

Quantifying the Losses from International Trade

PRELIMINARY AND WORK IN PROGRESS.

Spencer G. Lyon
New York University

Michael E. Waugh
New York University and NBER

December 2018

ABSTRACT

A popular narrative suggests that trade with China harmed the US economy in the 2000s—a rapid rise in import penetration, expanding trade deficits, and local-labor-market evidence on the negative impact of trade on wages and employment are all pieces of evidence pointing in this direction. We provide an interpretation of this evidence by developing a dynamic, standard incomplete market model with Ricardian trade and frictional labor markets. We calibrate our model to match the local-labor-market evidence and then study the transition dynamics of the economy in response to a “China Shock.” Despite being consistent with the evidence of Autor, Dorn, and Hanson (2013), the China shock induces a boom in the aggregate economy: a large increase in GDP and employment, a modest increase in consumption, and an improving trade deficit. Underlying the aggregate effects is much heterogeneity in the gains from trade; however, very few actually lose from trade.

Email: spencer.lyon@stern.nyu.edu, mwaugh@stern.nyu.edu. We would like to thank our discussant Martin Beraja, Jonathan Heathcote, and participants at the 2017 Annual Meeting of the Society for Economic Dynamics, Cowles Foundation International Trade Summer Conference, The Ohio State University, Federal Reserve Bank of Dallas, University of California-San Diego, Princeton University.

Trade generates winners and losers, but that the winners win more than the losers lose. This phrase is said often, but does little to assuage the concerns of people and policy makers about the forces of globalization—that the losses from trade are large and that there are insufficient mechanisms to insure against these losses. This paper takes one step toward evaluating these concerns by answering the question: How much do the losers lose from trade?

A popular narrative hints that the losses from trade were large. Stating in early 2000s and continuing into the beginning of 2008, US imports as a fraction of GDP increased by seventy percent—Chinese imports accounted for all of this growth. At the same time, the US trade deficit as a fraction of GDP expanded by two percentage points suggesting that growth in import penetration did not arrive with equivalent export and employment opportunities. Autor, Dorn, and Hanson (2013) provides evidence to support this idea. Using evidence from differential changes across local labor markets, they find that increased exposure of a labor market to Chinese trade resulted in reductions in labor income and labor force participation.

This paper provides a quantitative, structural interpretation of this narrative and the associated evidence. A structural evaluation is important for two reasons. The first reason relates to the research design behind the Autor, Dorn, and Hanson (2013) evidence. Approaches that exploit differential changes in local labor markets are only able to make statements about the differential effect—not levels. In case of Autor, Dorn, and Hanson (2013), this critique is relevant as an optimistic interpretation of the evidence is that all labor markets gained from trade with China, just some gained more than others. One contribution of this paper is to develop a model which can bridge the gap between the evidence of Autor, Dorn, and Hanson (2013) and the aggregate effects of trade exposure.

A second reason relates to normative issues. Even if there were losses in labor income and employment opportunities, open questions are: How large are the losses in welfare? Does policy have any ability to improve outcomes? The key issue here is the extent of market incompleteness and the ability of households to adjust to trade shocks. The standard gains from trade argument is that there exists a transfer scheme to compensate the losers from trade—i.e., even with losses in the labor market, transfers allow everyone to enjoy the aggregate gains from trade. Moreover, in this world, policy has no ability to improve outcomes. In contrast, if households are an island to oneself with no transfers or insurance markets, the welfare losses could be large and government interventions may improve welfare. A second contribution of this paper is examine the distributional effects of trade in a standard, intermediate setting—the standard incomplete markets model.

Our model builds on existing neoclassical trade theory and then departs from it in important ways. As in the model of Eaton and Kortum (2002), there is a continuum of goods with competitive producers who are heterogenous in productivity; comparative advantage determines

the pattern of trade. As in Lucas and Prescott (1974) the labor market is frictional and labor can only move across different goods producing markets (within a country) after paying some cost.

We depart from recent quantitative trade theory by studying an environment where insurance markets are incomplete. This is a model where households are an island to oneself and transfers or insurance markets do not exist, but as in the standard incomplete markets model (Huggett (1993), Aiyagari (1994)) households can self-insure by accumulating a non-state contingent asset.¹ In contrast to essentially all models of trade and labor market dynamics, this aspect of the model allows for the partial—but not complete—pass through of income shocks into consumption. This formulation provides a middle ground between a complete markets benchmark where the gains and losses from trade are shared equally and where all households are hand-to-mouth with no opportunities to smooth out shocks.

In addition to the traditional self-insurance channel, we endow households with two additional margins to mitigate shocks. First, we allow households to migrate and move to alternative labor markets. This margin has been an important dimension in models of trade and labor market dynamics (see, e.g., Kambourov (2009), Artuç, Chaudhuri, and McLaren (2010), Dix-Carneiro (2014), Caliendo, Dvorkin, and Parro (2015)). Second, households can exit the labor force and substitute into leisure. Our formulation of labor supply focuses on the extensive margin and follows the work of Rogerson (1988) and more closely Chang and Kim (2007).

Our analysis proceeds in several steps. First, we show how our model provides a structural interpretation of the local labor market regressions in Autor, Dorn, and Hanson (2013). The insight is the following. In a competitive labor market, wages reflect the value of the marginal product of labor. With CES demand, a log-linear link arises between the value of labor and the production of an island's good relative to national consumption—the home expenditure share which is *the* summary statistic for the exposure of a labor market to trade. Most importantly, the elasticity of substitution in the CES demand system controls how changes in the home share pass through to wages. Because the home share is directly related to Acemoglu, Autor, Dorn, Hanson, and Price (2016) IP measure and Autor, Dorn, and Hanson's (2013) IPW measure, their results when viewed through our model are informative about the elasticity of substitution in the demand system.

Our model also provides a motivation and rationalization for Autor, Dorn, and Hanson's (2013) instrumental variable strategy. In our model, unobserved to the econometrician is the local, physical marginal product of labor, i.e. the island level productivity shock. Through comparative advantage, this shock influences how exposed a labor market is to trade and, thus, it's

¹The quantitative nature of our analysis distinguishes this paper from earlier discussions about risk, incomplete markets, and the gains from trade, e.g., see Newbery and Stiglitz (1984) and Eaton and Grossman (1985). Furthermore the standard incomplete markets setting take the reason that insurance markets are missing as given; Dixit (1987, 1989a,b) shows how certain conclusion depend upon the modeling as to why these markets are missing.

omission biases the OLS projection of labor market outcomes on trade exposure. Under a small open economy assumption in our model, a valid instrument is another country's imports— it's orthogonal to the local productivity shock, but through world prices it's correlated with trade exposure. This is essentially the same instrument proposed in Autor, Dorn, and Hanson (2013).

These insights provide the foundation for the calibration of our model. The approach is to have the model match traditional aggregate and cross-sectional evidence and have key parameter values disciplined by the evidence from Autor, Dorn, and Hanson (2013) and migration evidence from Greenland, Lopresti, and McHenry (2017). The later is done by exploiting the points above by constructing data from our model and running the same, instrumented regressions of Autor, Dorn, and Hanson (2013) and Greenland, Lopresti, and McHenry (2017) on model generated data. Per our arguments above, this calibration approach pins down key parameters of interest: the elasticity of substitution in demand which controls how trade shocks pass through to wages; a household's value of non-market activity which controls how trade shocks affect participation; and a household's preference to move which partially shapes how trade shocks affect migration.

An important feature of our approach is that in our calibration and analysis we focus on transition paths. The expansion of Chinese trade is modeled as a change in the technological ability of the economy to import. Given this change we compute the transition path from the initial stationary equilibrium to the new stationary equilibrium. Several sets of moments, including those of Autor, Dorn, and Hanson (2013) and Greenland, Lopresti, and McHenry (2017), are targeted as the economy transits between stationary equilibrium.

The final step is to use the calibrated model as laboratory to answer two questions.

What were the aggregate effects of the China Shock? Over five years, the China shock leads to a two percent increase in GDP and a one percentage increase in consumption. About half of the increase in output is from standard, gains from trade effects; the remaining half comes from increases in labor force participation. In terms of employment, the China shock increased employment by about a million and a half jobs.

The increase in labor force participation and muted consumption response are due to the presence of incomplete markets and precautionary savings motives. An important mechanism in our model is that a reduction in trade costs increases the amount of uninsurable income risk in the model. The idea here is that while many households are not directly import exposed, a reduction in trade costs increases the likelihood that they will eventually become trade exposed and, thus, leads to an increase in income risk. This force increases precautionary savings motives (see, e.g., Huggett and Ospina (2001)) and, in turn, leads to an increase savings and labor supply to acquire those savings.

These findings are important relative to the "popular narrative" discussed above. First, the

aggregate labor supply response illustrates well the problems of interpreting approaches that exploit differential changes in local labor markets. By construction, our model is completely consistent with the local-labor market evidence of Autor, Dorn, and Hanson (2013)—that import exposure is correlated with declines in labor force participation. Yet—at the aggregate level—labor force participation is increasing in response to the trade shock.

A second issue relates to the trade deficit. The popular narrative suggests that if it becomes easier to import, then this leads to an expansion in the trade deficit. In our model, the trade deficit decreases—a reduction in the cost of importing, lead exports to expand by more than imports. This suggests that attributing an expansion of the US trade deficit in the early 2000s to increased trade with China is wrong and it fits well with the narrative that global financial forces advocated by Bernanke (2005) are a more likely culprit.

What were the welfare effects of the China Shock? On average, the gains from trade are a little less than one percentage point which more or less matches up with the average increase in real wages and output per worker which is a bit more than one percent. The interesting outcomes and that are model is designed to speak to are the distributional consequences which can be demarcated into three regions based on their initial exposure to trade. Those who are initially import exposed experience small welfare losses, though some experience small, positive gains. A few losses, minor gains, stands in sharp contrast to the harsh effects trade has on these labor market—wages falling by two percentage points over the first five years of the shock. This arises from households are actively using the margins they have access to to mitigate the effects of the shock.

Regions that are more export exposed gain the most, a bit less than two percent in consumption equivalents. These gains come purely through changes in the aggregate price index with real wages increasing by four percent after five years. Given that aggregate savings are rising, what is occurring is that households on islands of these types are taking some of the gains in the labor market and banking it for times when things are not so great.

Perhaps the most surprising finding is what occurs in the non-traded regions. In real terms, wages are essentially unaffected. While the real purchasing power is increasing, the value of the goods produced on these islands falls as consumers substitute into cheaper imported goods. Despite being isolated and producing goods that are not traded, they are still affected by import competition. In welfare terms, these households still gain about 0.75 percent in consumption equivalents. As mentioned, these gains are not from changes in real wages, but arise from the stochastic nature of comparative advantage in our model that at some point they will be export exposed at some point and reap the benefits.

Related Literature. 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), Coşar, Guner, and Tybout (2016)). We depart from this literature by studying an economy in which households face labor income shocks and incomplete markets. The cost of this departure is that we are unable to incorporate the the geographic and sectoral detail found in this work (see, e.g., Caliendo, Dvorkin, and Parro (2015)) due to computational complexities.

We view this as an important departure for several reasons. First, it because it breaks the link between labor market outcomes and welfare. And it breaks the link in non-obvious ways. For example, if households are sufficiently wealthy, they can smooth out the trade induced losses to labor income and, thus, the losses from trade are small. This is an example where households do have a mechanism to mitigate the losses from trade through self-insurance. However, if workers in import-competing sectors are already wealth poor, then the welfare losses from trade could be quite large. This is an example where the previously poor bear the burden of globalization and inequality is amplified.

Second, it opens the door to questions about government policy and the correct response when the economy is increasingly exposed to trade. As we show in our parallel work in Lyon and Waugh (2018), we find that progressivity should increase with openness to trade and that progressivity is an important tool to mitigate the negative consequences of globalization.

1. Motivating Facts

There is a body of evidence suggesting that trade with China harmed the US economy in the 2000s. This section discusses this evidence and it provides benchmarks that we will ask our quantitative experiment to replicate.

1.1. Macro Facts

Fact 1: Large rise in US Import Exposure. Figure 1 plots US imports of goods and services divided by GDP for the time period 1990 until the beginning of 2008. The red line is US imports in total, the blue line is all but China, thus the difference reflects Chinese imports relative to GDP. US import exposure grew dramatically during this time period from approximately 10 percent of GDP to near 17 percent of GDP by the beginning of 2008.

The blue line shows that non-Chinese US import exposure was roughly constant over this time at around 10 percent. What this implies that nearly **all** the rise in US import exposure during this time period came from an increase in Chinese import exposure. Moreover, most of this expansion was after China's accession to the World Trade Organization at the end of 2001.

Fact 2: An Expansion in the US Trade Deficit. Figure 2 plots US exports minus imports divided by GDP for the time period 1990 until the beginning of 2008. During this time period, the trade




Figure 1: US Imports




Figure 2: US Trade Deficit

deficit grew by about five percentage points in total; three percentage points prior to the large expansion in Chinese import expansion, and then two percentage points subsequently.

From a trade perspective, these two facts suggest that the “traditional” replacement of job opportunities did not occur with the rise in Chinese trade. That is import exposure was not offset by an increase export opportunities and, thus, was a drag on the US labor market in the 2000s as the facts below suggest. An alternative perspectives on Figure 2 is the “global savings glut” hypothesis put forth by Bernanke (2005). Developments in international credit markets during the 1990s and 2000s lead to an glut of foreign savings (some of which is Chinese related) and the US borrowed these funds and led to an expansion of the trade deficit.

1.2. Micro Facts

The next three facts focus on labor market outcomes from Autor, Dorn, and Hanson (2013). The idea behind Autor, Dorn, and Hanson (2013) is to exploit changes in the variation in trade exposure at the commuting zone level (see Tolbert and Sizer (1996)) and correlate it with changes in labor market outcomes such as earnings and labor force participation. Their measure of trade exposure is:

$$\Delta IPW_{uit} = \sum_j \left(\frac{L_{ijt}}{L_{it}} \right) \left(\frac{\Delta M_{ucjt}}{L_{ijt}} \right) \quad (1)$$

where u stands for united states, c stands for China, i is a commute zone, t is time, and j is industry. This measure essentially takes aggregate US imports from China M_{ucjt} for industry j and essentially apportions these imports to a commute zone based on that commute zones share of national employment in that industry. It then aggregates across industries for that commute zone.

Given this measure of trade exposure, they estimate its affect on labor market outcomes with the following empirical specification:

$$\Delta L_{it} = \gamma_t + \beta \Delta IPW_{uit} + \text{controls}_i + \epsilon_{it}. \quad (2)$$

A key issue is that the error term may incorporate factors that are simultaneously changing a commute zone’s trade exposure and labor market outcomes. As our model makes clear (see Section 4.3), local productivity shocks would be a key threat to identification, i.e. they would both change the pattern of trade and labor markets outcomes. To avoid this problem, Autor, Dorn, and Hanson (2013)) estimate (2) using an other countries imports from China as an instrument.

In the facts below, we report estimates after standardizing the ΔIPW_{uit} . That is we demeaned

Table 1: ADH Evidence: Labor Market Outcomes and Trade Exposure

	Δ Labor Earnings	Δ NILF
Standardized Δ IPW	-4.30 [-6.62, -2.00]	1.11 [0.52, 1.72]

Note: Values in brackets report 95-5 confidence intervals. Variable definitions are as follows. Δ Labor Earnings is average household “wage and salary” income per adult variable from the Census and ACS; the units are in decadal, percent changes. Δ NILF corresponds to the change in the not in labor force share. Δ IPW is standardized by netting out the mean and dividing by the standard deviation.

this measure and divided by its standard deviation. As in Autor, Dorn, and Hanson (2013) all labor market outcomes are converted to 10 year changes.

Fact 3: Import Exposure Decreased Household Income The first column in Table 1 reports Autor, Dorn, and Hanson’s (2013) estimate of the response of commute zone household level income to trade exposure. What this coefficient means is that a one standard deviation increase in trade exposure reduced wage growth by four percent over 10 years. To put this in context, average wage growth over the period 2000-2007 was only about six percent (converted to 10 year change). In other words, these estimates suggest that moderate exposure to trade eliminates 2/3 of expected wage growth.

Fact 4: Import Exposure Increased Non Participation The second column in Table 1 reports the response of commute zone nonparticipation in the labor force in response to trade exposure. The magnitude of this coefficient means is that a one standard deviation increase in trade exposure reduced participation in the labor market by 1.11 percentage points. Average non-participation across commute zones, across all time periods was about 25 percent.

Fact 5: Muted Migration Response. The first column in Table 2 reports the estimate of the response of commute zone population to trade exposure. What this coefficient means is that a one standard deviation increase in trade exposure reduced a commute zone’s population by -0.05 log points, moreover this effect is not statistically different than zero. This suggests that trade exposure did not induce households to reallocate across labor markets as standard trade theory would predict.

Greenland, Lopresti, and McHenry (2017), however, call into question this particular finding of ADH. For our purposes, an important issue is the treatment of pre-trends of population growth in commute zones. In the ADH specification, regional controls (e.g. midwest, pacific northwest, etc) are included and thus, control in broad terms for different population trends across regions.

Table 2: ADH and GLM Evidence: Migration and Trade Exposure

	ADH Δ Population	GLM, Δ Population
Standardized Δ IPW	-0.05 [-1.51, 1.41]	-1.43 [-3.33, 0.48]

Note: Values in brackets report 95-5 confidence intervals. Variable definitions are as follows. Δ Population corresponds with the log change in population. Δ IPW is standardized by netting out the mean and dividing by the standard deviation. Greenland, Lopresti, and McHenry (2017) (GLM) replace ADH regional controls with agged population growth at the commute zone level.

Greenland, Lopresti, and McHenry (2017) suggest instead using lagged population growth at the commute zone level. Visual inspection of pre-trends and regression evidence suggest pre-trends at the commute zone are an issue.

The second column of Table 2 reports the effects when regional controls are replaced with lagged population. Here the the point estimate increases by an order of magnitude, e.g. from essentially zero to -1.40 log points. That is a one standard deviation increase in trade exposure reduces a commute zones population by almost one and a half log points. The dramatic change in point estimates suggest that prior to the ADH window of 1990, commute zones that would eventually experience growth in trade exposure, were experiencing population growth in the 1980s. Not taking this fact into account biased the ADH estimates upward. With that said, at conventional levels of statistical significance, the point estimate is not statistically different from zero.

Greenland, Lopresti, and McHenry (2017) raise two other issues that lead to finding out migration from trade exposure that lead to similar findings Table 2. One issue is the use of Census data versus nationally representative samples, use of Census data delivers similar point estimates in Table 2 but stronger and significant effects for the young vs. old. A second issue is the time window, extending the window to 2010 leads to statistically significant effects with point estimates comparable to Table 2. In sum, relative to the body of evidence provide by Greenland, Lopresti, and McHenry (2017), ADH's specification is the only specification that finds no migration response.

1.3. Discussion

Taken together, these facts suggest a compelling narrative: At the macro level, there was a large increase in import exposure from China and no corresponding increase in export opportunities.

And at the micro level, the evidence suggest this wreaked havoc on labor market in the United States—households absorbed loses in labor income, stopped participating in the labor market, and with some migration from trade exposed regions.

How to interpret this evidence? First, there is an issue about the interpretation of the cross-sectional estimates in (1) and how to map them into aggregate conclusions. The model that we develop below plays an important role here. And as we describe in Section 4 the model bridges this gap, clarifies what these estimates are informative about, and provide key inputs into our qualitative analysis.

Second, a model is needed to jump from changes in income and non-participation to statements about welfare. Our presumption is that while households have limited access to direct insurance markets, they do have a myriad of ways to smooth out negative labor market outcomes. One is self-insurance, another is that leisure is a substitute for leaving the labor market, a third is migration. Thus, we develop a model that allows us to entertain multiple mechanisms for which households can mitigate the negative outcomes that may arise from international trade.

A final issue relates to the trade deficit. Traditional trade theory's perspective on this suggests that increases in import penetration from China did not come with increases in export and employment opportunities for those displaced from trade. This perspective ignores the idea that a rising trade deficit implies households are increasing their consumption above and beyond their savings. In other words, the of China came with an expansion of both intra- and inter-temporal trading opportunities.

2. Model

Here, we describe a model of international trade with households facing incomplete markets and frictions to move across labor markets. The first subsection discusses the production structure; the second subsection discusses discusses the households.

Below, since we focus on the perspective of one country, country subscripts are omitted unless necessary. Similarly, time subscripts are omitted unless necessary.

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

Competitive firms produce intermediate goods with linear production technologies,

$$q(\omega) = z(\omega)\ell, \quad (3)$$

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

$$\log z_{t+1} = \phi \log z_t + \epsilon_{t+1}, \quad (4)$$

where ϵ_{t+1} is distributed normally with mean zero and standard deviation σ_ϵ . The innovation ϵ_{t+1} is independent across time, goods, and countries.

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

Transporting intermediate goods across countries is costly. Specifically, consumers and firms face iceberg trade costs when importing and exporting their products. We allow for the import and export cost to differ with $\tau_{im} > 1$ being the cost to import a good from abroad and $\tau_{ex} > 1$ being the cost an export faces to ship goods onto the world market.

Intermediate goods are aggregated by a competitive final-goods producer who has a standard CES production function:

$$Q = \left[\int_0^1 q(\omega)^\rho d\omega \right]^{\frac{1}{\rho}}, \quad (5)$$

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 $\theta = \frac{1}{1-\rho}$.

2.2. Households

Within a country, there is a continuum of infinitesimally small households of mass L . Each household is infinitely lived and maximizes expected discounted utility

$$E \sum_{t=0}^{\infty} \beta^t \left\{ \log(c_t) - B \frac{h_t^{1-\gamma}}{1-\gamma} + \nu_t^i \right\}, \quad (6)$$

where E is the expectation operator and β is the subjective discount factor. Period utility depends on both consumption of the final good, the disutility of labor, and a preference shock ν_t^i . The preference shock ν_t^i is indexed by i which corresponds with choice of the household to move

or not. As in Artuç, Chaudhuri, and McLaren (2010) and Caliendo, Dvorkin, and Parro (2015), this preference shock is independently and identically distributed across time and is distributed Type 1 extreme value distribution with scale parameter σ_ν .

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 in which it can work. Given their current location, households can choose to work, to move and work someplace else in the future, and to 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 in which the household resides. In the following presentation, we normalize the value \bar{h} equal to one. If a household does not work, it receives home production wh . The value of home production partially determines the value of being out of the labor force and hence, the elasticity of labor supply on the extensive margin.

Households can move to an alternative intermediate-goods sector ω' at some cost. Paying m in units of the final good allows the household to change where it can work in next period. We assume that the new location is a random labor market. Moving also triggers the realization of the ν preference shock associated with the move.

Households residing in a intermediate-goods location face labor income risk associated with fluctuations in local productivity and 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 . We treat R as exogenous and not solved for in equilibrium. An interpretation is that this country faces a large supply of assets at this rate. Households face a lower bound on asset holding $-\bar{a}$, so agents can acquire debt up to the value \bar{a} .

State Variables. The individual state variable of a household are its location, asset holdings a , and preference shocks ν^i . The island-level state variable is the domestic productivity state and world price state. The aggregate state is a distribution over island-level state variables and asset holdings.

Let us expand on this a bit more. The wage per efficiency unit that a household receives is an important island-level object impacting individual decisions. The wage per efficiency unit depends on the value of the marginal product of labor on that island. The marginal product depends on a country's productivity level. The "value" part depends on (i) the world price and (ii) the labor supply decisions of households residing on the island. Given our preference specification in (6), households' labor supply decisions depend on the distribution of asset holdings within the island. Thus, this is where the aggregate state matters for island-level outcomes.

The presentation below depicts a stationary equilibrium. That is, the aggregate state—the distribution over island level states and assets holdings—is constant.² Thus, to conserve on notation, we only carry around the households specific state variables: its own asset holdings, preference shocks, and island-level state variables associated with its location. In particular, let \mathbf{s} denote the domestic productivity and world price combination associated with that island. Furthermore, because the CES aggregator is symmetric over varieties, it is sufficient to index islands by their productivity and world price state. The wage per efficiency unit a household earns is $w(\mathbf{s})$.

Budget Constraints. Given the description of the environment, the budget constraints are as follows. For households that are working, the household's period t budget constraint (all denominated in units of the final good) is

$$a_{t+1} + c_t + \iota_{m,t}m \leq Ra_t + w_t(\mathbf{s}), \quad (7)$$

where the left-hand side are expenditures on new assets, consumption, and possibly moving costs with $\iota_{m,t}$ being an indicator function equaling one if a household moves and zero otherwise. The right-hand side are income payments from asset returns and labor income.

If a household is not working, then the budget constraint is modified to ensure that home production is not used to accumulate assets or pay for moving costs. In this case, a non-working households budget constraint is

$$c_t \leq w_h + |Ra_t - a_{t+1} - \iota_{m,t}m|^+, \quad (8)$$

where the right-hand side is home production plus any positive income from net asset holdings, net of moving costs. What this formulation insures is the home prosecution is just that. It can not be used for purchases in the market of assets of moving costs.

Recursive Formulation. The recursive formulation of the household's problem is

$$V(a, \mathbf{s}, \nu) = \max [V^{s,w}, V^{s,nw}, V^{m,w}, V^{m,nw}]. \quad (9)$$

that is a discrete choice among 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. Unpacking each of these four options is the following. The value of staying and working

²Given an island with state \mathbf{s} , denote the measure of agents with asset holdings a as $\lambda(\mathbf{s}, a)$. Stationarity implies that this value is constant.

is

$$V^{s,w}(a, \mathbf{s}, \nu) = \max_{a' \geq -\bar{a}} [u(Ra + w(\mathbf{s}) - a') - B + \nu^s + \beta EV(a', \mathbf{s}', \nu')], \quad (10)$$

where u is the utility value over consumption and ν^s is the preference shock associated with staying in its current location. The value of staying and not working is

$$V^{s,nw}(a, \mathbf{s}, \nu) = \max_{a' \geq -\bar{a}} [u(w_h + |Ra - a'|^+) + \nu^s + \beta EV(a', \mathbf{s}', \nu')]. \quad (11)$$

The value of moving and working is

$$V^{m,w}(a, \mathbf{s}, \nu) = \max_{a' \geq -\bar{a}} [u(Ra + w(\mathbf{s}) - a' - m) - B + \nu^m + \beta V^m(a')], \quad (12)$$

where there are several key distinctions relative to (10). First, the moving cost, m is paid. Second, ν^m is the preference shock associated with moving. Third, the continuation value is $V^m(a')$ or the value associated with a move. Finally, the value of moving and not working is

$$V^{m,nw}(a, \mathbf{s}, \nu) = \max_{a' \geq -\bar{a}} [u(w_h + |Ra - a' - m|^+) + \nu^m + \beta V^m(a')]. \quad (13)$$

3. Equilibrium

We close the model by focusing on a small open economy equilibrium. The small open economy assumption is that there is no feedback from home country actions into world prices.³

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

$$\log p_w(\omega)_{t+1} = \phi \log p_w(\omega)_t + \epsilon(\omega)_{t+1}, \quad (14)$$

where $\epsilon(\omega)_t$ is distributed normally with mean zero and standard deviation σ_w and is independent of the innovation to the home country's productivity ϵ_t . We express these prices in units of the numeraire, which we take to be the final good in the home country.

A Note on Notation. We denote $\pi(\mathbf{s})$ as the stationary distribution of productivity states and world prices induced by (4) and (14). And denote $\mu(\mathbf{s})$ as the measure of households working on an island with state \mathbf{s} . This value is defined later in (27) and integrates over the labor supply choice of households (which depends upon their individual states and preference shocks).

³Relative to the trade and labor market dynamics 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)) of being relatively simple, yet allows us to be specific about the interaction between trade flows and capital flows.

3.1. 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.

Final Goods Production. The final-goods producer's problem is:

$$\max_{q(\mathbf{s})} P_h Q - \int p(\mathbf{s})q(\mathbf{s})\pi(\mathbf{s})d\mathbf{s}, \quad (15)$$

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

$$q(\mathbf{s}) = \left(\frac{p(\mathbf{s})}{P_h} \right)^{-\theta} Q. \quad (16)$$

where Q is the aggregate demand for the final good; P_h is the price associated with the final good which will be carried around briefly, but is ultimately normalized to the value one.

Intermediate Goods Production. The intermediate-goods-producer's problem is

$$\max_{q(\mathbf{s}), \ell(\mathbf{s})} p(\mathbf{s})q(\mathbf{s}) - w(\mathbf{s})\ell(\mathbf{s}) \quad (17)$$

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

$$w(\mathbf{s}) = p(\mathbf{s})z \quad (18)$$

or the value of the marginal product of labor. Only at the wage in (18) 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 $p(\mathbf{s})$, $\tau_{im}p_w$, p_w/τ_{ex} . The final-goods producer purchases intermediate goods from the low-cost supplier. This decision gives rise to three cases with three different market-clearing conditions: if the good is non-traded; if the good is imported; and if the good is exported.⁴

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: $\frac{p_w}{\tau_{ex}} < p(\mathbf{s}) < \tau_{im}p_w$. That is, from the home country's perspective, it is optimal to source the good domestically and not optimal for the home

⁴This is more nuanced than the standard formulation in Eaton and Kortum (2002) due to the frictional labor market. In our model, 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.

country to export the product.

In this case, the market-clearing condition is:

$$\left(\frac{p(\mathbf{s})}{P_h}\right)^{-\theta} Q = z(\mu(\mathbf{s})/\pi(\mathbf{s})) \quad (19)$$

or that domestic demand equals production. The left-hand part is demand and the right-hand side is supply. That is the 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 $p(\mathbf{s}) = \tau_{im}p_w$. Why? If the price were lower, then it would not be imported. If the domestic price were higher, then the good will be imported with not domestic production and, thus, the prevailing domestic price will equal the imported price. With frictional labor markets, there may be some domestic production so the quantity of imports is

$$\underbrace{\left(\left(\frac{\tau_{im}p_w}{P_h}\right)^{-\theta} Q\right) - z(\mu(\mathbf{s})/\pi(\mathbf{s}))}_{\text{imports}} > 0. \quad (20)$$

That is home demand (net of home production) is met by imports of the commodity. Rearranging gives

$$\left(\left(\frac{\tau_{im}p_w}{P_h}\right)^{-\theta} Q\right) = z(\mu(\mathbf{s})/\pi(\mathbf{s})) + \text{imports}(\mathbf{s}) \quad (21)$$

or domestic demand equals domestic production plus imports.

- **Exported.** If the good is exported, then the prevailing price must be $p(\mathbf{s})\tau_{ex} = p_w$. Why? If the home price were larger, then the good 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. Finally, note that only the trade cost, not the tariff, matters here. At this price, the quantity of exports is

$$\underbrace{\left(\frac{p_w/\tau_{ex}}{P_h}\right)^{-\theta} Q - z(\mu(\mathbf{s})/\pi(\mathbf{s}))}_{-\text{exports}} < 0 \quad (22)$$

or domestic demand net of production which should be negative, implying that the coun-

try is an exporter. Rearranging gives

$$\left(\frac{p_w/\tau_{ex}}{P_h}\right)^{-\theta} Q = z(\mu(\mathbf{s})/\pi(\mathbf{s})) - \text{exports}(\mathbf{s}) \quad (23)$$

or domestic demand equals domestic production minus exports.

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

$$Q = C = \int_{\mathbf{s}} \int_a \int_{\nu} c(\mathbf{s}, a, \nu) \lambda(\mathbf{s}, a, \nu) d\nu da d\mathbf{s}, \quad (24)$$

where $c(\mathbf{s}, a, \nu)$ is the consumption policy function that satisfies the households' problem, and $\lambda(\mathbf{s}, a, \nu)$ is the mass of consumers with state \mathbf{s} , asset holding a , and preference shock ν (defined below in (??)). This relationship says that household-level consumption—aggregated across all households—must equal the aggregate production of the final good Q .

Market-clearing conditions for the intermediate goods in (19),(21), (23) and the aggregate final good in (24) summarize the equilibrium relationship on the production side of the economy.

3.2. 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— $\{ \iota_m(\mathbf{s}, a, \nu), \iota_n(\mathbf{s}, a, \nu), g_a(\mathbf{s}, a, \nu) \}$ —as the move, work, and asset policy functions that satisfy the households' problem in (9).

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

$$\lambda'(\mathbf{s}', a', \nu') = \phi(\nu') \int_{\mathbf{s}} \int_{\nu} \int_{a: a' = g_a(\mathbf{s}, a, \nu)} \lambda(\mathbf{s}, a, \nu) (1 - \iota_m(\mathbf{s}, a, \nu)) \pi(\mathbf{s}', \mathbf{s}) + \lambda(\mathbf{s}, a, \nu) \iota_m(\mathbf{s}, a, \nu) \bar{\pi}(\mathbf{s}') da d\nu d\mathbf{s}. \quad (25)$$

Equation (25) says that in the next period, the mass of households with asset holding a' in state \mathbf{s}' with preference shock ν' equals several terms. First, the probability that preference shocks ν' are realized where ϕ is the probability density function associated with the Type 1 extreme value distribution. The second term is 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 (25). The third

term is 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 (25). The probability, $\bar{\pi}(\mathbf{s}')$, is given by the moving protocol—i.e., random assignment across islands according to the invariant distribution associated with $\pi(\mathbf{s}', \mathbf{s})$. All of this is conditional on those households that choose asset holdings equal to a' , this is denoted by the conditionality under the innermost integral sign. And integrated across preference shocks and current island state \mathbf{s} .

Population and Labor Supply. Given a distribution of households, we define the population of households on islands with state \mathbf{s} as

$$\int_{\nu} \int_a \lambda(\mathbf{s}, a, \nu) da d\nu. \quad (26)$$

In words, the population is found by just integrating over the mass across asset holdings and preference shocks.⁵ Then the supply of labor to intermediate good producers with productivity state \mathbf{s} is,

$$\mu(\mathbf{s}) = \int_{\nu} \int_a \iota_n(\mathbf{s}, a, \nu) \lambda(\mathbf{s}, a, \nu) da d\nu, \quad (27)$$

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 (19)-(23).

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

$$\mathcal{A}' = \int_{\nu} \int_a \int_{\mathbf{s}} g_a(\mathbf{s}, a, \nu) \lambda(\mathbf{s}, a, \nu) d\mathbf{s} da d\nu, \quad (28)$$

where $g_a(\mathbf{s}, a, \nu)$ is the policy function describing asset holdings tomorrow, given the states today. A couple of points about this are warranted. First, this is in aggregate—some households may have positive holdings, while 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).

3.3. 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.

⁵This calculation should be distinguished by the population on an individual island. This latter value is found by dividing 26 by the measure of that island type.

A Stationary Small Open Economy (SSOE) Equilibrium. Given world prices $\{p_w, R\}$, a stationary Small Open Economy Equilibrium is domestic prices $\{p(\mathbf{s})\}$, policy functions $\{g_a(\mathbf{s}, a, \nu), \iota_n(\mathbf{s}, a, \nu), \iota_m(\mathbf{s}, a, \nu)\}$, and a probability distribution $\lambda(\mathbf{s}, a, \nu)$ such that

- i Firms maximize profits, (15) and (17) ;
- ii The policy functions solve the household's optimization problem in (9);
- iii Demand for the final and intermediate goods equals production, (19), (20), (22) and (24);
- iv The probability distribution $\lambda(\mathbf{s}, a, \nu)$ is a stationary distribution associated with $\{g_a(\mathbf{s}, a, \nu), \iota_m(\mathbf{s}, a, \nu), \pi(\mathbf{s}', \mathbf{s}, \nu)\}$. That is, it satisfies

$$\lambda(\mathbf{s}', a', \nu') = \phi(\nu') \int_{\mathbf{s}} \int_{\nu} \int_{a: a'=g_a(\mathbf{s}, a, \nu)} \lambda(\mathbf{s}, a, \nu) (1 - \iota_m(\mathbf{s}, a, \nu)) \pi(\mathbf{s}', \mathbf{s}) + \lambda(\mathbf{s}, a, \nu) \iota_m(\mathbf{s}, a, \nu) \bar{\pi}(\mathbf{s}') da d\nu ds. \quad (29)$$

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 (16) and wages (18) at which firms are willing to produce. The second bullet point (ii) says that households are optimizing.

At a superficial level, bullet (iii) says that demand must equal supply. It's meaning, however, deeper. The households' choices of the 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 (iv) requires stationarity. Specifically, the distribution of households across productivity and asset states is not changing. Mathematically, this means that distribution $\lambda(\mathbf{s}, a, \nu)$ must be such that when plugged into the law of motion in (25), the same distribution is returned.

Finally, note that there is no requirement that the asset market clears—i.e., that (28) 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. This implies that trade need not balance, as the trade imbalance will reflect asset income on foreign assets and the acquisition of assets. After adjusting for moving costs, this implies that the current account and capital account are always zero in a stationary equilibrium, but that trade may be imbalanced.

Computation. Computing a stationary equilibrium for this economy deserves some discussion. First, this economy is unlike standard incomplete markets models in which only one or two prices (e.g., one wage per efficiency unit and/ or the real interest rate) must be solved for. In

contrast, we must solve for an equilibrium *function* $p(\mathbf{s})$. Thus, the iterative procedure is to (i) guess a price function; (ii) solve the household’s dynamic optimization problem; (iii) construct the stationary distribution $\lambda(\mathbf{s}, a, \nu)$; (iv) check whether markets clear; and (v) update the price function. See, e.g., Krusell, Mukoyama, and Şahin (2010), who solve a similar problem.

Second, while the household’s problem contains three state variables, the i.i.d. assumption on ν allows us to integrate out the preference shock, compute choice probabilities, and then work directly on the distribution across islands and asset states, i.e. $\lambda(\mathbf{s}, a)$.⁶ The Type 1 extreme value distribution allows us to perform this integration step in closed form. In addition to allowing our model to match certain targets, these preference shocks also make the aggregate economy continuous in its price (modulo the discussion below) and parameter space which facilitates the quick computation of solutions to a stationary equilibrium.

Third, an important observation is that the inequalities in (20) and (22) impose additional structure on an equilibrium. The observation is that when domestic demand and supply are not equal, the price in those markets must respect bounds on international arbitrage. This implies that the problem of finding a price function consistent with a stationary equilibrium can be represented as a mixed complementarity problem (see, e.g., Miranda and Fackler (2004)). Appendix B provides a complete description of our solution procedure and links to our code repository.

4. Model Properties

This section describes some qualitative properties of the model. Below we focus on three issues (i) the pattern of trade across labor markets (ii) how trade exposure affects wages and how our model relates to the empirical approach/specification of Autor, Dorn, and Hanson (2013). Finally, we use these results to motivate our quantitative exercise.

4.1. Trade

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

$$\omega(\mathbf{s}) := \frac{p(\mathbf{s})z\mu(\mathbf{s})}{p(\mathbf{s})z\mu(\mathbf{s}) + p(\mathbf{s})\text{imports}(\mathbf{s}) - p(\mathbf{s})\text{exports}(\mathbf{s})}. \quad (30)$$

What does equation (30) represent? The denominator is the value of domestic consumption: everything domestically produced plus imports minus exports. The numerator is production. The interpretation of (30) is how much of domestic consumption at the island/variety level the home country is producing. This is similar to the micro-level “home share” summary statistic

⁶To see this point, simply examine (25) and first integrate out the ν ’s. Integrating the individual policy functions over the ν s are the choice probabilities of a particular action, given state \mathbf{s}, a , which take the familiar logit form given the distributional assumption.

emphasized in Arkolakis, Costinot, and Rodríguez-Clare (2012). As we discuss below, this statistic (i) provides a clean interpretation of a labor market’s exposure to trade and (ii) is tightly connected with local labor market wages.

Figure 3 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 regions where the home share lies below one, demand is greater than supply, and, hence, goods are being imported. This region naturally corresponds to the situation with low world prices or low home productivity—i.e. the economy has a comparative disadvantage in producing these commodities.

Second, in the regions where the home share lies above one, supply is greater than demand, and, hence, goods are being exported. This region corresponds to high world prices or high home productivity. In other words, this is where the country has a comparative advantage and is an exporter of the commodities.

Third, there is the “table top” region in the middle, where the home share equals one. Hence, this is the region where the goods are non-traded. Exactly like the inner, non-traded region in the Ricardian model of Dornbusch, Fischer, and Samuelson (1977), the reason is trade costs. In this region, world prices and domestic productivity are not high enough for a producer to be an exporter of these commodities given trade costs. Furthermore, world prices and domestic productivity are not low enough to merit importing these commodities either. Thus, these goods are non-traded.

Finally, unlike Dornbusch, Fischer, and Samuelson (1977) or Eaton and Kortum (2002), it is important to reflect on the stochastic nature of this economy. While the stationary equilibrium of the economy leads to the stationary pattern of trade seen in Figure 3, individual islands transit between different states (world prices and domestic productivity). For example, an island may be an exporter, but given a sequence of bad productivity shocks, the island will stop exporting and maybe even become an importer of a commodity it once exported.

4.2. Trade and Wages

One can connect the pattern of trade across islands/labor markets in Figure (3) with the structure of wages in the economy. As we show in the Appendix, real wages in a market with state variable \mathbf{s} equal

$$w(\mathbf{s}) = \omega(\mathbf{s})^{\frac{1}{\theta}} \hat{\mu}(\mathbf{s})^{\frac{-1}{\theta}} z^{\frac{\theta-1}{\theta}} C^{\frac{1}{\theta}}. \quad (31)$$

Here $\omega(\mathbf{s})$ is the home share defined in (30); $\hat{\mu}(\mathbf{s}) = \frac{\mu(\mathbf{s})}{\pi(\mathbf{s})}$ is the number of labor units; z is domestic productivity; C is aggregate consumption.




Figure 3: Trade: Home Share, $\omega(s)^{\frac{1}{\theta}}$




Figure 4: Wages

Equation (31) connects the trade exposure measure in (30) with island-level wages. 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 means that wages are higher.

While this looks like the “micro-level” analog of the aggregate result of Arkolakis, Costinot, and Rodríguez-Clare (2012) it is different in one important respect: the micro-level wage response to micro-level trade exposure to trade takes the exact opposite sign.

Figure 4 illustrates these observations by plotting the logarithm of pre-tax wages by world price and home productivity so it exactly matches up with Figure 3 . As equation (31) makes clear, there is a tight correspondence between wages and the home share in Figure 3. As in Figure 3, 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 way to understand this result is as follows: wages reflect the value of the marginal product of labor. In import competing islands, 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 islands. Finally, the center region is where commodities are non-traded. Here, the gradient of wages very much mimics the increase in domestic productivity. In contrast, where goods are imported or exported, the wage gradient mimics the the change in world prices.

Again, it is important to reflect on the stochastic nature of this economy. While the stationary equilibrium of the economy results in a stationary distribution of wages, individual islands (and households living on those islands) transit between different states (world prices and domestic productivity). For example, an island may be an exporter with households receiving high wages, but given a sequence of bad productivity shocks, the island will stop exporting, and household wages will fall.

Finally, equation (31) connects with the aggregate gains from trade. Any change in aggregate trade exposure will also change in aggregate consumption, i.e. the C term. That is all workers benefit from the “aggregate gains to trade”, but the island-level incidence will vary with its trade exposure and may mitigate or completely offset the aggregate benefits from trade.

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

The preceding results relate closely to the empirical specification and evidence of Autor, Dorn, and Hanson (2013) and Acemoglu, Autor, Dorn, Hanson, and Price (2016) that link changes in trade exposure with labor market outcomes such as wages (see Section IV.B of Autor, Dorn, and Hanson (2013)). To do illustrate the connection, start with (31) and take log differences across

time yielding

$$\Delta \log w(\mathbf{s}) = \frac{1}{\theta} \Delta \log (\omega(\mathbf{s})/\hat{\mu}_h(\mathbf{s})) + \underbrace{\frac{1}{\theta} \Delta \log C}_{\gamma_t} + \underbrace{\Delta \log \left(z \frac{\theta-1}{\theta} \right)}_{\epsilon_{s,t}}, \quad (32)$$

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 aggregate consumption and (iii) the change in location-specific productivity.

Equation (32) is closely related to the empirical specification of Autor, Dorn, and Hanson (2013) (see equation (5)). Autor, Dorn, and Hanson (2013) relate various labor market outcomes at the commute zone level to commute-zone-level measures of trade exposure. Put in their terms, our theory connects changes in wages on the left hand side with trade exposure, an aggregate effect (which would be picked up by the constant/or time effect), and the error term reflects unobserved commute-zone-level productivity shocks.

Consistent with their arguments, equation (32) makes clear that an instrumental variable strategy is necessary to identify the causal effect of trade exposure on wages. Commute-zone-level productivity shocks are unobserved, but correlated with trade exposure and, thus, trade exposure could increase either because of changes in world prices or domestic productivity.

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 that the stochastic process in (14) that is assumed to be orthogonal to z .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, but correlated with world prices. This, in fact, is quite similar to the instrument proposed in Autor, Dorn, and Hanson (2013).

4.4. Labor Supply

Figure 9 plots the policy functions for labor supply in the calibrated economy. Because there are multiple dimensions of heterogeneity, this plots the policy functions for a given z and then illustrates how labor supply varies across asset holdings and trade exposure. The distribution of asset holdings (for that z) is plotted to the right. Darker colors mean that more are likely to not work, lighter colors mean that working is stronger.

Either by looking across panels or within a panel across trade exposure, households are gen-

⁷Moreover, the model makes clear that one should be concerned, in general equilibrium, that a change in domestic productivity would feed into world prices and, thus, invalidate this strategy.

erally more likely to work the better labor market conditions. For example, the colors in the medium z panel are generally lighter than those in the low z case. Or when making comparison within a z panel as trade exposure lessens (and wages improve) the colors become lighter, households are generally participating more.

With that said, there are economically interesting non-linearities that arise because of wealth effects and the ability to use migration as an insurance mechanism. For example, focus on the most trade exposed area of the low z panel. In this labor market, the wealthiest households are not participating. Like in Chang and Kim (2007), wealthy households have a high reservation wage because they can simply enjoy leisure and consume of their large stock of assets. And because the low z , high trade exposed region will have a low wage, these households do not participate.

However, participation is not monotone in wealth. For the low z island, at the top of the bottom half of the asset distribution, participation is near one hundred percent. What makes this region special is households there are close to the minimum amount that they need to migrate. Thus, these households participate in the labor market to work, save, and preserve the option of moving to a better market. Households are working to self insure, but not through asset accumulation, but through migration.

Below the bottom half of the asset distribution households stop participating. These households don't have and can not acquire the resources to migrate, the labor market is poor, thus their best option is simply to drop out of the labor force and receive leisure and home production. Home production is critical here as without it, households would prefer to work and, in turn, be able to consume.

4.5. Migration

Figure 6 illustrates how migration varies across states. Similar to the labor supply figures, this plots the policy functions for a given z and then illustrates how the moving choice varies across asset holdings and trade exposure. Again, the distribution of asset holdings is plotted to the right.

Current income, through the z and p_w shocks, play a very strong role in shaping migration. In low z islands, and in particular high trade exposed islands, the desire to move is very strong. In contrast, on medium z islands, the desire to migrate is far more modest. The key idea here is that the migration motive in our model is for insurance. That is, households undertake costly moves to escape negative labor market conditions in one location and capture favorable labor market conditions in another location.

Asset holdings, however, constrain who can move and who cannot. In the bottom of both panels, this can be seen in the white sections with no migration for all levels of trade exposure.


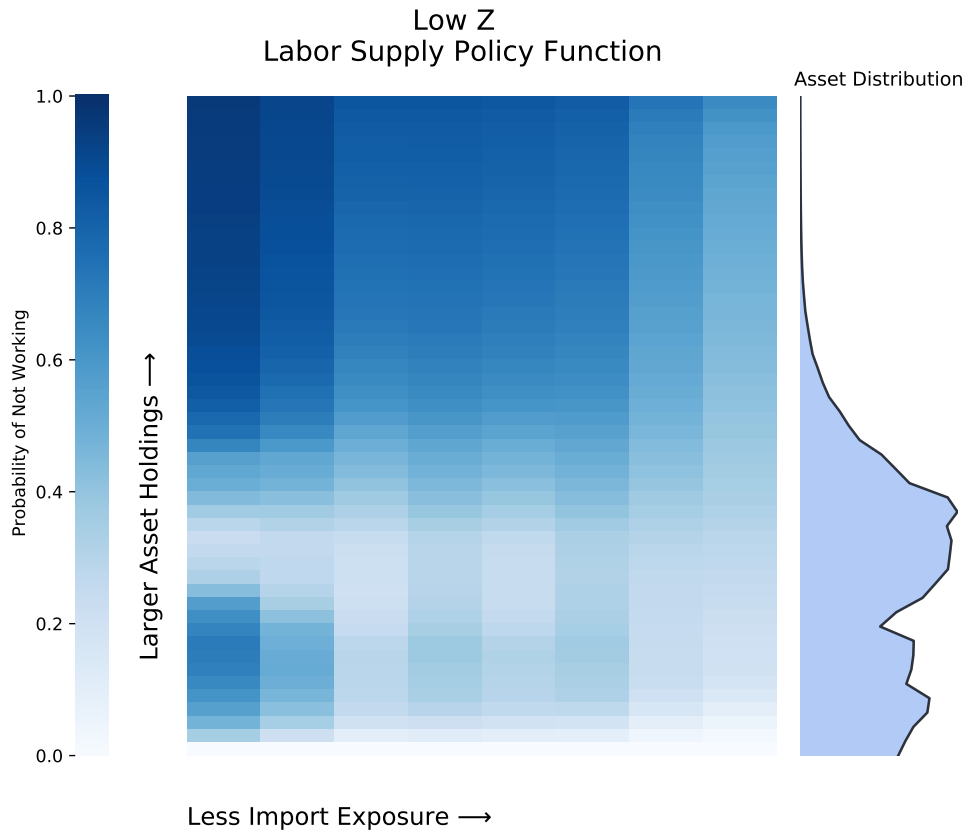


Figure 5: Labor Supply by Assets and Trade Exposure

Households in this region do not have the financial means to migrate.

Asset holdings and income shocks interact through a “double whammy” effect—those with the strongest desire to move are likely not to have the means to do so. In the low z and high trade exposed islands, the desire to out migrate is the strongest. However, because those on low z islands likely have experienced a sequence of negative shocks their asset holdings are likely insufficient to afford a move. This can be seen by noting the large mass of households with asset holdings in the region where moves are not feasible. Thus, the double whammy of wanting to move, but can’t move. In contrast, consider the medium z panel. Here the desire to move is modest even for those islands who are the most trade exposed. But the households on these islands are also much wealthier and have greater means to afford a move. In other words, these guys don’t want to move, but they could if they wanted.

These observations lead to a couple of issues that are worth pointing out. First, market incompleteness is interacting with the migration decision. From a positive stand point, this interaction is one reason why our model is able to match the low migration response in Table 2. From a normative stand point, this interaction suggests that the welfare costs of trade exposure for certain households may be quite high.

Second, this behavior is very different than in the models of trade and migration in the literature, e.g., Artuç, Chaudhuri, and McLaren (2010), Caliendo, Dvorkin, and Parro (2015). In previous work, the migration cost is purely in utility terms and...

5. Calibration

While illustrative, the previous section left open several questions. While our model will predict heterogeneous responses, how the aggregate gains from trade offsets these responses is not clear. Second, our model has several margins which households can potentially partially offset reductions in wages — borrowing, labor supply adjustments, and moving. The quantitative analysis below explores these issues.

This section outlines our calibration approach. We divide our calibration approach into essentially two steps. First, there is a subset of parameters that are determined outside of the model and based on prior evidence. Second, the remaining parameters that calibrated so that our economy replicates aggregate and micro-level facts about the US economy prior to the China shock and then the economy’s response to the China shock. The latter point makes the calibration procedure difficult and non-standard in that we are matching moments as the economy transitions and adjusts to China—this is not a calibration based on steady state to steady state responses.


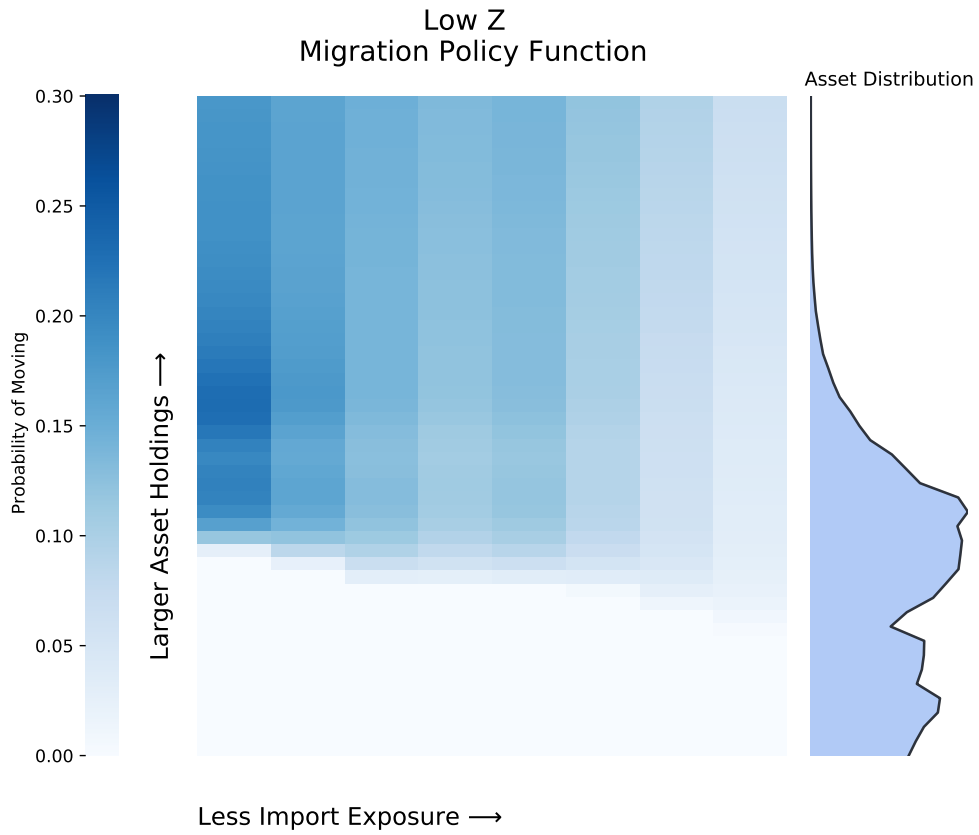


Figure 6: Migration by Assets and Trade Exposure

Table 3: Predetermined Parameter Values

Parameter	Value	Target Moment/Notes
Discount Factor, β	0.95	—
World Interest Rate, R	1.02	—
Persistence of z and p_w process	0.95	—

5.1. Predetermined Parameters

Time and Geography. The time period is set to a year. Given the time period we set discount factor equal to 0.95. This value is at the top end with those used in Krueger, Mitman, and Perri (2016). Geographically, in our model, there is an abstract notion of an island, households living on that island, and working within its local labor market. Per the discussion in Section 4.3, we want to tightly connect our model’s implications with empirical evidence of Autor, Dorn, and Hanson (2013). Thus, we will think of the empirical counterpart to an island as a Commuting Zone (see Tolbert and Sizer (1996)) and as used in Autor, Dorn, and Hanson (2013)).

Productivity and World Price Process. The productivity process in (4) and (14) leaves three parameters to be calibrated: $\{\phi, \sigma_z, \sigma_w\}$. The parameters controlling the volatility are internally calibrated and described below. The parameter controlling the persistence of the shocks is externally calibrated: we set the persistence parameter ϕ to 0.95. With that said, a key issue in this class of models is how persistent the shocks are and, more specifically for our question, the permanence of the change in comparative advantage. This is important in that it will affect how insurable or uninsurable these shocks are. We speculate that the results of Krishna and Senses (2014) and Hanson, Lind, and Muendler (2015) speak to these dynamics of comparative advantage, as well.

The final world price that we must calibrate is the gross real interest rate, R in the initial stationary equilibrium. We set this equal to 1.02 which corresponds with a two percent annual interest rate.

5.2. Calibrated Parameters

We calibrate the remaining parameter values so that moments in the initial stationary equilibrium and moments along the transition to a new stationary equilibrium match moments about the period prior to and post China’s rise. In particular, we think of the initial stationary equilibrium period as corresponding with the first time period of 1990 to 2000 used in Autor, Dorn, and Hanson (2013); the China shock period (i.e. the transition) corresponds with the second

time period or 2000-2007.

The following parameters are chosen: the disutility of work, home production, the migration cost, the borrowing constraint, pre-China shock trade cost, post-China shock trade cost, and the demand elasticity; call this parameter vector $\Theta = \{B, \bar{a}, m, \tau_{im}, \tau_{ex}, \sigma_z, \sigma_w, \tau'_{im}, h, \theta, \sigma_\nu\}$. This parameter vector is then chosen to minimize the distance between eight moments in the model and eight moments in the data. Below we describe the eight moments and the parameters they are most tightly linked with.

- **Labor force participation.** We target a labor force participation rate of 67 percent which corresponds with the average value across the period of 1990-2000 in US data. We target this value in the initial stationary equilibrium. This moment is most informative about the disutility of work, B .
- **Fraction of households with households have zero or negative wealth.** Krueger, Mitman, and Perri (2016) report from the Survey of Consumer Finances that approximately 40 percent of households have zero or negative wealth. We target this value in the initial stationary equilibrium. This moment is most informative about the borrowing constraint, \bar{a} .
- **Migration rate.** We use the the IRS migration data which uses the address and reported income on individual tax filings 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 commuting zone at a yearly frequency. We target this value in the initial stationary equilibrium. This moment is most informative about the migration cost, m
- **Trade volumes pre China's rise.** In the initial stationary equilibrium, we target an initial import to GDP ratio of thirteen percent. This latter value is consistent with the degree of openness seen in Figure 1 in the late 1990s prior to the acceleration of Chinese trade. This moment is most informative about the initial import and export trade cost (which we assume take on the same value initially).
- **Standard deviation of growth rates in commute zone level labor earnings.** In the initial stationary equilibrium, we target standard deviation of growth rates in commute zone level labor earnings of 6.5 percent. This value is measured in the data by using the decennial Census data from Autor, Dorn, and Hanson (2013) and focusing on the period between 1990 and 2000. One short-cut that we take is that given in the closed economy version of this model and given a ρ and θ , we can determine, in closed form, the volatility of growth in labor earnings simply by picking σ_z . Thus, we directly calibrate this value prior to computing the stationary equilibrium.

- **Trade volumes post China’s rise.** Along the transition path, we target an import to GDP ratio of seventeen percent seven years after the change in policy. This latter value is consistent with the degree of openness seen in Figure 1 in 2007. This moment is most informative about the final trade cost.
- **Aggregate gains from trade.** Given the long-run increase in trade from the bullet above, we can use the formula of Arkolakis, Costinot, and Rodríguez-Clare (2012) to impute the aggregate gain in output associated with a trade elasticity of four (Simonovska and Waugh (2014)). We then ensure that our model has the aggregate responses in output as the Arkolakis, Costinot, and Rodríguez-Clare (2012) would formula implies. While the Arkolakis, Costinot, and Rodríguez-Clare (2012) does formula does not hold in our model, we are using their formula as a model diagnostic to ensuring that the level of the gains from trade are the same as in a standard, representative agent trade model.

This moment is most informative about the volatility of world prices, σ_w . This may seem odd, but read on. The insight here is that variation in world prices determines how elastic aggregate trade flows are to a change in trade frictions. For example, if world prices are very dispersed, then large changes in trade frictions are necessary to generate large changes in trade flows. In contrast, if there is little variation in world prices, then small changes in trade frictions will generate large changes in trade flows. This insight is analogous to the behavior of the Eaton and Kortum (2002) model where the extent of technology heterogeneity controls how elastic trade flows are to changes in trade costs.

- **Wage and labor force participation elasticity from Autor, Dorn, and Hanson (2013).** Specifically, we target the elasticities described in Table 1. These moments are most informative about θ and w_h . The logic about the relationship between the Autor, Dorn, and Hanson’s (2013) wage elasticity result and θ and follows from the discussion in 4, that is the elasticity of wages to changes in trade exposure is related to demand elasticity.

The home production parameter w_h is related to the labor force participation elasticity of Autor, Dorn, and Hanson (2013). The idea is that home production controls the opportunity cost of working in the market and, thus, it controls how households substitute in and out of the labor force in response to shocks. Given that Autor, Dorn, and Hanson (2013) have identified the elasticity of labor supply to a trade shock, this will inform our home production parameter.

- **Migration elasticity from Greenland, Lopresti, and McHenry (2017).** We target the Greenland, Lopresti, and McHenry (2017) migration elasticity described in Table 2. This moments is most informative about the scale of the preference shock σ_ν . For example, if the variance of the preference shocks are large, then large movements in the value of moving relative to staying are required to induce households to move. On the other hand, if

the variance of the shocks are small, then small changes in the relative value of moving to staying will induce large numbers of households to move. Hence, the migration elasticity of Greenland, Lopresti, and McHenry (2017) helps pin this parameter down.

5.3. Implementing the China Shock

Main exercise focuses on a change in the the ability to import goods, i.e., a reduction in τ_{im} . Mechanically, we implement the change in the following way: In year one, the a new path of trade costs are announced and implemented. The new, announced path of new trade costs those that linearly decrease from τ_{im} to τ'_{im} over seven years. The idea here is to generate the gradual rise in trade as in the data and to mimic the narrative around the change trade policy with China's accession to the WTO and granting of permanent normal trade relations by the United States. As mentioned above, the exact level of τ'_{im} is chosen so that after seven years from announcement, imports to GDP resemble that seen in the US in 2007.

5.4. The Procedure

To implement the calibration, we work through the following steps. Much of this is done in a simultaneous manner, but we describe the core steps to facilitate how we map our model into the data.

Step 1. Guess the parameters $\sigma_w, \theta, w_h, \sigma_v$.

Step 2. We pick the parameters $\{B, \bar{a}, m, \tau_{im}, \tau_{ex}, \sigma_z\}$ so that the initial stationary equilibrium (pre-China shock period) replicates the labor force participation rate, migration rate, net worth of households, volatility of labor earnings, and the initial volume of trade.

Step 3. We pick the parameter τ'_{im} and then compute the new stationary equilibrium. We compute the transition path. That is starting from the initial stationary distribution we change the trade friction and compute the transition path to the new stationary distribution. We then check that seven years after the change, the volume of trade equals seventeen percent. We compute the long run aggregate gain in output and compare it to the value predicted by the Arkolakis, Costinot, and Rodríguez-Clare (2012) formula.

Step 4. Given the transition path, we simulate data sets analogous to those in Autor, Dorn, and Hanson (2013) and estimate the wage, labor force participation, and migration elasticities with respect to the trade exposure metric. In particular, we constructing data analogs in our model as they are constructed in Autor, Dorn, and Hanson (2013) and estimate (2) on model generated data with a time effect. Simulated trade flows for another country (same sequence of $p_w s$, different sequence of $z s$) is the instrument.

Given the difference between the model and data elasticities, return to **Step 1.** and update the parameters $\sigma_w, \theta, w_h, \sigma_v$.

Table 4: Calibration: Parameters and Aggregate Moments

Parameter	Value	Target	Data	Model
Disutility of work, B	1.05	Aggregate participation rate	66	66
Migration Cost, m	1.75	CMZ. migration rate	3	3
Borrowing Limit, $-\bar{a}$	0.84	% Households with ≤ 0 net worth	40	40
Pre-China Trade Cost (τ_{ex}, τ_{im})	1.16	1990s Imports/GDP	13	13
Post-China Trade Cost (τ'_{im})	1.37	2007 Imports/GDP	16.2	15.4
Std. Dv. of z (σ_z)	0.032	Std. Dev. in CMZ earnings	7	9
Std. Dv. of p_w (σ_w)	$1.64 \times \sigma_z$	Predicted ACR Gains	1.6	1.8

Note: All moments are reported in percent. Migration cost and borrowing limit parameters are reported as a fraction of output per worker.

5.5. Calibration Results

Table 4 reports the calibration results parameters and the associated aggregate moments. Overall, the model is very flexible and is easily able to exactly fit most moments. The only moments in which there is some minor difficulty are those that depend upon the transition path. At the current stage, this appears to be largely due to time constraints and the solver needing more time to find the exact parameterizations.

Figure 7 plots the model's prediction for imports. Similar to the data, the model generates rising trade exposure over the early 2000s. As we discuss below, one feature of the model is that once the path of trade costs are announced, consumption contracts by a small amount. Since demand for all goods (including imported) are contracting, imports fall by a slight amount in the first couple of years and then slowly rise.

Table 5 reports the micro-moments, model prediction/fit, and the resulting parameter values. The first row reports the estimates from Autor, Dorn, and Hanson (2013) and Greenland, Lopresti, and McHenry (2017), the second row reports the results from the calibrated model. There are a couple of points to mention. First, our model is able to come in the close vicinity the micro-moments. All point estimates from the model lie within the 95-5 confidence bands of the data. Our interpretation of this (along with extensive sensitivity analysis) is that results of Autor, Dorn, and Hanson (2013) are informative about structural parameters that govern the




Figure 7: Rising Trade Exposure: Model and Data

underlying data generating process.

With that said, the model is experiencing difficulty in exactly nailing these moments. The difficulty appears to be in the strong substitution pattern between not working and migrating. In the model, households are unseizing these two margins to avoid the negative consequences of the shock. And when one of the margins looks more or less favorable, the model implied elasticities change in different ways.

A final point of note is that the estimated demand elasticity, θ , is relatively large. The reason follows from the discussion in Section 4: The elasticity of wages to changes in trade exposure is related to one over the demand elasticity. So if the demand elasticity is high, then changes in trade exposure will weakly pass-through into wages. Hence, what this is telling us vis-a-vi the Autor, Dorn, and Hanson (2013) evidence, is that pass-through was not that strong.

6. The China Trade Shock

This section analyzes the affects of a change in the ability to import goods on the economy at the macro and micro level.

Table 5: Calibration: Micro Moments and Parameters

	Δ Labor Earnings	Δ NILF	GLM Δ Population
Data	-4.30 [-6.62, -2.00]	1.11 [0.52, 1.72]	-1.43 [-3.33, 0.48]
Model	-4.10	1.24	-1.92
	Demand elasticity θ	Home production w_h	Moving shock variance σ_ν
Parameter Values	9.53	0.22	0.96

Note: Values in brackets report 95-5 confidence intervals.

6.1. The China Shock at the Micro level

As a first step to understanding the aggregate and welfare effects of the shock, we take a deeper look “under the hood”, we look at how real wages, labor supply, and consumption respond at the island level.

Real Wages. Figure 8 plots the log change in real wages across islands six years after the shock. The islands are ordered from the left to right by initial trade exposure with those to the left (and red colored) being import exposed. The pattern of change looks a bit bizarre, but with a little effort it makes complete sense and leads to a couple surprising insights.

First consider the exporting region (blue, positive bars to the right). These islands experience large gains in real wages—nearly four percent. Where does the gain come from and why is it same across exporting islands? The reason that real wages increase is through increases in real purchasing power, i.e. the decline in the CES price index. And this effect is common across islands, thus all exporting regions receive the same proportional gain. To see this point, recall that wages in this region are such that

$$w(\mathbf{s}) = \frac{p_w z}{\tau_{ex} P}.$$

For a given island type, neither the world price, the z , or the exporting trade cost changes. Thus, only the CES price index P changes and these is why those in the exporting region gain.

Now consider the importing region (red, positive bars to the right). This region experiences large declines in real wages, nearly two percent. Again, why the loss and why is it same across import exposed islands? The loss is from the decline in world prices (inclusive of trade costs)




Figure 8: Real Wages

net of the gain in purchasing power. Combining the iceberg trade cost assumption and the decline in the CES price index implies that the change is the same across import-exposed islands. To see this point, wages are:

$$w(\mathbf{s}) = \frac{\tau_{im} p_w z}{P}.$$

Two forces work in the opposite direction: one is the decline in the trade cost (effective world prices) and the second is the rise in purchasing power (decline in P). Given the discussion above, the fall in the CES price index can be inferred by the effects on the exporting region, a four percent decline. This implies that trade costs fall by six percent and, on net, households in these regions experience a two percentage point drop in real wages.

Finally, consider the non-traded region (black, in the middle). Here wages are relatively unchanged. This may seem strange—we just argued that real purchasing power is increasing and these goods are not exposed to trade. Why are real wages not increasing? The issue is that the reduction in the CES price index from trade leads consumers to substitute away from these non-traded goods. And since labor is quasi-fixed, this necessitates a fall in the price of the lo-

cally produced good to maintain demand and clear the labor market. Thus, the islands nominal wages fall by the nearly the same exact amount as the improvement in purchasing power.

To mechanically see this, examine the labor market clearing condition on islands where the good is non-traded and the price is domestically determined:

$$\left(\frac{p(\mathbf{s})}{P}\right)^{-\theta} C = z(\mu(\mathbf{s})/\pi(\mathbf{s}))$$

Abstract from any change in aggregate demand C and treat labor supply as fixed. In this case, $p(\mathbf{s})$ must **decrease** to clear the goods market by the exact same amount as the decrease in the aggregate price index. And since wages in this region are $w(\mathbf{s}) = \frac{p(\mathbf{s})z}{P}$, the total effect on wages is no or a small change because of aggregate demand effects or changes in labor supply. Finally, note that this mechanism does not operate in the exporting region because excess product is exported; there is no need for demand to match up with supply in these islands.

Several, more points regarding this. The surprising part is that the only segment of the population that actually gains from an import shock is the exporting region. The standard intuition typically is that the non-traded region will gain as well through increased purchasing power. With frictional labor markets, this logic is incomplete and implies that this benefit will not accrue towards them. This also implies that the CES price index is changing by a very large margin and more than the output response (as would be the case in Arkolakis, Costinot, and Rodríguez-Clare (2012)). Prices in both the imported and non-traded region are falling which covers about 75 percent of the population.

Finally, these magnitudes depend very much on the calibrated parameter θ . Recall, that the evidence of Autor, Dorn, and Hanson (2013) pushed our estimation routine to find a very high value for the demand elasticity, around ten. As we argued above, a high demand elasticity implies that changes in trade exposure weakly pass-through into wages. This is why wages in this region only fall by two percent. However, a high demand elasticity also implies that substitution effects are very strong and thus (i) real wages for those not exposed to trade change very little and (ii) all the gains from trade accrue to those who are export exposed. That is the Autor, Dorn, and Hanson (2013) evidence implies that while the direct losses are small, the gains from trade are narrowly concentrated.

How Households Respond. Before discussing the responses of labor supply or consumption, it's important to understand the nature of the shock. After the announcement, the change in trade costs is gradual and perfectly foreseeable. Second, the local productivity shock and world price process are highly persistent, thus a households current location likely to be the same in the next couple of years. Together, this implies households see the bad (or good) outcomes coming their way in Figure 8 and react appropriately.

For households that are initially import exposed, they are likely to see a drop in real wages. In other words, they are facing a negative, back-loaded income shock, i.e., stuff is good today, but foreseeably bad in the future. Standard intertemporal smoothing logic implies that these households will want to do things to transfer income from today into the future. In contrast, for those that are initially export exposed, they see increases in real wages coming and will want to do things to transfer this gain in future income into benefits today.

Labor Supply. Figure 9 illustrates the change in labor force participation across islands/labor markets after the shock. To orient yourself, the x-axis (lower left) reports the periods after the shock. The y-axis (lower right) reports trade exposure. The vertical, z-axis reports the island level change in labor force participation.

Those islands that are initially import (red and to the left on the y-axis) exposed—and likely to become increasingly so—immediately increase their labor supply by a substantial margin, almost four percentage points for the most trade exposed. Households on these islands see the negative effects coming and labor supply to fulfill their intertemporal savings motives, i.e., work today to build up savings today to facilitate the smoothing out of the future negative income shock in the future. As the shock progresses, this desire partially dissipates and labor force participation is one percent higher for import exposed islands.

In contrast, those that are export exposed (blue and to the right on the y-axis) decrease their labor supply. Households in these locations see the gains from trade coming their way and their motives are to transfer some of this new found wealth into benefits today. One way they do this is by substituting into leisure today. And this effect dissipates as households eventually increase their labor supply as the benefits from participating in the labor market become increasingly generous. That is because of the rise in wages seen in Figure 8, their opportunity cost of not working increases and this induces participation.

Consumption/Savings. Figure 10 illustrates the change in island-level savings rates after the shock. Everything here is the same as Figure 9, except now the vertical, z-axis reports the change in the island level savings rate. The savings rate is simply the one minus island level consumption relative to island level labor income (wages multiplied by labor supply). This is crude and hides a lot of heterogeneity, but provides a picture of the change in savings behavior.

In many ways, this figure mimics the patterns in labor supply. Import exposed islands (red and to the left on the y-axis) are facing a negative, back-loaded income shock and want to find way to transfer income today to mitigate the shock in the future. They initially increase savings by between six and eight percent for those that are most exposed. In contrast, export exposed islands who see increases in real wages coming increase their consumption today and decrease their savings to transfer future income into benefits today.

Long Run Labor Supply and Consumption/Savings. Figures 9 and 10 also illustrate that these




Figure 9: Model: Change in Labor Supply, Across Islands




Figure 10: Model: Change in Savings, Across Islands

patterns reverse as the economy approaches the new stationary equilibrium. Export exposed islands tend to participate more and save more; import exposed islands tend to participate less and save less. Moreover, it's the islands in the middle (which are non-traded regions) which work and save at the highest rates. This behavior is about precautionary motives. An important mechanism in the new stationary equilibrium is that a reduction in trade costs increases the amount of uninsurable income risk in the model. The idea here is that while many households are not directly exposed, a reduction in trade costs increases the likelihood that they will eventually become trade exposed and, thus trade leads to an increase in income risk. This force increases precautionary savings motives (see, e.g., Huggett and Ospina (2001)) and, in turn, leads to an increase in labor supply to acquire those savings and a higher savings all together.

6.2. From Micro to Macro

How do the micro-level outcomes aggregate? The top panel of Figure 11 illustrates the behavior of aggregate consumption, labor supply, and output both before and after the trade shock.

Focus on labor supply first. Figure 11 illustrates that there is temporarily a small decrease in labor supply and output during the announcement period. Recall from Figure 9 is that those import exposed increase their labor supply, those who are export exposed decrease their labor supply. The latter more than offsets the former and thus overall leads to a small decline in aggregate labor supply.

Very quickly, however, this pattern reverses and is a large increase in labor supply—about one percentage point. Along the path, this is because those import exposed are slowly decreasing their participation and the export exposed are increasing their participation as they smooth out the shock. The medium/long-term increase in labor supply arises from precautionary savings motives of several sub-groups. About sixty percent of the one percentage point increase is due increased participation on islands that are export exposed; the remaining forty percent is from islands that are interior and not trade exposed.

All of these effects translate into a temporary fall and then sustained increase in output. After six years, the output gain is about 2 percentage points larger after seven years and exceeds that of labor supply by one percentage point through aggregate gains from trade.

The aggregate consumption response is quite similar to the labor supply response. Initially, aggregate consumption falls by a small amount and then rises almost one for one with labor supply. Like the labor supply response, this is all being driven by intertemporal motives of households at the micro-level to smooth out the shock.

These observations have the implications that that the trade deficit is **decreasing**. Figure 12 illustrates this point by showing that after an initial decrease, there is an improvement of the trade deficit by nearly one percentage point of GDP, seven years after the shock. The improve-





Figure 11: Model: Consumption, Output, Labor Supply



41
Figure 12: Trade Defect

ment in the trade deficit and the consumption response and are intimately related. A traditional open economy macro perspective is to view this result through the identity that output minus consumption (aggregate savings) equals trade deficit. In response to the trade shock, output increases, consumption not as much thus the trade deficit must improve.

It is also consistent with a traditional international trade perspective, i.e., that the trade deficit equals exports minus imports. First, a muted consumption response implies that demand for all goods (foreign and those domestically produced) is muted. Thus, imports do not increase as much as they would otherwise since aggregate demand is muted. Second, because the relative price of domestically produced goods is rising, relative demand for non-traded and exported goods falls. Given that production is increasing through the increase in labor supply, exports must increase as product not consumed domestically goes abroad. Thus, exports increase more than imports, and the trade deficit improves.

These findings are important relative to the “popular narrative” discussed in the introduction. First, the aggregate increase in labor supply is consistent with Autor, Dorn, and Hanson (2013) finding that increases in trade exposure lead to reductions in labor force participation. By construction, our model matches their evidence as shown in Table 5. However, as mentioned before, this evidence has little content when speaking about aggregate effects. In this case in particular, the aggregate response is in the exact opposite direction of a naive interpretation of Autor, Dorn, and Hanson (2013) labor force participation evidence.

A second issue relates to the behavior of the trade deficit. The popular narrative suggests that if it becomes easier to import, then this leads to an expansion in the trade deficit. In our model, the trade deficit decreases—a reduction in the cost of importing, lead exports to expand by more than imports. This suggests that attributing an expansion of the US trade deficit in the early 2000s to increased trade with China is wrong and it fits well with the narrative that global financial forces advocated by Bernanke (2005) are a more likely culprit.

6.3. Welfare

How does this all add up in terms of welfare? Let’s first take stock. Section 6.2 alluded to the fact that those who could engaged in pre-cautionary measures in reaction to the shock. Section 6.1, in general, confirm the idea that those who were import exposed suffer losses in earnings, but the effects on others—non-traded and exporting regions—are far more nuanced.

How We Measure Welfare. First, we measure welfare and report numbers at the island level. So the idea here is take the welfare of the “average” household on island s and compute the lifetime consumption equivalent that would make that “average” household indifferent between living in the world prior to the shock and immediately after.

To formalize this, define expected lifetime utility on island in the initial stationary equilibrium

as

$$W(\mathbf{s}; \chi(\mathbf{s})) = \int_a \int_\nu V(a, \mathbf{s}, \nu; \chi(\mathbf{s})) \lambda(\mathbf{s}, a, \nu) d\nu da, \quad (33)$$

where χ is the consumption equivalent parameter. Welfare immediately after the shock is

$$W'(\mathbf{s}) = \int_a \int_\nu V'(a, \mathbf{s}, \nu) \lambda'(\mathbf{s}, a, \nu) d\nu da. \quad (34)$$

Thus, our measure of welfare is the value of $\chi(\mathbf{s})$ which must make $W(\mathbf{s}; \chi(\mathbf{s})) = W'(\mathbf{s})$. One point to notice that this measure hides the heterogeneity of households within an islands. Because each island has its own distribution of assets, there is variation in how much a household on that island wins (or loses). By integrating over asset holdings within an island, we are not exploring this variation.

Welfare Effects of the Shock Table 6 summarizes our findings in addition to reporting the change in real wages; Figure 13 illustrates the full range of outcomes by plotting the welfare gains for each island type.

On average, the gains are not out-of-line with what a simple calculation from Arkolakis, Costinot, and Rodríguez-Clare (2012) would suggest. A little less than one percentage point which more or less matches up with the average change in real wages which is a bit more than one percent.

Table 6: Welfare

		Welfare (Baseline)	Δ Real Wages
Average		0.94	1.06
Initial Exposure	Import Exposed	0.19 [0.09]	-2.19 [0.09]
	Non-Traded	0.75 [0.68]	0.34 [0.68]
	Export Exposed	1.64 [0.25]	3.99 [0.25]

Note: Values are lifetime consumption equivalents; values in brackets report the share of the population in that category.

The interesting outcomes are with respect to the distributional consequences. First, as Figure 13 shows, those who are initially import exposed experience welfare losses, though some ex-




Figure 13: Welfare Across Islands

perience very small, but positive gains. A few losses, minor gains, stands in sharp contrast to the harsh effects trade has on these labor market outcomes—wages falling by two percentage points over the first five years of the shock. As alluded to above, this suggests that these households are actively consumption smoothing to mitigate the effects of the shock.

Surprisingly those that are in non-traded regions experience substantial welfare gains, about 0.75 percent increase in welfare. Again, this stands in sharp contrast to the non-effects trade had on these labor market with wages essentially unchanged. The issue here is that these islands are reaping dynamic gains from trade. While their current island is not trade exposed, given the stochastic nature of comparative advantage they will be export exposed at some point and reap the benefits. In other words, while I might not benefit today, my future self, children, etc. benefit from the opportunities that trade.

Finally, those in the exporting regions gain the most, a bit less than two percent. Recall several points discussed above. First, these gains are purely through changes in the price index and, hence, a rise in real wages. Second, the change in real wages does not map one for one into welfare for several reasons. These are transitory gains relative to the average and, thus, a portion

of this is simply saved for possible bad times in the future. Moreover, relative bad times in the future will come and, thus, their invariable position is merely transitory.

Several open questions remain. As the discussion above suggested there are two forces shaping these gains and making them different from the labor market effects. One force is that households have mechanisms to self-insure and mitigate bad outcomes. A second force is that these gains reflect future changes in the comparative advantage and trade exposure. Disentangling these effects is on the plate for future outcomes.

7. Conclusion

Lots more work todo!

References

- ACEMOGLU, D., D. AUTOR, D. DORN, G. H. HANSON, AND B. PRICE (2016): "Import competition and the great US employment sag of the 2000s," *Journal of Labor Economics*, 34(S1), S141–S198.
- AIYAGARI, S. R. (1994): "Uninsured idiosyncratic risk and aggregate saving," *The Quarterly Journal of Economics*, 109(3), 659–684.
- 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.
- 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.
- BERNANKE, B. (2005): "The global saving glut and the US current account deficit," Discussion paper.
- 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.
- COŞAR, A. K., N. GUNER, AND J. TYBOUT (2016): "Firm dynamics, job turnover, and wage distributions in an open economy," *The American Economic Review*, 106(3), 625–663.
- DIX-CARNEIRO, R. (2014): "Trade liberalization and labor market dynamics," *Econometrica*, 82(3), 825–885.
- DIXIT, A. (1987): "Trade and insurance with moral hazard," *Journal of international economics*, 23(3-4), 201–220.
- (1989a): "Trade and Insurance with Adverse Selection," *The Review of Economic Studies*, 56(2), 235–247.
- (1989b): "Trade and Insurance with Imperfectly Observed Outcomes," *The Quarterly Journal of Economics*, 104(1), 195–203.
- 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 G. M. GROSSMAN (1985): "Tariffs as Insurance: Optimal Commercial Policy When Domestic Markets Are Incomplete," *The Canadian Journal of Economics / Revue canadienne d'Economique*, 18(2), 258–272.
- EATON, J., AND S. KORTUM (2002): "Technology, geography, and trade," *Econometrica*, 70(5), 1741–1779.
- GREENLAND, A., J. LOPRESTI, AND P. MCHENRY (2017): "Import competition and internal migration," .
- HANSON, G. H., N. LIND, AND M.-A. MUENDLER (2015): "The dynamics of comparative advantage," 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.
- HUGGETT, M., AND S. OSPINA (2001): "Aggregate precautionary savings: when is the third derivative irrelevant?," *Journal of Monetary Economics*, 48(2), 373–396.
- 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.
- 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.
- KRUSELL, P., T. MUKOYAMA, AND A. ŞAHIN (2010): "Labour-market matching with precautionary savings and aggregate fluctuations," *The Review of Economic Studies*, 77(4), 1477–1507.
- 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 (2018): "Redistributing the Gains From Trade Through Progressive Taxation," .
- MIRANDA, M. J., AND P. L. FACKLER (2004): *Applied computational economics and finance*.
- MORTEN, M. (2016): "Temporary migration and endogenous risk sharing in village india," Discussion paper, National Bureau of Economic Research.
- NEWBERY, D. M., AND J. E. STIGLITZ (1984): "Pareto inferior trade," *The Review of Economic Studies*, 51(1), 1–12.

- ROGERSON, R. (1988): "Indivisible labor, lotteries and equilibrium," *Journal of monetary Economics*, 21(1), 3–16.
- SIMONOVSKA, I., AND M. E. WAUGH (2014): "The Elasticity of Trade: Estimates and Evidence," *Journal of International Economics*, 92(1), 34–50.
- TOLBERT, C. M., AND M. SIZER (1996): "US commuting zones and labor market areas: A 1990 update," .

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. Note that in all of this derivation, we normalize the price of the final good to one.

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. The latter equals aggregate payments to labor in the production of all intermediate goods.

$$Y = \int_{\mathbf{s}} w(\mathbf{s})\mu(\mathbf{s}) \quad (35)$$

Examining (??) and (35) allows us to connect aggregate income with consumption. Specifically, by integrating over the consumers budget constraint (ignoring home production), then noting how total value added must equal wages, one can substitute (35) into (??), then we arrive at the following:

$$Y = C - RA + \mathcal{A}' + \int_a \int_{\mathbf{s}} m_{\mathcal{L}_m}(\mathbf{s}, a)\lambda(\mathbf{s}, a) \quad (36)$$

so aggregate income equals consumption (private and public) minus (i) returns on assets (ii) new purchases of assets (iii) plus moving costs.

This basically says that income/production must equal consumption net of income not associated with production (i.e. returns on assets) 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

$$Y = \int_{\mathbf{s}} p(\mathbf{s})z\mu(\mathbf{s}) \quad (37)$$

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 (19), imports in (21) and exports in (23) and

the equating final demand with consumption we have

$$Y = C + \int_{\mathbf{s}} p(\mathbf{s}) \text{exports}(\mathbf{s}) - \int_{\mathbf{s}} p(\mathbf{s}) \text{imports}(\mathbf{s}). \quad (38)$$

Or GDP equals consumption (market) 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 (36) and (38) we get the following relationship

$$\begin{aligned} Y - C &= \int_{\mathbf{s}} p(\mathbf{s}) \text{exports}(\mathbf{s}) - \int_{\mathbf{s}} p(\mathbf{s}) \text{imports}(\mathbf{s}), \\ &= -r\mathcal{A} + (\mathcal{A}' - \mathcal{A}) + \int_a \int_{\mathbf{s}} ml_m(\mathbf{s}, a) \lambda(\mathbf{s}, a), \end{aligned} \quad (39)$$

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 moving costs). 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) moving costs would show up as the “balancing item.”

To see this, consider the special case where moving costs are zero, no home production, and tariffs are zero. Then we have the relationship

$$Y - C = \int_{\mathbf{s}} p(\mathbf{s}) \text{exports}(\mathbf{s}) - \int_{\mathbf{s}} p(\mathbf{s}) \text{imports}(\mathbf{s}) = -r\mathcal{A} + (\mathcal{A}' - \mathcal{A}). \quad (40)$$

Here if exports are greater than imports, then this implies that the households 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, but that trade may be imbalanced.

B. Computational Appendix

This section describes our computational approach. Related materials are posted at https://github.com/mwaugh0328/redistributing_gains_from_trade which retains code for our companion paper in Lyon and Waugh (2018). The code is presented in two different languages: Matlab and Julia. The implementation in both languages follows the same core steps, but the details are slightly different. The discussion below follows the Matlab code. For a detailed explanation of how the Julia code works, see the file `julia/README_Julia.md`.

2.1. Computing and Solving the Stationary Equilibrium

Below we describe how to compute and solve the stationary equilibrium of the model.

- We approximate the continuous asset, productivity, and world price states by discretization. The asset space follows a non-uniform grid with grid points clustered near the borrowing constraint. The number of grid points was set to 50; the results are not sensitive to increases in this number. We use the method of Rouwenhorst to discretize the productivity and world price process. We use 10 states for productivity and world prices each, thus there are 100 different states \mathbf{s} .
- Guess a proposed price function $\hat{p}(\mathbf{s})$.
- Compute wages.
- Solve the households problem in (9). This is performed using value function iteration.

One technique we use to facilitate finding a solution to the equilibrium is to “smooth” out the discrete choice problem. We do this by assuming that there are additive logistically distributed preference shocks with parameter b_{smth} and these preference shocks are independently distributed across each choice. These enter into the choice problem in (9) by adding onto each option. What this gives rise to is, for each asset holdings state and state \mathbf{s} , there will be a non-zero mass of households choosing all options. The probabilities take the familiar logit form. We tune the parameter b_{smth} to ensure that it is small and not affecting the economics of the problem, but at the same time ensure that we find a solution.

Smoothing in this manner is important as it facilitates the use derivative based solvers in finding an equilibrium $p(\mathbf{s})$. This in turn results in a dramatic speed-up in the computation of an equilibrium. See, e.g., Morten (2016) who employs a similar approach.

- Given the policy functions associated with the solution to the problem in (9), we compute the stationary distribution over assets and states \mathbf{s} , i.e. $\lambda(a, \mathbf{s})$ which is of size 50 (for each

asset state) and 100 for each state \mathbf{s} . This process is sped up using sparse matrices, see the code `island_invariant.m` for details.

- Given the stationary distribution, we can compute excess demand functions for all islands \mathbf{s} which also must respect the inequalities implied by (20) and (22). These conditions imply that the problem of finding a price function consistent with a stationary equilibrium can be represented as a mixed complementarity problem (see, e.g., Miranda and Fackler (2004)). To smooth out the nondifferentiability issues with the complementarity problem we pass the excess demand functions through Fischer's function. Again, see Miranda and Fackler (2004).
- Update $\hat{p}(\mathbf{s})$ and proceed until convergence criteria are met.

In solving for the equilibrium, we employed derivative based solvers. One solver that we found much success with is the `c05qc` solver from Numerical Algorithms Group. MATLAB's `fsolve` with central finite differences performed well too.

Part of our calibration approach employs the following technique. Let Θ be the parameter vector we chose to match some moments in, say the initial stationary equilibrium. We then *jointly* solved for $\{p(\mathbf{s}), \Theta\}$ in one step. That is we asked the algorithm described above to find a price vector and set of parameters such that (i) equilibrium conditions are satisfied and (ii) model implied moments match our target empirical moments. This avoided a more standard, but time consuming approach of guessing a Θ , solving for an equilibrium, updating Θ , etc. Extensive sensitivity analysis found no issues surrounding multiplicity of equilibrium.

2.2. Computing the Transition Path

To be completed soon.