$

The discussion above completely summarizes the demand and supply conditions that must be met in all intermediate goods markets.

The Final Good and Market Clearing. The final good's producer sells the final good to consumers. Thus, we have the following market clearing condition

$$Q_h = C_h + G = \int_z \int_a c(\mathbf{s}, a)\lambda(\mathbf{s}, a) + G, \tag{23}$$

where $c(\mathbf{s}, a)$ is the consumption policy function which satisfies the households problem and $\lambda(\mathbf{s}, a)$ is the mass of consumers with state $\mathbf{s}$ and asset holding a (defined below in (24)) . This relationship says household-level consumption—aggregated across all households— plus government consumption must equal the aggregate production of the final good Q_h.

Market clearing conditions for the intermediate goods in (18),(20), (22), and the aggregate final good in (23) summarize the equilibrium relationship on the production side of the economy.

2.2. Household Side of the Economy

The households in the economy make choices about where to reside, how much to work, and how much to consume. Here we describe the equilibrium conditions associated with these choices. In the discussion below, we define the following functions $\{ \iota_m(\mathbf{s}, a), \iota_n(\mathbf{s}, a), g_a(\mathbf{s}, a) \}$ as the move, work, and asset policy functions that satisfy the households problem in (12).

Population and Labor Supply. We define the probability distribution of households, across assets and states, as $\lambda_h(\mathbf{s}, a)$. Furthermore, define the probability distribution of households in the next period as $\lambda_h'(\mathbf{s}, a)$. The distribution of households evolves across time according to the following law of motion

$$\lambda_h'(\mathbf{s'}, a') = \int_{a:a'=g_a(\mathbf{s},a)} \lambda_h(\mathbf{s}, a)(1 - \iota_m(\mathbf{s}, a))\pi(\mathbf{s'}, \mathbf{s}) + \lambda_h(\mathbf{s}, a)\iota_m(\mathbf{s}, a)\bar{\pi}(\mathbf{s'}). \tag{24}$$

Equation (24) says the following. Next period, the mass of households with asset holding a' in state $\mathbf{s'}$ equals

- The mass of household that do not move multiplied by the transition probability that $\mathbf{s}$ transits to $\mathbf{s'}$. This is the first term in equation (24). Plus...
- The mass of households that do move multiplied by the probability that they end up in state $\mathbf{s'}$. This is the second term in equation (24). The probability, $\bar{\pi}(\mathbf{s'})$, is given by the moving protocol.
- All of this is conditional on those households who choose assets holdings equal to a'. This denoted by the conditionality under the integral sign.

Given a distribution of households, the supply of labor to intermediate good producers with productivity state $\mathbf{s}$ is,

$$\int_a \iota_n(\mathbf{s}, a)\lambda_h(\mathbf{s}, a) = \mu_h(\mathbf{s}), \tag{25}$$

which is the size of the population residing in that market multiplied by the labor supply policy function and integrated over all asset states. This then connects the supply of labor with production in (18)-(22).

Asset Holdings and Consumption. The distribution of asset holdings and consumption take the following form. Next period, aggregate net-asset holdings are

$$\mathcal{A}' = \int_a \int_z g_a(\mathbf{s}, a)\lambda_h(\mathbf{s}, a). \tag{26}$$

We make a couple of points about this. First, this is in aggregate. However, recall that some households in the home country may have positive holdings; others may have negative holdings. Second, net asset holdings must always be claims on foreign assets since there is no domestic asset in positive supply (such as capital).

Using the definition in (26) and the focus on a stationary equilibrium (so $\mathcal{A}' = \mathcal{A}$) we can work from the consumers budget constraint and derive aggregate consumption:

$$C_h = -\mathcal{A}' + R\mathcal{A} + \int_a \int_z \left\{ \tilde{w}(\mathbf{s})\iota_n(\mathbf{s}, a) + w_n(1 - \iota_n(\mathbf{s}, a)) - m\iota_m(\mathbf{s}, a) \right\} \lambda_h(\mathbf{s}, a). \tag{27}$$

In words, aggregate consumption equals net asset purchases (the first two terms) plus wage income and home production net of moving costs.[5]

2.3. Government

We assume that the government runs a balanced budget. Thus, government spending must equal total tax revenues collected or

$$G = \int_a \int_z T(w(\mathbf{s}))\iota_n(\mathbf{s}, a)\lambda_h(\mathbf{s}, a), \tag{28}$$

which says that government spending must equal the tax revenues conditional on working and then integrates over all markets and asset states.

What does the government do in our economy? Spending levels G and the tax progressivity τ_p parameter are exogenously given. The government then picks the average tax rate, λ, such that (28) holds.

2.4. A Stationary Small Open Economy (SSOE) Equilibrium

Given the equilibrium conditions from the production and household side of the economy, we define a "Stationary Small Open Economy (SSOE) Equilibrium" equilibrium.

A Stationary Small Open Economy (SSOE) Equilibrium. Given world prices $\{p_w, R\}$ and government policy $\{G, \tau_p\}$, a stationary Small Open Economy Equilibrium is domestic prices $\{p_h(\mathbf{s}), P_h\}$, tax rate λ, policy functions $\{\ g_a(\mathbf{s}, a), \iota_n(\mathbf{s}, a), \iota_m(\mathbf{s}, a)\ \}$, a probability distribution $\lambda_h(\mathbf{s}, a)$ such that

i Firms maximize profits, (14) and (16) ;

[5] To understand this a bit better, note the connection with (27) and the closed, endowment economy in Huggett (1993); aggregate asset holdings are in zero net supply, thus the first term on the right-hand-side is zero, there is no moving, and there is no labor supply choice. Thus, aggregate consumption must equal aggregate labor income in Huggett's (1993) economy.

ii The policy functions solves the household's optimization problem in (12);

iii Demand for the final and intermediate goods equals production, (18)-(21) and (23);

iv The government budget is balanced (28);

v The probability distribution $\lambda_h(\mathbf{s}, a)$ is a stationary distribution associated with $\{ g_a(\mathbf{s}, a), \iota_m(\mathbf{s}, a), \pi(\mathbf{s'}, \mathbf{s}) \}$. That is it satisfies

$$\lambda_h(\mathbf{s'}, a') = \int_{a:a'=g_a(\mathbf{s},a)} \lambda_h(\mathbf{s}, a)(1 - \iota_m(\mathbf{s}, a))\pi(\mathbf{s'}, \mathbf{s}) + \lambda_h(\mathbf{s}, a)\iota_m(\mathbf{s}, a)\bar{\pi}(\mathbf{s'}). \tag{29}$$

The idea behind the equilibrium definition is the following. The first bullet point (**i**) gives rise to the equilibrium conditions for the demand of intermediate goods in (15) and wages (17) at which firms are willing to produce. The second bullet point (**ii**) says that households are optimizing.

At a superficial level, bullet (**iii**) simply says that demand must equal supply. It is, however, deeper. The choices of the household matter for both the demand and the supply side. Specifically, it requires that prices (and hence wages) must *induce* a pattern of (i) consumption and (ii) labor supply such that demand for goods equals the production of goods.

Bullet point (**v**) requires stationarity. Specifically, the distribution of households across productivity and asset states is not changing. Mathematically, this means that distribution $\lambda_h(\mathbf{s}, a)$ must be such that when plugged into the law of motion in (24), the same distribution is returned.

Finally, note that there is no requirement that the assets market clear, i.e., that (26) equals zero. This is an aspect of the small open economy assumption. At the given world interest rate R, the assets need not be in zero net supply. As we discuss below, this implies that trade need not balance as the trade imbalance will reflect asset income on foreign assets and the acquisition of assets.

Computationally, finding a solution to this economy is straightforward: guess a set of prices, solve the household problem and construct the stationary distribution, then check that the market clearing conditions are satisfied. The open-economy aspect of this economy (that domestic demand and supply may not be equal, but that prices must respect specified bounds) means that the market clearing conditions can be represented as a mixed complementarity problem (see, e.g., Miranda and Fackler (2004)).

3. Model Properties

This section describes some qualitative properties of the model. It borrows from our own parallel work in Lyon and Waugh (2017) that focuses in detail on the workings of the model. Below

we focus on two issues (i) the pattern of trade across labor markets and (ii) how trade changes the distribution of wages.

3.1. Trade

To illustrate the pattern of trade, first define the following statistic:

$$\omega(\mathbf{s}) := \frac{p_h(\mathbf{s}) z_h \mu_h(\mathbf{s})}{p_h(\mathbf{s}) z_h \mu_h(\mathbf{s}) \bar{h} \;\; + \;\; p_h(\mathbf{s}) \text{imports}(\mathbf{s}) - p_h(\mathbf{s}) \text{exports}(\mathbf{s})}, \tag{30}$$

What does equation (30) represent? The denominator is the value of domestic consumption (everything domestically produced plus imported net of exports). The numerator is production. So the interpretation is how much of domestic consumption (at the variety/island level) is the home country producing. This is essentially the micro-level "home share" summary statistic emphasized in Arkolakis, Costinot, and Rodríguez-Clare (2012). As we discuss below, this statistic provides (i) a clean interpretation of a labor markets explore to trade and (ii) is tightly connected with local labor market wages.

Figure 1 plots the home share (raised to the power of inverse θ) by world price and home productivity. There are three regions to take note of: where goods are imported, exported, and non-traded. First, in the region where world prices are low or home productivity is low, this is where the country has a comparative disadvantage at producing these commodities. Thus, the country is predominately an importer of these commodities. And for the lowest regions, the country does not produce any of these commodities. Second, the region where world prices are high or home productivity is very high is the opposite. This is where the country has a comparative advantage and is an exporter of the commodities.

Finally, there is the "table top" region in the middle where the home share equals one for moderate world prices and home productivity. This is the region where the goods are non-traded (home demand equals home supply). Exactly like the inner, non-traded region in the Ricardian model of Dornbusch, Fischer, and Samuelson (1977), the reasons is because of trade costs. In this region, the home country finds that it is best to source the goods from home and not export.

3.2. Trade and Wages

Following our work in Lyon and Waugh (2017), one can connect the pattern of trade across islands/labor markets in 1 with the structure of wages in the economy. More specifically, as we show in the Appendix and in Lyon and Waugh (2017), one can derive (pre-tax) real wages in a market with state variable $\mathbf{s}$ equal

$$w(\mathbf{s}) = \omega(\mathbf{s})^{\frac{1}{\theta}} \hat{\mu}_h(\mathbf{s})^{\frac{-1}{\theta}} z_h^{\frac{\theta-1}{\theta}} C_h^{\frac{1}{\theta}}. \tag{31}$$

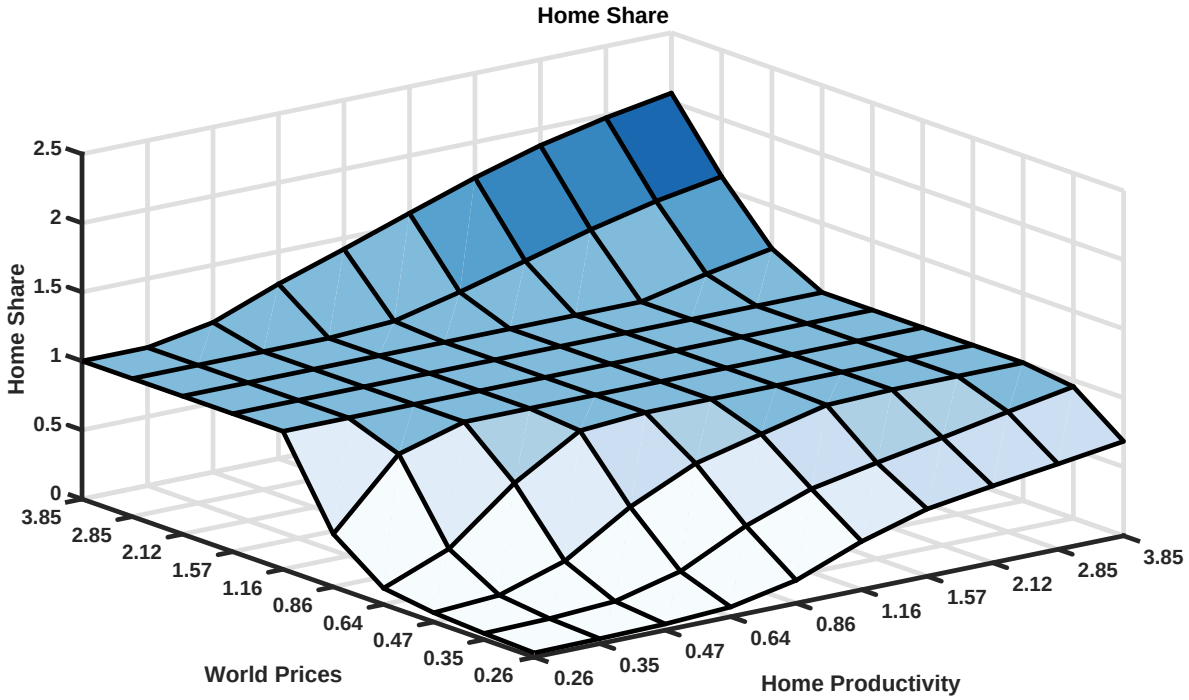


Figure 1: Trade: Home Share, $\omega(\mathbf{s})^{\frac{1}{\theta}}$

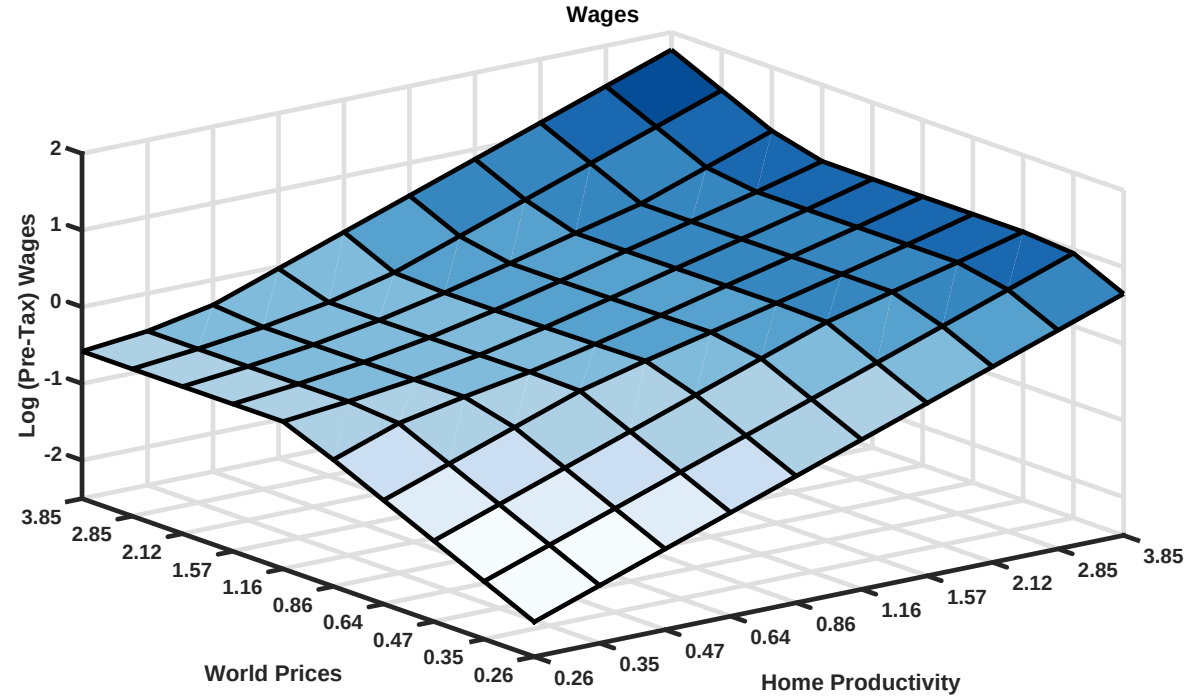


Figure 2: Wages (Pre-tax): Open Economy

where $\hat{\mu}_h(\mathbf{s}) = \frac{\mu_h(\mathbf{s})\bar{h}}{\pi(\mathbf{s})}$ is the number of labor units per firm in market $\mathbf{s}$. Equation (40) connects the home-share measure in 30 with wages. The key thing to notice is that a **smaller** home share implies that wages are **lower** with elasticity $\frac{1}{\theta}$. This means that, if imports (relative to domestic production) are **larger**, then **wages in that labor market are lower**. Similarly, a larger home share—which could be above one if an exporter—, are wages are larger. This is the exact opposite of the aggregate result of Arkolakis, Costinot, and Rodríguez-Clare (2012).

Figure 2 further illustrates these observations by plotting the logarithm of (pre-tax) wages by world price and home productivity. As equation 40 makes clear, note the close correspondence between wages and the home share in Figure 1. Like in Figure 1, there are three regions to take note of.

The first region is where import competition is prevalent (low world prices or low home productivity), wages are low. A simply way to understand this result is the following: wages reflect the value of the marginal product of labor. In import competing sectors, trade results in lower prices and, hence, lower wages. The second region is where exporting is prevalent. Exporting regions are able to capture high world prices and, thus, wages are high in these sectors. Finally, center region is where commodities are non-traded. Here the gradient of wages very much mimics the the increase in domestic productivity. In contrast, where goods are imported or exported, the wage gradient mimics the the change in world prices.

4. Quantitative Analysis

We first describe the calibration (or how we are currently doing it). We essentially use the calibration results from our parallel work in Lyon and Waugh (2017). This is still in progress. In the results below, we simply use existing parameter values from various literatures and/or guesstimates. Below is the discussion that summarizes our approach from Lyon and Waugh (2017) to calibrate the model.

4.1. Calibration

Preferences. Given our specification in (5) there are three parameters to calibrate: B which controls the disutility of working, the the discount factor β, and the utility preference paraemter σ. The disutility term B is picked to target aggregate labor force participation rates. We follow the strategy of Chang and Kim (2007) (who have essentially the same specification) and target a employment to population ratio of sixty percent.

The discount factor is typically calibrated to match certain features of the aggregate wealth distribution. We plan on doing so, but currently set the discount factor equal to 0.95.

We currently set $\sigma = 1$, which corresponds to $u(c) = \log(c)$.

Endowments and Financial Constraints: There are two parameters controlling households endowments: a time endowment and home production. And one parameter controlling the borrowing limit on the household $\bar{a}$.

The time endowment is simply a normalization. We choose $\bar{h} = 1$. Home production is more complicated, but important. In the current calibration, we set the value of home production equal to zero. This is a simplification due to time constraints, but should led to more conservative estimates of how tax progressivity should change. The reason is that value of home production affects the reservation wage of the households and, thus, controls how elastic labor supply is to various shocks (including tax changes). Thus, by setting this value to zero we are picking the most inelastic labor supply specification.

This observation motivates our currently-in-progress calibration strategy. More specifically, we use the evidence from Autor, Dorn, and Hanson (2013) on how labor supply responds to changes in trade-induced wage changes to estimate the home production value. This procedure is currently in progress.

The borrowing limit parameter is calibrated to match properties of the aggregate wealth distribution. Krueger, Mitman, and Perri (2016) report from the Survey of Consumer Finances that approximately forty percent of households have zero or negative wealth. Thus, we chose the borrowing limit so that the model replicates this fact.

Productivity and World Price Process. As specified the productivity process in (2) and (13) leave three parameters to be calibrated: $\{\phi, \sigma_z, \sigma_w\}$, the parameter controlling the persistence of the shocks and the size of the innovations.

Currently, we simplify the process even further and restrict the standard deviation of innovations to productivity to be the same size as the standard deviation of innovations to world prices. Given this restriction, we simply use existing estimates of labor income processes to discipline these parameters. Specifically, we use the estimates from Kaplan (2012) which imply a value of 0.95 for ϕ and a value of 0.026 for σ_z.

We are still working on this as well, but a can make a few comments. Given the relationship between wages and productivity described above, much of the economy will not be import or export exposed and, thus, wages will simply mimic the productivity process z adjusted for the value of θ (which we do above). Thus, it is natural to have our model to replicate these features of the data.

With that said, a key issue is how persistent are shocks that change the nature of comparative advantage. This is important in that it will affect how insurable or not these shocks are and, thus, the desirability of progressive taxation and social insurance. We currently speculate that the results of Krishna and Senses (2014) may be informative; additionally the results of Hanson,

Lind, and Muendler (2015) also speak to the dynamics of comparative advantage as well.

These parameters also play a deeper and less obvious in determining trade elasticities. Much like in the model of Eaton and Kortum (2002), the extent of technology heterogeneity controls how elastic trade flows are to changes in trade costs. Furthermore, per the insights of Arkolakis, Costinot, and Rodríguez-Clare (2012), these parameters will control the aggregate gains from trade as well.

Migration Cost and Location Choice: Currently, the migration cost is picked to match aggregate data on migration rates across commuting zones. We use the the IRS migration data which uses the address and reported income on individual tax fillings to track how many individuals move in or out of a county. We compute that a bit over three percent of households move across a commute zone at a yearly frequency. We pick the migration cost to target this value.

A related issue is the specification of where moving households end up. In the results below, we use the random labor market specification. That is upon moving, a household will end up in a random labor market, with the distribution function being the invariant distribution of labor markets.

We are currently exploring alternatives to calibrating both these pieces of the model. In particular, focusing on the trade specific evidence as to migration in response to trade exposed areas (Autor, Dorn, and Hanson (2013) and Greenland, Lopresti, and McHenry (2016)) and focusing on the worker-level responses observed in Autor, Dorn, Hanson, and Song (2014).

Tax Function and Government Spending. We set government spending, G to be twenty percent of GDP. This is consistent with National Accounts data for the US over the past forty years. Given that we calibrate the model to the US economy, we must take a stand on the current progressiveness of the tax code. We use the estimates from Heathcote, Storesletten, and Violante (2014) and set τ_p equal to 0.18.

Demand Elasticity θ. In the current results now, we simply take a stand on the this parameter. Consistent with a wide range of estimates we set its value equal to four.

We are currently pursuing alternatives to estimating this parameter. In particular, the result in 40 shows how θ affects the intensity of pass through of changes in trade exposure to wages within a labor market. This suggests the empirical strategy of using the evidence in Autor, Dorn, and Hanson (2013) as an empirical moment for the model to target.

Trade Costs. Finally, we set the iceberg trade cost to target an aggregate import to GDP ratio of ten percent. This is consistent with the aggregate level of trade in in the mid 1990's, pre-China WTO accession.

We have two other extensions that are currently in progress. One is an explicit treatment of

tariffs as a policy instrument. The second is an asymmetric change in trade costs. For example, an explicit reduction in market access to the US without a corresponding reduction in market access to the rest of the world.

5. How Does Optimal Progressivity Vary with Openness?

Given the calibrated model, we ask several questions of it: First, what is the optimal policy and what trade-offs does the policy maker face? Second, how do the gains from moving to an optimal policy change with increased openness to trade? Third, how does optimal policy change? Finally, how to these conclusions depend upon the particular calibration (e.g. labor supply and migration elasticities).

Below, we describe the social welfare function the policy maker is maximizing and then answer each of these questions.

5.1. Social Welfare Function

We focus on a utilitarian planner which places equal weight on households within the domestic economy. That is

$$W(\tau_p,\ \tau) = \int_z \int_a V(a,\ \mathbf{s}) \lambda_h(\mathbf{s}, a). \tag{32}$$

or the value function of an agent integrated over markets and asset states with respect to the stationary distribution of households across those states $\lambda_h(\mathbf{s}, a)$ (thus no weight is given to foreign agents). Here we explicitly index social welfare by the tax progressivity parameter and by the trade cost. Thus, given the social welfare function in (32), the optimal degree of progressivity is

$$\tau_p^*(\tau) = \arg\max \ \ W(\tau_p,\ \tau). \tag{33}$$

That is, the tax progressivity parameter that maximizes social welfare. Here we make explicit that the tax progressivity parameter depends upon the trade cost or the extent to which the economy is open or closed. When presenting results, we convert units into consumption equivalent values. That is the permanent increase in consumption that must be allocated to make an agent indifferent between living in the baseline economy and an economy with an alternative progressivity parameter τ_p.

5.2. Optimal Progressivity and the Insurance, Migration Trade-off

What is the optimal policy and what trade-offs does the planner face? We explore this question by traceing out social welfare for different levels of τ_p or tax progressivity. All other parameters

are fixed. Furthermore, we report welfare in consumption equivalent values relative to the baseline economy (thus, when $\tau_p = 0.18$, the welfare gain equals zero since this is the baseline economy).

Figure 3 shows that our economy generates a "Laffer-like" curve—social welfare displays an inverted "U" shape as progressivity varies. In our calibration, the US economy lies to the left of the optimal policy. In particular, the optimal progressivity is found to be 0.28 versus current progressivity 0.18 as measured by Guner, Kaygusuz, and Ventura (2014). **This implies that something map this into a tax system average tax rates...**

With that said, welfare gains are very small from a move to optimally progressive system. Less than one tenth of a percent. In other words, the costs of not having an optimal system are small. This observation is not generically true. As we discuss below, on key issue about increased trade exposure is that these costs start to grow.

The reason behind the arc like shape in Figure 3 is the tradeoff between distorting labor supply and migration and, hence, reducing the size of the "pie" versus providing better social insurance.

To show that insurance improves with tax progressivity, Figure 4 plots the regression coefficient from the projection of household consumption on household pre-tax labor income. This coefficient is analogous to the tests of risk sharing in Cochrane (1991) and Townsend (1994). If the coefficient is zero, then interpretation is that the household is completely insured against household-level income shocks (as in the complete markets allocation). The key observation from Figure 4 shows that pass-through of income to consumption systematically declines as the tax system becomes more progressive. To flip it around, "insurance" is better with a more progressive tax code. This is welfare improving.

Gains from insurance are offset by losses in economic efficiency. In our model, the loss in efficiency is *not* about labor supply. Labor supply changes by less than half a percentage point over the whole span of tax progressivity parameters on the x-axis in Figure 3. This largely follows from our choice to model labor supply as being purely on the extensive margin and we set the outside option of not being in the labor force to zero. Labor supply the micro level is inelastic given the choice set of the household level. Labor supply at the macro level is also very inelastic since the extensive margin responds very little as well.

What does reduce economic efficiency is the reduction of migration and allocative efficiency. Figure 5 plots the percentage point difference in migration rates relative to the baseline economy. When moving to a flat tax to a progressivity parameter of 0.30 migration rates fall by nearly one percentage point. Relative to the baseline with a three percent migration rate, this means that migration rates are falling by a third.

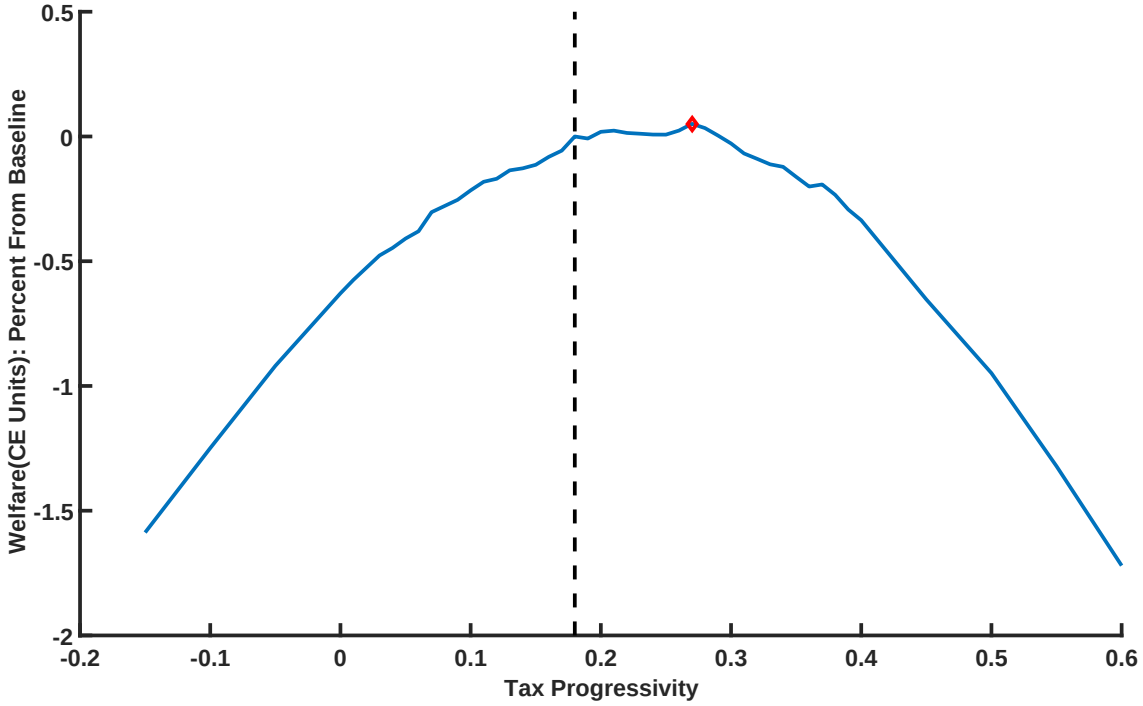


Figure 3: Social Welfare at Calibrated Model

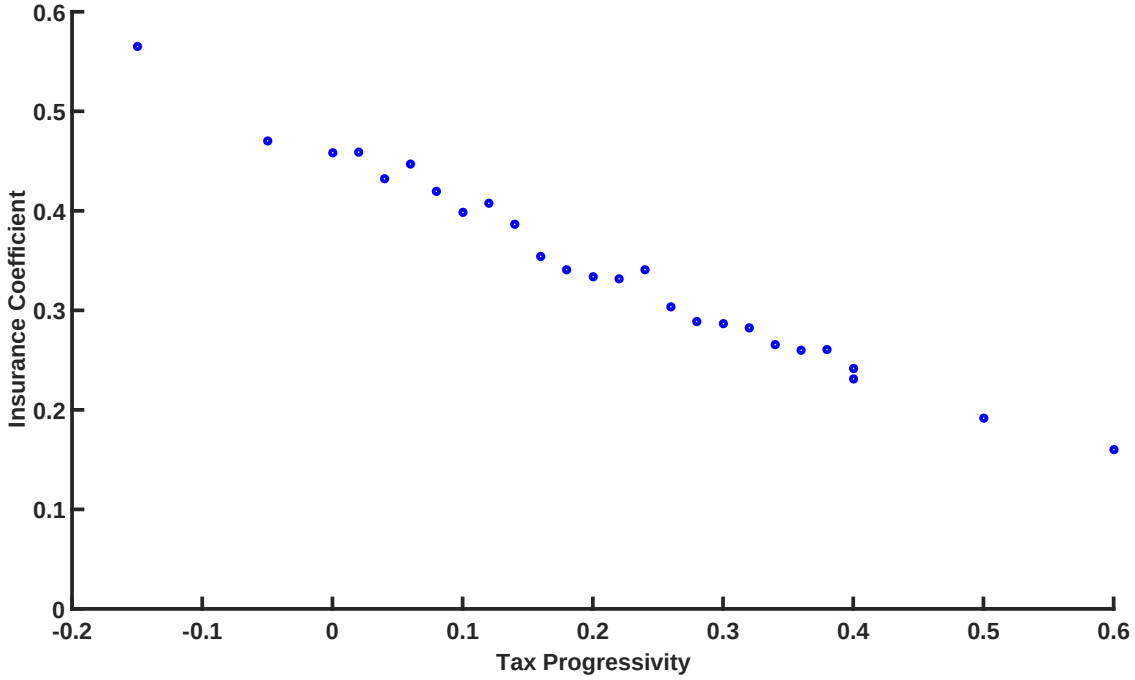


Figure 4: Insurance and Tax Progressivity

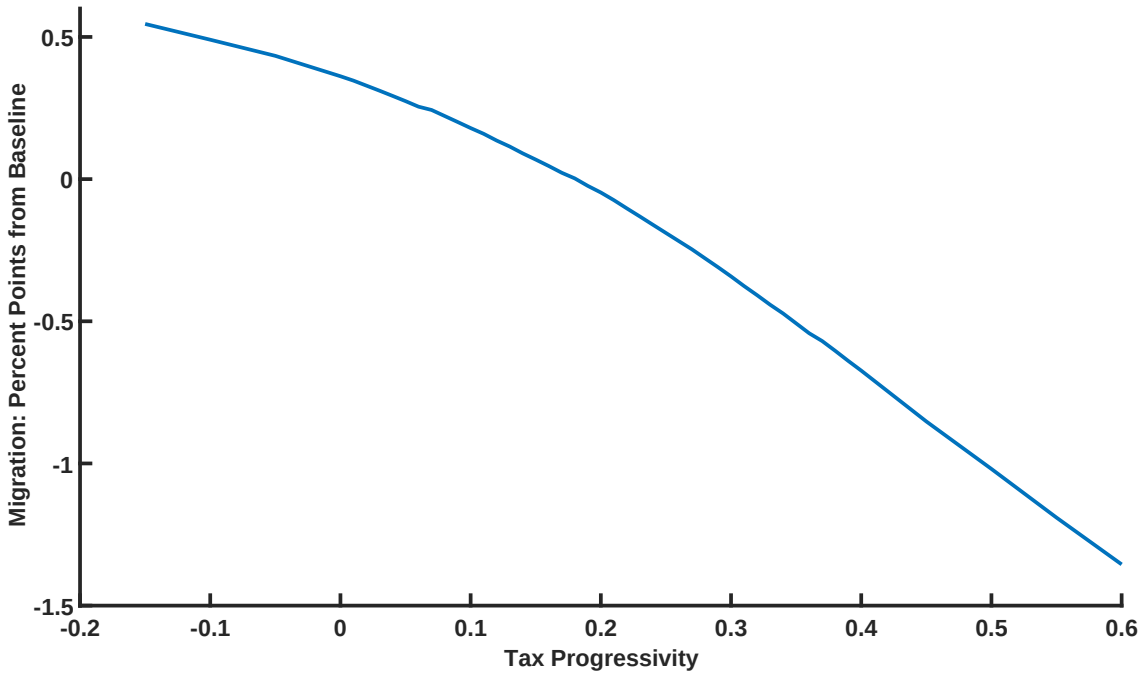


Figure 5: Migration and Tax Progressivity

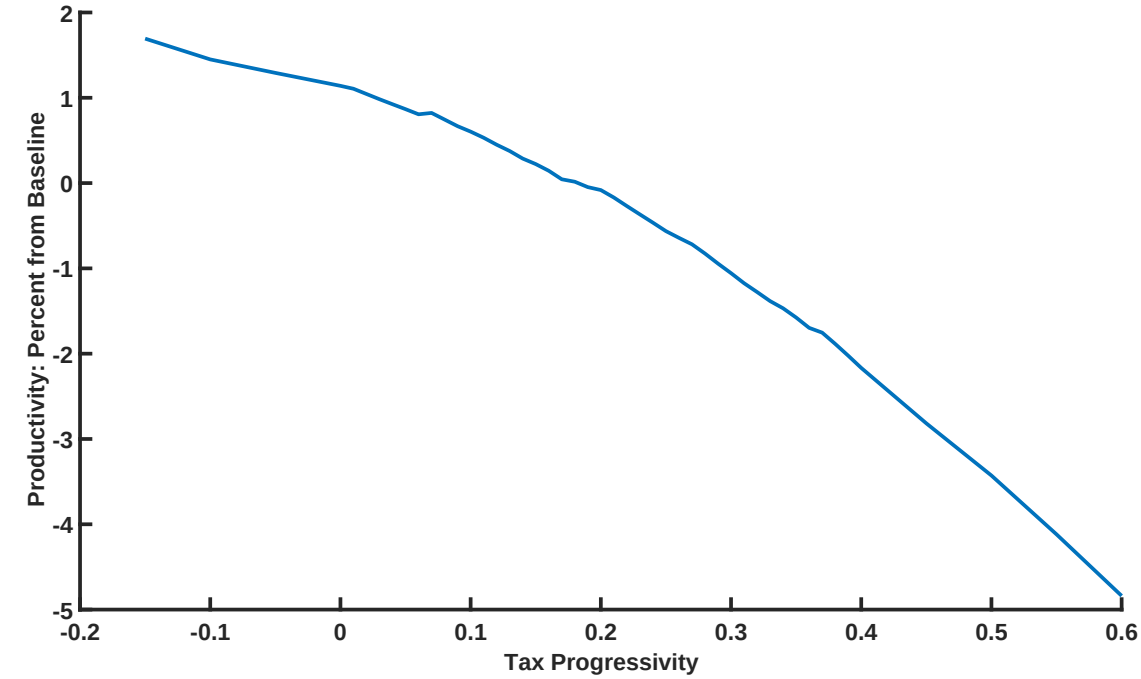


Figure 6: Aggregate Productivity and Progressivity

The reason for this observation is that migration is playing a substitute for insurance. The motive to migrate in this model is to move in response to negative economic shocks and attempt to smooth consumption but finding a better income realization. As the tax system becomes more progressive, this provides better insurance and, hence, households chose to seek less an alternative (and costly) form of insurance, i.e. migration.

Migration is important for allocative efficiency. Absent labor market frictions, the ideal allocation would have households locate in the most productive, high comparative advantage sectors of the economy. Higher migration rates help achieve this. However, as migration rates decline, allocative efficiency starts to fall and aggregate productivity falls.

Figure 6 illustrates this point showing how output per worker falls by large amounts. For example, a move from a flat tax to a progressivity parameter of 0.30 lowers aggregate productivity by nearly three percent—about two years of economic growth at current rates in the US.

To summarize: Figure 3 illustrates how social welfare changes with the progressivity in the current calibration of our model. As in previous analysis, this Laffer-like curve reflects a tradeoff between insurance and economic efficiency. What is unique about our setting is that the reduction in economic efficiency is not primarily about labor supply issues, but migration issues and allocative efficiency that arises as labor becomes less mobile as the tax system becomes more progressive.

5.3. Optimal Progressivity and Trade

how does optimal policy change with increased openness to trade? Figure 7 compares several different regimes. Table 1 reports optimal values for the same different levels of openness. We find that the tax system should generally become more progressive as the economy becomes more open. However, only at relatively large levels of openness (at least for the US, but comparable to economies such as Canada and Mexico) do we see quantitatively large, prescribed changes in the optimal policy.

A different question is how does openness change the costs from moving away from optimal (or current) policy? This question is about concavity of the "Laffer-like" curve in social welfare and less about how optimal policy changed. A key feature of Figure 7 is that the curve becomes more concave. This implies that the costs associated with moving a way from the optimal or current policy become larger as the economy become more open. For example, a move to a flat tax system when the economy is very open would lead to a two and half percent decrease in welfare where as at the baseline the move to a flat tax system would lead to a half a percent decrease in welfare.

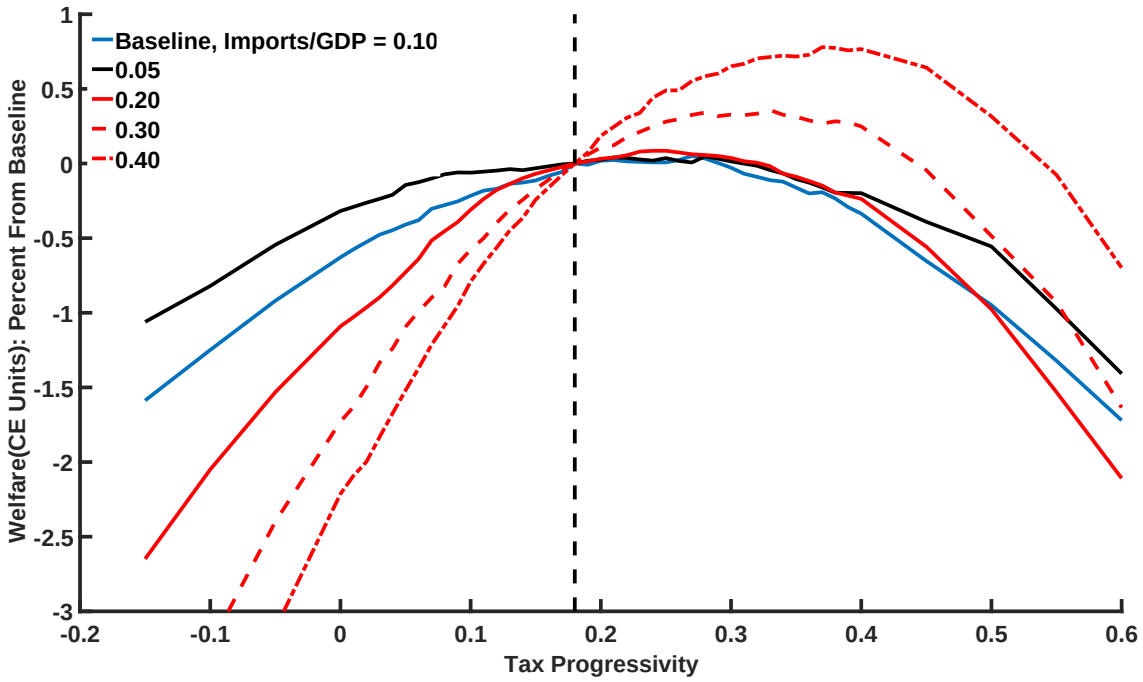


Figure 7: Welfare and Progressivity

6. Conclusion

Still work in progress.

Table 1: Openness and Optimal Progressivity

Imports/GDP	Optimal Progressivity	Gains to Moving to Optimum	Gains to Moving to Flat
0.05	0.273		
0.10	0.207		
0.20	0.243		
0.30	0.331		
0.40	0.371		

References

ANTRÀS, P., A. DE GORTARI, AND O. ITSKHOKI (2016): "Globalization, Inequality and Welfare," Discussion paper, National Bureau of Economic Research.

ARKOLAKIS, C., A. COSTINOT, AND A. RODRÍGUEZ-CLARE (2012): "New Trade Models, Same Old Gains?," *American Economic Review*, 102(1), 94–130.

ARTUÇ, E., S. CHAUDHURI, AND J. MCLAREN (2010): "Trade shocks and labor adjustment: A structural empirical approach," *The American Economic Review*, 100(3), 1008–1045.

ATTANASIO, O., AND S. J. DAVIS (1996): "Relative Wage Movements and the Distribution of Consumption," *Journal of Political Economy*, 104(6), 1227–1262.

AUTOR, D., D. DORN, AND G. H. HANSON (2013): "The China Syndrome: Local Labor Market Effects of Import Competition in the United States," *The American Economic Review*, 103(6), 2121–2168.

AUTOR, D. H., D. DORN, G. H. HANSON, AND J. SONG (2014): "Trade adjustment: Worker-level evidence," *The Quarterly Journal of Economics*, 129(4), 1799–1860.

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

CALIENDO, L., M. A. DVORKIN, AND F. PARRO (2015): "Trade and labor market dynamics," .

CHANG, Y., AND S.-B. KIM (2007): "Heterogeneity and aggregation: Implications for labor-market fluctuations," *The American Economic Review*, 97(5), 1939–1956.

COCHRANE, J. H. (1991): "A simple test of consumption insurance," *Journal of political economy*, pp. 957–976.

CONESA, J. C., AND D. KRUEGER (2006): "On the optimal progressivity of the income tax code," *Journal of Monetary Economics*, 53(7), 1425–1450.

DIX-CARNEIRO, R. (2014): "Trade liberalization and labor market dynamics," *Econometrica*, 82(3), 825–885.

DORNBUSCH, R., S. FISCHER, AND P. A. SAMUELSON (1977): "Comparative advantage, trade, and payments in a Ricardian model with a continuum of goods," *The American Economic Review*, 67(5), 823–839.

EATON, J., AND S. KORTUM (2002): "Technology, geography, and trade," *Econometrica*, 70(5), 1741–1779.

EATON, J., AND H. S. ROSEN (1980): "Optimal redistributive taxation and uncertainty," *The Quarterly Journal of Economics*, 95(2), 357–364.

GREENLAND, A., J. LOPRESTI, AND P. MCHENRY (2016): "Import Competition and Internal Migration," .

GUNER, N., R. KAYGUSUZ, AND G. VENTURA (2014): "Income taxation of US households: Facts and parametric estimates," *Review of Economic Dynamics*, 17(4), 559–581.

HANSON, G. H., N. LIND, AND M.-A. MUENDLER (2015): "The dynamics of comparative advantage," Discussion paper, National Bureau of Economic Research.

HEATHCOTE, J., K. STORESLETTEN, AND G. L. VIOLANTE (2014): "Optimal tax progressivity: An analytical framework," Discussion paper, National Bureau of Economic Research.

HUGGETT, M. (1993): "The risk-free rate in heterogeneous-agent incomplete-insurance economies," *Journal of economic Dynamics and Control*, 17(5), 953–969.

KAMBOUROV, G. (2009): "Labour Market Regulations and the Sectoral Reallocation of Workers: The Case of Trade Reforms," *Review of Economic Studies*, 76(4), 1321–1358.

KAPLAN, G. (2012): "Inequality and the life cycle," *Quantitative Economics*, 3(3), 471–525.

KRISHNA, P., AND M. Z. SENSES (2014): "International Trade and Labour Income Risk in the U.S.," *The Review of Economic Studies*, 81(1), 186–218.

KRUEGER, D., K. MITMAN, AND F. PERRI (2016): "Macroeconomics and Household Heterogeneity," *Handbook of Macroeconomics*, 2, 843–921.

LUCAS, R. E., AND E. C. PRESCOTT (1974): "Equilibrium search and unemployment," *Journal of Economic Theory*, 7(2), 188–209.

LYON, S., AND M. WAUGH (2017): "Quantifying the Losses from International Trade," .

MIRANDA, M. J., AND P. L. FACKLER (2004): *Applied computational economics and finance*.

MIRRLEES, J. (1974): "Notes on welfare economics, information and uncertainty," *Essays on economic behavior under uncertainty*, pp. 243–261.

MIRRLEES, J. A. (1971): "An exploration in the theory of optimum income taxation," *The review of economic studies*, 38(2), 175–208.

TOWNSEND, R. M. (1994): "Risk and insurance in village India," *Econometrica: Journal of the Econometric Society*, pp. 539–591.

VARIAN, H. R. (1980): "Redistributive taxation as social insurance," *Journal of public Economics*, 14(1), 49–68.

A. Connection with National Accounts

This section connects these equilibrium relationships to national income and product accounts (NIPA). This will help facilitate an understanding of the connection between trade imbalances and household's consumption-savings decisions.

The Income Side of NIPA. It is first useful to start from an income side measure of production (or GDP) in our economy. Given competition, the value of aggregate production of the final good must equal aggregate payments for intermediate goods plus. The latter equals aggregate payments to labor in the production of all intermediate goods.

$$P_h Y_h = \int_{\mathbf{s}} P_h w_h(\mathbf{s}) \mu_h(\mathbf{s}) \tag{34}$$

Examining (27) and (34) allows us to connect aggregate income with consumption

$$P_h Y_h = P_h C_h + P_h G - P_h R \mathcal{A} + P_h \mathcal{A}' - P_h \int_a \int_z w_n (1 - \iota_n(\mathbf{s}, a)) \lambda_h(\mathbf{s}, a) + P_h \int_a \int_z m \iota_m(\mathbf{s}, a) \lambda_h(\mathbf{s}, a) \tag{35}$$

so aggregate income equals consumption (private and public) minus (i) returns on assets (ii) new purchases of assets (iii) home production and (iv) plus moving costs. This basically says that income/production must equal consumption net of income not associated with production (i.e. returns on assets and home production) plus "investment" in assets and moving costs. For example, if consumption is larger than income one reason is that (in aggregate) households (on net) are borrowing from abroad ($\mathcal{A}' < 0$).

Production Side of NIPA. The the value of aggregate production of the final good must equal the value of intermediate goods production

$$P_h Y = \int_{\mathbf{s}} p_h(\mathbf{s}) z \mu_h(\mathbf{s}) \tag{36}$$

which we can then connect with the expenditure side of GDP through the market clearing conditions for intermediate goods and final goods. Specifically, by connecting the production side with the demand side for non-traded goods in (18), imports in (20) and exports in (22) and the equating final demand with consumption we have

$$P_h Y = P_h C_h + P_h G + \int_z \text{exports}(\mathbf{s}) - \int_z \text{imports}(\mathbf{s}). \tag{37}$$

Or GDP equals consumption (private and public) plus exports minus imports.

Savings, Trade Imbalances, and Capital Flows. Finally, we can connect the income side and

the production side the national accounts to arrive at a relationship between asset holdings and trade imbalances. By working with both (35) and (37) we get the following relationship

$$P_h Y - P_h C_h - P_h G = \int_z \text{exports}(\mathbf{s}) - \int_z \text{imports}(\mathbf{s}), \tag{38}$$

$$= - P_h r \mathcal{A} + P_h(\mathcal{A}' - \mathcal{A}) - P_h \int_a \int_z w_n(1 - \iota_n(\mathbf{s}, a))\lambda_h(\mathbf{s}, a) + P_h \int_a \int_z m\iota_m(\mathbf{s}, a)\lambda_h(\mathbf{s}, a),$$

where r is the net real interest rate. This relationship says the following: aggregate savings equals the trade imbalance. And this, in turn, we can connect the trade balance with the savings decisions of the households. That is the trade balance equals payments on net asset holdings plus net change in asset holdings (adjusted for home production and moving costs).[6] That last statement is essentially a statement about international capital flows. To see this, consider the special case where moving costs are zero and no home production. Then we have the relationship

$$P_h Y - P_h C_h - P_h G = \int_z \text{exports}(\mathbf{s}) - \int_z \text{imports}(\mathbf{s}) = -P_h r \mathcal{A} + P_h(\mathcal{A}' - \mathcal{A}). \tag{39}$$

Here if exports are greater than imports, then this implies that the households in the home country are doing several things. The trade surplus may reflect that households (on net) are making debt payments ($r\mathcal{A}$ is negative). Second, the trade surplus may reflect that the households (on net) are acquiring foreign assets ($\mathcal{A}' - \mathcal{A}$ is positive). Finally, note that in a stationary equilibrium, the trade imbalance only reflects payments from foreign asset holdings. This implies that the current account and capital account are always zero in a stationary equilibrium.

B. Connection with Autor, Dorn, and Hanson (2013)

This discussion is borrowed from Lyon and Waugh (2017) but presented here for convenience.

This section connects wages with trade exposure in a structural way. There are basically two insights that deliver the key result in Proposition 1. First, in a competitive environment, wages reflect the value of the marginal product of labor. This implies that any change in wages through trade exposure works through prices—the "value" part. Second, the CES production structure tightly links changes in prices and changes in quantities. Together, these two insights provide a link between wages and a quantity based measure of trade exposure much like in Autor, Dorn, and Hanson (2013).

[6] To map this into Balance of Payments language: the trade imbalance plus foreign income payments is the current account; the capital account is the net change in foreign asset holdings; then (we suspect) home production and moving costs would show up as the "balancing item."

To arrive at this conclusion, take the relationships in (18), (20), (21) and "combine" them by noting that demand must simply equal domestic production plus imports minus exports. Then manipulating the demand curve to arrive at a relationship between prices and the share of domestic production relative to consumption. Then connecting this relationship with the wage we have Proposition 1 summarized below.

Proposition 1 (Wages and Trade Exposure) *Real wages in a market with state variable* $\boldsymbol{s}$ *equal*

$$w(\boldsymbol{s}) = \omega(\boldsymbol{s})^{\frac{1}{\theta}} \hat{\mu}_h(\boldsymbol{s})^{\frac{-1}{\theta}} z_h^{\frac{\theta-1}{\theta}} C_h^{\frac{1}{\theta}}. \tag{40}$$

where

$$\omega(\boldsymbol{s}) := \frac{p_h(\boldsymbol{s}) z_h \mu_h(\boldsymbol{s})}{p_h(\boldsymbol{s}) z_h \mu_h(\boldsymbol{s}) \bar{h} \;+\; p_h(\boldsymbol{s}) \text{imports}(\boldsymbol{s}) - p_h(\boldsymbol{s}) \text{exports}(\mathbf{s})}, \tag{41}$$

which is the production of goods relative to consumption or the "home share," and $\hat{\mu}_h(\boldsymbol{s}) = \frac{\mu_h(\boldsymbol{s})\bar{h}}{\pi(\boldsymbol{s})}$ *is the number of labor units per market.*

Lets talk through Proposition 1 in the following way. First, what does equation (41) represent? The denominator is the value of consumption (everything domestically produced plus imported net of what I shipped out via exports). Then the numerator is production. So the interpretation is how much of domestic consumption is the home country producing. This is essentially the micro-level "home share" summary statistic emphasized in Arkolakis, Costinot, and Rodríguez-Clare (2012).

Equation (40) connects the home-share measure with wages. The key thing to notice is that a **smaller** home share implies that wages are **lower** with elasticity $\frac{1}{\theta}$. This means that, if imports (relative to domestic production) is large, then wages in the production of that commodity must be lower. Similarly, a larger home share—which could be above one if an exporter—, are wages are larger.

To understand this, consider the case when imports are increasing (relative to domestic production) in a market, i.e. the home share is falling. What is going on is that the increase in imports (holding all else fixed) implies that prices must be falling. Because the wage equals the value of the marginal product of labor, then the wage must be falling. The CES production structure just ties the price with the home share through the parameter θ in a very succinct manner.

Finally, notice the role that θ plays. The parameter θ is critical in determining how much productivity risk is associated with a market. For example, if $\theta = 1$ (i.e. the production technology is Cobb-Douglas), then wages become independent of productivity.

Now we can connect this with wage regression performed in ADH that linked changes in

wages with changes in trade exposure. To do so, start with (40) and take log differences across time

$$\Delta \log w(\mathbf{s}) = \frac{1}{\theta}\Delta \log \left(\omega(\mathbf{s})/\hat{\mu}_h(\mathbf{s})\right) + \frac{1}{\theta}\Delta \log C_h + \Delta \log \left(z_h^{\frac{\theta-1}{\theta}}\right), \tag{42}$$

which says that the change in wages across locations is summarized by (i) trade exposure via the change in per-worker home share, (ii) the change in location-specific productivity and (iii) aggregate consumption.

Equation (42) makes clear that an instrumental variable strategy is necessary to identify the causal effect of trade exposure on wages. Generally, domestic productivity is unobserved to the econometrician, but it is correlated with home trade share. That is imports could (and hence the home share could decrease) increase because either world prices changed or domestic productivity changed. Our notation makes this clear in that ω is a function of $\mathbf{s} = \{z_h, p_w\}$.

The structure of the model suggests several instrumental variable strategies. One valid instrument would be to use the world price (if observed) directly. The world price is orthogonal to domestic productivity (the exclusion restriction), yet correlated with the home trade share. The exclusion restriction follows from our small open economy assumption and the specification the stochastic process in (13) that is assumed to be orthogonal to z_h.[7]

An alternative strategy would be to use another country's imports as an instrument. Another country's imports would be orthogonal to the home country's productivity (the exclusion restriction), but correlated with world prices. This, in fact, is quite similar to the instrument proposed in ADH. That is ADH use Chinese imports to non-US advanced economies as an instrument to identify the causal effect of Chinese imports on US wages.

There are several empirical issues in going from the specification in (42) to that used in ADH. First, they only have industry level imports at the aggregate level and a geographic region has multiple industries operating within it. In other words, a commuting zone more closer to a collection of islands in our model. If one made the assumption that one industry corresponds with one region, then their ΔIPW measure would be imports per worker. So similar to $\Delta \log \left(\omega(\mathbf{s})/\hat{\mu}_h(\mathbf{s})\right)$ but operating directly with the level of imports rather than the shares. We explore this distinction in future work.

[7] This discussion makes clear that in general equilibrium, one should be concerned that a change in domestic productivity would feed into world prices and, thus, invalidate this strategy.