

Quantifying the Losses from International Trade

Spencer G. Lyon
New York University

Michael E. Waugh
New York University and NBER

March 2019

ABSTRACT

Did trade with China harm the US economy in the 2000s? A popular narrative suggests that the rapid rise in Chinese import penetration lead to an expanding trade deficit and negative impacts on wages and employment within narrowly defined labor markets. We provide an alternative interpretation of this evidence by developing a dynamic, standard incomplete market model with Ricardian trade and frictional labor markets and calibrated to match local-labor-market evidence. Despite being consistent with the evidence of Autor et al. (2013), rising trade exposure induces a boom: a increase in GDP and employment, a modest increase in consumption, and an improving trade deficit. While some labor markets experience large losses in labor income; in welfare terms, very few actually lose as many are able to smooth out the shock.

Email: spencer.lyon@stern.nyu.edu, mwaugh@stern.nyu.edu. We would like to thank our discussants Martin Beraja, Jonathan Heathcote, Kerem Cosar 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, Federal Reserve Bank of Philadelphia International Trade Workshop.

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 suggests 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 seven percentage points—Chinese imports accounted for all of this growth. During this same time period, the US trade deficit as a fraction of GDP expanded by two percentage points. Often these facts are linked with the suggestion that rising import penetration caused the deterioration of the trade deficit as Chinese trade did not arrive with equivalent export opportunities. Moreover, these forces led to bad labor market outcomes. The evidence in Autor et al. (2013) supports this interpretation: using 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 evaluates this narrative through the lens of a dynamic, standard incomplete market model with Ricardian trade and frictional labor markets which is disciplined by the local-labor-market evidence. We use this setting as a laboratory to ask questions regarding the aggregate effects and the heterogeneity in the gains or losses from the China Shock.

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 “on their own” as they face technology and trade related shocks. 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. 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, households have two additional margins to mitigate shocks. First, we allow households to migrate and move to alternative labor mar-

kets. This margin has been an important dimension in models of trade and labor market dynamics (see, e.g., Artuç et al. (2010), Caliendo et al. (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 that our model provides a structural interpretation of the local labor market regressions in Autor et al. (2013). The CES demand system for island-level goods yields a log-linear link between island level wages and the production of an island's good relative to national consumption—the good level “home expenditure share.” This measure is *the* summary statistic for the exposure of a labor market to trade and is directly related to Acemoglu et al. (2016) IP measure and Autor et al.'s (2013) IPW measure. Most important, the estimate of how trade exposure passes through to wages is directly related to the elasticity of substitution in the CES demand system.

Our model also provides a motivation and rationalization for Autor et al.'s (2013) instrumental variable strategy. In our model, unobserved to the econometrician is the local island level productivity shock. Through Ricardian comparative advantage, this shock influences how exposed a labor market is to trade and, thus, its omission biases the OLS projection of labor market outcomes on trade exposure. Under a small open economy assumption, a valid instrument is another country's imports—it's orthogonal to the local productivity shock, but correlated with trade exposure through world prices. This is essentially the same instrument proposed in Autor et al. (2013).

These insights provide the foundation for the second step of our analysis: the calibration of our model. Our approach is to discipline key parameter values with the labor market evidence from Autor et al. (2013) and migration evidence from Greenland et al. (2017). This done by simulating our model and running the same, instrumented regressions used in the data on model outcomes. Per our arguments above, this 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 labor market participation; and a household's preference to move which influences how trade shocks affect migration behavior.

A difficulty, but unique feature of our calibration approach is the focus on outcomes along transition paths. We model the shock behind the expansion of Chinese trade as a change in the cost to import goods. Given this change we compute the transition path from the initial stationary equilibrium to the new stationary equilibrium. Along the transition path, parameters of the model are chosen to such the model implied moments match up with empirical moments.

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

the aggregate effects and the heterogeneity in the gains and losses from the China shock.

What were the aggregate effects of the China Shock? At an aggregate level, the China shock leads to a two percent increase in GDP over five years. About half of the increase in GDP is from standard, gains from trade effects. The remaining half comes from increases in labor force participation. Specifically, the China shock *increased* employment by about a million and a half jobs. Aggregate consumption increases by one percentage point and, thus, from the national savings identity the implication is that trade with China leads to an *improvement* in the trade deficit.

Aggregate increases in labor force participation, mild increases in consumption, improvements in the trade deficit arise from changes in the consumption-savings behavior of households. 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. Thus, precautionary savings motives kick in (see, e.g., Huggett and Ospina (2001)). For many households, this leads to an increase labor supply to acquire earnings (participation increases) and an increase in how much of those earnings are put aside for a rainy day (thus, a mild increases in consumption). These changes in participation and consumption behavior lead to an improvement in the trade deficit with increased trade from China.

Relative to the popular narrative discussed above, our interpretation is very different. For example, while our model is consistent with the local-labor market evidence of Autor et al. (2013) at the micro level, the China shock increases labor force participation. The disconnect arises exactly because of the nature of the research design behind difference-in-difference approaches—the aggregate labor supply response is not identified, only differential effects are.

A second distinction relative to the popular narrative is that the trade shock (only on the import side) induces an expansion in exports and improving trade deficit. The behavior of the trade deficit is different than popularly interpreted, i.e., if it becomes easier to import, then this leads to an expansion in the trade deficit. Moreover, the behavior is quite different than the observation that increased import penetration did not arrive with equivalent export opportunities. Our interpretation is that our model measures the effect of only a trade shock while abstracting from other shocks were occurring during this same time. As argued in Bernanke (2005), we are keenly interested in the global decline in real interest rates which occurred at nearly the same time goods trade with China expanded. Our intuition is that a decline in real interest rates will offset precautionary motives, lead to a deteriorating trade deficit.

What were the welfare effects of the China Shock? Those who are initially import exposed—and who have the most potential to lose—experience a near zero change in welfare. This stands in sharp contrast to the harsh effects trade has on these labor markets with real wages falling by

two percent over the first five years of the shock. The disconnect between wage outcomes and welfare arises from households actively using margins of adjustment to mitigate the shock. These households increase their labor supply to build up savings and smooth out the foreseen, future negative income shock from China.

The gains to households in non-traded regions are nuanced. Surprisingly, these households experience essentially no change in real wages. The issue is that two forces counteract each other: while the real purchasing power increases for these households, the value of the goods produced by these regions also falls as consumers substitute into cheaper imported goods. Thus, real wages are unchanged for households in non-traded regions. In welfare terms, these households still gain about 0.75 percent. These are dynamic gains which arise from the stochastic nature of comparative advantage—that at some point they will be in a situation that comparative advantage favors them and they reap the benefits.

Households in regions with a comparative advantage—the export exposed—gain the most, a bit less than two percent increase in welfare. These gains come through changes in the aggregate price index with real wages increasing by four percent after five years. However, given the stochastic nature of comparative advantage, these gains in real wages are transitory and, thus, a portion of this is saved for bad times in the future.

The evidence of Autor et al. (2013) is central to disciplining these numbers. We illustrate this point by exploring alternative calibrations with counterfactually larger wage effects. Consistent with our theoretical arguments, this calibration results in a smaller elasticity of substitution across goods and hence stronger pass-through of trade shocks into wages. A larger wage elasticity implies that the losers from trade lose much more both in real wages and in welfare terms. In other words, the data could have pushed us to a conclusion with larger losses from trade, but it instead points us towards the conclusion that trade had at most small negative welfare effects.

Related Literature

Our work builds on a literature that uses cross-regional variation to speak to aggregate outcomes. Obviously our context regards the impact of trade and leans on the results of Autor et al. (2013), but this approach has found wide applicability. For example, the role of housing (see, e.g., Mian et al. (2013) and Mian and Sufi (2014)) in prorogating the recession starting in 2008, or Nakamura and Steinsson (2014) on the effects of government spending shocks.

Our quantitative strategy uses estimates off cross-regional variation to discipline the aggregate responses from structural models. This approach of mixing reduced form approaches and structural models is growing with examples being Nakamura and Steinsson (2014) for government spending shocks, Beraja et al. (2016) for the role of wage rigidities, Acemoglu and

Restrepo (2017) on robots, Jones et al. (2018) look at household leverage. In our setting, the modeling ingredients and mechanisms that we focus on are strongly informed by the evidence of Autor et al. (2013) and Greenland et al. (2017). In a similar spirit, through a procedure which is essentially simulated method of moments, we ask our model match the cross-regional empirical outcomes. Moreover, our model provides a clear interpretation about how these empirical moments map into model parameters.

While quantifying the aggregate effects of a trade shock is important, it is also central to pinning down the distributional effects. A weakness of the research design in Autor et al. (2013) and related evidence is that these estimates are only able to make statements about differential effects—not levels, see, e.g. the discussion in Nakamura and Steinsson (2018). In the case of trade, this critique is particularly pertinent as an optimistic interpretation of the evidence is that all labor markets gained from trade with China, just some gained more than others. Substantively, the disconnect between estimated differential effects and levels shows up prominently in the result that the China shock caused an increase employment by a million and a half at the aggregate level.

Our work is also closely related to growing body work on trade and labor market dynamics (see, e.g., Kambourov (2009), Artuç et al. (2010), Dix-Carneiro (2014), Coşar et al. (2016)). Substantively, Galle et al. (2017) and Caliendo et al. (2015) are closely related papers which focus on both the aggregate and distributional effects in the context of the Autor et al. (2013) evidence. Our key departure from these papers is the focus on an incomplete market setting. There are important costs associated with our approach. We are unable to incorporate the the geographic and sectoral detail found in Caliendo et al. (2015) or Galle et al. (2017). We do not have the clean mathematical expressions relating labor allocations and expenditure shares to welfare as in Galle et al. (2017).

With that said, there are important benefits that motivate our model’s setting. The first reason relates to normative issues. Even if there were losses in income and employment opportunities, open questions are: How large are the losses in welfare? We believe that issue here is the extent of market incompleteness and the ability of households to adjust to trade shocks.¹ 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. In contrast, if households are an island to oneself with no transfers or insurance markets, the welfare losses could be large. Thus, an important contribution of this paper is examine the distributional effects of trade in a standard, intermediate setting—the standard incomplete markets model.

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) and Dixit (1987, 1989a,b).

The second reason incomplete market setting also opens the door to questions about government policy which far understudied within this trade and labor market dynamics literature. As an economy is increasingly exposed to trade does policy have an ability to improve outcomes? Can tariffs improve welfare? Our model provides a natural setting to entertain these types of questions. As one example, our parallel work in Lyon and Waugh (2018) explores these issues and finds that a progressive tax system is an important tool to mitigate the negative consequences of trade.

1. Motivating Facts

Below we discuss macro- and micro-level facts behind the expansion of US-Chinese trade. None of these facts are new, but they are important to state as they help motivate our modeling and calibration strategy.²

1.1. Macro Facts

Fact 1: Large rise in US import exposure to China; really large post 2001. Figure 1 plots US imports of goods and services divided by GDP for the time period 1990 until the beginning of 2008. Total US imports is the top red line and the blue line excludes China, thus, the difference reflects the contribution of Chinese imports. 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. 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 came from an increase in Chinese import exposure.

The dashed black vertical line indicates the timing of the accession of China to the World Trade Organization and the granting Permanent Normal Trade Relations (PNTR) by the United States at the end of 2001. Prior to this change in trade policy, import penetration from China grew at a moderate pace of about a quarter of a percentage point per year. After this period, growth in import penetration from China accelerated dramatically.

Fact 2: An expansion in the US trade deficit. Figure 2 plots the US trade deficit measured as exports minus imports divided by GDP. During this time period, the trade deficit grew by about five percentage points in total; three percentage points prior to the rise in Chinese import penetration, and then two percentage points subsequently.

Together, **Fact 1** and **Fact 2** suggest that Chinese trade did not come with the replacement of job opportunities through a rise in exports. This point has been emphasized by Autor et al. (2013) as to why the rise of China may have forcefully affected the US labor market in the 2000s.

²An interactive python notebook that replicates these facts and facilitates exploration is posted at https://github.com/mwaugh0328/ADH_replication.




Figure 1: US Imports




Figure 2: US Trade Deficit

1.2. Micro Facts

The next three facts focus on labor market outcomes from Autor et al. (2013) and migration responses in Greenland et al. (2017). These studies 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. The main 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 takes aggregate US imports from China M_{ucjt} for industry j and 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 outcome Y_{it} with the following empirical specification:

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

One important issue with this specification is that the error term embeds factors that simultaneously change a commute zone's trade exposure and labor market outcomes. As our model makes clear (see Section 4), local productivity shocks is a key threat to identification as it would change that commute zones comparative advantage and labor markets outcomes. To avoid this problem, Autor et al. (2013) estimate (2) using other countries imports from China as an instrument which would be correlated with US trade exposure, but orthogonal to local productivity shocks.

In the facts below, we report estimates after standardizing the ΔIPW_{uit} . That is we demeaned this measure and divided by its standard deviation. As in Autor et al. (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 et al.'s (2013) estimate of the response of commute zone, average household labor income per adult 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 ten year change).

Fact 4: Import exposure increased non-participation. The second column in Table 1 reports the response of commute zone non-participation in the labor force in response to trade expo-

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.

sure. 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. This coefficient says that a one standard deviation increase in trade exposure reduced a commute zone’s population by -0.05 log points and it’s 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 et al. (2017), however, call into question this finding. For our purposes, an important issue is the treatment of pre-trends of population growth in commute zones. In the Autor et al. (2013) specification, regional controls (e.g. midwest, pacific northwest, etc) are included and, thus, control for different population trends across regions. This does not control, however, for pre-trends at the commute zone level. As Greenland et al. (2017) show, visual inspection and regression evidence at the commute zone suggest pre-trends are an issue and they at the commute zone suggest using lagged population growth at the commute zone level to control for these trends.

The second column of Table 2 reports the effects when regional controls are replaced with lagged population. The point estimate increases by an order of magnitude—from essentially zero to -1.40 log points. The dramatic change in point estimates suggest that prior to the 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 Autor et al. (2013) estimates upward. With that said, at conventional levels of statistical significance, the point estimate is not statistically different from zero.

Greenland et al. (2017) raise two other issues support the view that there was a migration re-

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 et al. (2017) (GLM) replace ADH regional controls with agged population growth at the commute zone level.

sponse. The first issue is the use of Census data versus nationally representative samples; using 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 et al. (2017) suggests that there was a 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 we use these estimates as inputs into our quantitative 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 changes in participation can mitigate and assist self-insurance, a third is migration. Thus, our model that allows us to entertain multiple, standard mechanisms for which households can mitigate the negative shocks.

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

This section 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 et al. (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,

$$q(\omega) = z(\omega)\ell, \tag{3}$$

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

$$\log z_{t+1} = \phi \log z_t + \epsilon_{t+1}, \tag{4}$$

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 produce 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 index by i which corresponds with choice of the household to move or not. As in Artuç et al. (2010) and Caliendo et al. (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 a 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.³ 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

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

³Given 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.

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 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.⁴

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).

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

⁴Relative to the trade and labor market dynamics literature, this is similar to the second specification solved in Artuç et al. (2010). Moreover, it has the advantage (say, relative to Caliendo et al. (2015)) of being relatively simple, yet allows us to be specific about the interaction between trade flows and capital flows.

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.⁵

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 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

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

⁵This 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.

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 country 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.⁶ Then the supply of labor to intermediate good producers with productivity

⁶This 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.

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.

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

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

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 et al. (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. $\lambda(\mathbf{s}, a)$.⁷ 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 struc-

⁷To 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.

ture 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 et al. (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 emphasized in Arkolakis et al. (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,




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




Figure 4: Wages

this is the region where the goods are non-traded. Exactly like the inner, non-traded region in the Ricardian model of Dornbusch et al. (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 et al. (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.

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 et al. (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 et al. (2013)

The preceding results relate closely to the empirical specification and evidence of Autor et al. (2013) and Acemoglu et al. (2016) that link changes in trade exposure with labor market outcomes such as wages (see Section IV.B of Autor et al. (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 et al. (2013) (see equation (5)). Autor et al. (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 expo-

sure 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 .⁸ 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 et al. (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 generally 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.

⁸Moreover, 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.





Figure 5: Labor Supply by Assets and Trade Exposure

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. 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.

While illustrative, the previous section left open several questions. While our model predicts





Figure 6: Migration by Assets and Trade Exposure

heterogenous responses, how the aggregate gains from trade offsets these responses it 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.

5. Calibration

This section outlines our calibration approach. Our calibration approach asks our model to replicate 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.

5.1. Empirical Analogs and Predetermined Parameters

Time, Geography, and the Interest Rate The time period is set to a year. Given the time period we set discount factor equal to 0.95. This is value is at the top end with those used in Krueger et al. (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 calibrate our model to the empirical evidence of Autor et al. (2013). Thus, our empirical counterpart to an island is a Commuting Zone. 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 model’s parameters 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 et al. (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 $\Theta = \{B, \bar{a}, m, \tau_{im}, \tau_{ex}, \phi, \sigma_z, \sigma_w, \tau'_{im}, h, \theta, \sigma_\nu\}$. 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 .

- **Net worth of households.** Krueger et al. (2016) report from the Survey of Consumer Finances find that share of aggregate net-worth held by the bottom 40 percent of households is essentially zero. We target this value in the initial stationary equilibrium by computing the share of net worth held by the bottom 40th percentile of households in our model. 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 in the early 2000s. This moment is most informative about the initial import and export trade cost (which we assume to be the same initially).
- **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, τ'_{im} .
- **Variance in commute-zone-level labor earnings.** In the initial stationary equilibrium, we target standard deviation of log labor earnings, across commute zones/islands, of 0.0368. This value is measured in the data by using the decennial Census data from Autor et al. (2013) and focusing on the period between 1990 and 2000. Combined with the other parameters controlling the stochastic process for world prices, this moment is informative σ_z .
- **Dynamics of comparative advantage.** In the initial stationary equilibrium, we calibrate the model to match evidence from Hanson et al. (2015) on the dynamics of comparative advantage. Hanson et al.’s (2015) find that of sectors in the top five percentile of exports, only forty percent of these comparative advantaged sectors remain in the top five percent after twenty years. We model match this moment in the stationary equilibrium by computing what fraction of exporting islands/goods in the top five percent, remain there after twenty years. This is moment is very informative about ϕ , the persistence of z and p_w .

- **Aggregate trade elasticity.** Given the long-run increase in trade, we can measure the implied trade elasticity, i.e. how much trade changes given the change in trade costs. We then ask that our model's trade elasticity equals four (Simonovska and Waugh (2014)).

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 observation 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 et al. (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 et al.'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 et al. (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 et al. (2013) have identified the elasticity of labor supply to a trade shock, this moment informs our home production parameter.

- **Migration elasticity from Greenland et al. (2017).** We target the Greenland et al. (2017) migration elasticity described in Table 2. This moment is most informative about the scale of the preference shock σ_v . 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 et al. (2017) helps pin this parameter down.

5.3. Implementing the China Shock

Our 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 is a 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

Table 3: Calibration: Parameters and Aggregate Moments

Parameter	Value	Target	Data	Model
Disutility of work, B	1.05	Aggregate participation rate	66	66
Migration Cost, m	2.46	CMZ. migration rate	3	3
Borrowing Limit, $-\bar{a}$	0.65	Net worth, bottom 40 percentile	0	0
Pre-China Trade Cost (τ_{ex}, τ_{im})	1.31	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	Var. in Log CMZ earnings	0.037	0.037
Autocorr. z, p_w (ϕ)	0.98	Exports in top 5% in $t + 20$	40	40
Std. Dv. of p_w (σ_w)	0.07	Trade Elasticity	4	4.6

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

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 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 $\theta, w_h, \sigma_\nu, \sigma_w$.

Step 2. We pick the parameters $\{B, \bar{a}, m, \tau_{im}, \tau_{ex}, \phi, \sigma_z\}$ so that the initial stationary equilibrium (pre-China shock period) replicates the labor force participation rate, migration rate, net worth of households, dynamics of comparative advantage, variance 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 et al. (2012) formula.

Step 4. Given the transition path, we simulate data sets analogous to those in Autor et al. (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 et al. (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 which mimics the instrument strategy of Autor et al. (2013) in the model.

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

5.5. Calibration Results

Table 3 reports the calibration results parameters and the associated aggregate moments. Figure 7 plots the model's prediction for imports relative to the data. The vertical, black dashed line indicates when the new path of trade costs is announced within the model. Similar to the data, the model accurately generates rising trade exposure over the early 2000s. Overall, the model is very flexible and is easily able to fit most moments.

Table 4 reports the micro-moments, model prediction/fit, and the resulting parameter values. The first row reports the estimates from Autor et al. (2013) and Greenland et al. (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 is that results of Autor et al. (2013) are informative about structural parameters that govern the underlying data generating process; Section 6.4 illustrates this point more substantively.

A second 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 et al. (2013) evidence, is that pass-through was not that strong. In section 6.4 we explore alternative perturbations of the data and show how larger wage elasticities would imply lower demand elasticities.

The other parameters are not abnormal. Here the home production is about thirty percent of average output/labor earnings. The moving shock variance is low relative to other papers such as Artuç et al. (2010) or Caliendo et al. (2015). For example, Artuç et al. (2010) find estimates of the standard deviation of the moving preference shock which is around four times larger than




Figure 7: Rising Trade Exposure: Model and Data

our estimate. This observation is important in the sense that in our model there other economic motives for moving (or not) and hence there is less burden on the preference shocks to drive moves.

6. The China Trade Shock

Armed with the calibrated model which is replicating the aggregate changes in trade and micro-level evidence, we use the as laboratory to answer questions regarding the aggregate effects and the heterogeneity in the gains and losses from the China shock. We answer these questions by first studying behavior at the micro-level, then aggregate, and finally address the welfare questions we care about.

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” at how real wages, labor supply, and consumption respond at the island level.

Table 4: Calibration: Diff-in-Diff 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	-5.66	1.15	-1.57
	Demand elasticity θ	Home production w_h	Moving shock variance σ_ν
Parameter Values	10.062	0.18	0.95

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

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 several 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 only reason is through increases in real purchasing power, i.e. the decline in the CES price index. And because this effect is common across all exporting regions, they 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, only the CES price index P changes. The exporting cost is not changing, nor are world prices, thus this 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. 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) net of the gain in purchasing power. To see this point, wages for a given island type are:

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

Two forces work in the opposite direction: one force is the decline in the trade cost—effective




Figure 8: Real Wages

world prices are falling. The second force is the rise in real purchasing power, again, the decline in the CES price index. 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). In this region real 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 locally 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 point, examine the labor market clearing condition for islands where

the good is non-traded:

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

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 zero. Quantitatively, we find small changes because of aggregate demand effects or changes in labor supply.

This observation defies the standard intuition that those in non-trade-exposed labor markets should benefit from trade through increased purchasing power. With frictional labor markets, this logic is incomplete and it implies that this benefit will not accrue towards them as consumer substitute away from non-traded goods. This logic also implies that the CES price index is changing by a margin larger than the output response. Why? 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 parameter θ . Recall, that the evidence of Autor et al. (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. 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. In sum, the Autor et al. (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 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., ok stuff today, but foreseeably bad in the future. Standard intertemporal smoothing logic implies that these households will want transfer income from today into the future. In contrast, for those that are initially export exposed, they see increases in real wages coming. That is they are facing a positive, back-loaded income shock, i.e., ok stuff today, but foreseeably good in the future. Thus, these households 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 exposed (red and to the left on the y-axis) initially increase their labor supply by a substantial margin, almost four percentage points for the most exposed. Households on these islands see the negative effects coming and use labor supply to fulfill their intertemporal savings motives. That is work hard today to build up savings and facilitate the smoothing out of the foreseen, future negative income shock. 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 on these islands see the gains from trade coming their way. And thus they transfer some of this new found wealth into benefits today by substituting into leisure. And this effect dissipates as households eventually increase their labor supply as the benefits from participating in increasingly generous labor market.

Consumption/Savings. Figure 10 illustrates the change in island-level savings rates. Everything here is the same as Figure 9, except now the vertical, z-axis reports the percentage point 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 it provides a picture of the change in savings behavior.

In many ways, Figure 10 mimics the patterns in labor supply. Import exposed islands are facing a negative, back-loaded income shock and want to find ways to transfer income today into the future to mitigate the foreseen, future negative income shock. In contrast, export exposed islands who foresee increases in real wages increase their consumption today and decrease their savings to transfer future benefits into today.

Long Run Outcomes. Figures 9 and 10 also illustrate that these patterns reverse as the economy approaches the new stationary equilibrium. Export exposed islands tend to participate more and save more; import exposed islands tend participate less and save less. Moreover, the islands in the middle (which are non-traded regions) which work and save at the highest rates.

Precautionary motives drive this behavior. In the model, 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




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




Figure 10: Model: Change in Savings, Across Islands

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 leads to an increase in labor supply to acquire income and a higher savings rate.

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 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 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, six years after the shock. The improvement 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





Figure 11: Model: Consumption, Output, Labor Supply



41
Figure 12: Trade Defect

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. Given that production is increasing through labor supply, exports must increase as product not consumed domestically goes abroad. Thus, exports increase more than imports and the trade deficit improves. This mechanism is essentially the same as “venting-out” mechanism explored in Almunia, Antràs, Lopez-Rodriguez, and Morales (2018) where essentially excess supply—in our case caused by intertemporal motives—is exported.

Discussion. Relative to the “popular narrative” discussed in the introduction, these findings are important for several reasons. First, the aggregate response of labor supply and income moves in the **opposite direction** of a naive interpretation of Autor et al. (2013) labor force participation and income evidence. There is no inconsistency here since, By construction, our model is matches their evidence as shown in Table 4. This finding clearly illustrates how diff-in-diff evidence has little content when speaking about the magnitude and direction of aggregate outcomes.

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. An the missing insight that our model clarifies is that an expansion of trade changes the savings behavior of households which leads to a smaller trade deficit. Our model’s emphasis on the savings behavior of households fits well with the narrative that global financial forces (as advocated by Bernanke (2005)) are a more likely culprit.

An oft asked question or comment goes along the lines that “well the behavior of the trade deficit is counterfactual.” This is true, but that’s not the point. We are using the model to measure effect of the change in the ability to import goods from abroad on aggregate outcomes. We abstract from important changes in the economy, in particular, the large fall in real interest rates during the same exact time period as the acceleration of trade with China. In the model, a decline in real interest rates will offset these precautionary effects and lead to a deteriorating trade deficit. But this is a separate point with regards to understanding how the change in the ability to import goods (statically) affects aggregate outcomes.

6.3. Welfare: Winners and Losers from the China shock?

How does this all add up in terms of welfare? Let’s first take stock. Section 6.1 confirms the idea that those who were import exposed suffered losses in the labor market, but the effects on others—non-traded and exporting regions—are far more nuanced. Section 6.2 suggests that households’ faced with these changes in the labor market engaged in precautionary measures

in reaction to the shock. This section measures welfare and examine the extent to which some households lost or won from trade exposure.

How We Measure Welfare. First, we measure welfare and report numbers at the island level, asset level. To formalize this, define expected lifetime utility for a household on island \mathbf{s} with assets a in the initial stationary equilibrium as

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

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

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

The value of $\chi(\mathbf{s}, a)$ which makes $W(\mathbf{s}, a; \chi(\mathbf{s}, a)) = W'(\mathbf{s}, a)$ is our measure of welfare gains or loses. In some of the measures below, we aggregate to the island by integrating over the consumption equivalents according to the appropriate distribution.

Welfare Effects of the Shock Figure 13 illustrates the full range of outcomes by plotting the welfare gains for islands and asset holdings. Table 5 summarizes our findings in addition to reporting the change in real wages;

On average, the gains are not out-of-line with what a simple calculation from Arkolakis et al. (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.

The interesting outcomes are with respect to the distributional consequences. First, as Figure 13 shows, those who are initially import exposed gain the least and some actually loose. In fact those who loose the most are the households with the least asset holdings and in the most exposed islands. These are the households who have few mechanisms to smooth out the shock, in particular, assets to pay for a move away from these islands. Thus, the model does have households who are “stuck” in a loose sense and suffer. With that said, the population who happen to loose is about half of a percent and the average loss for the losers is about -0.20 percent in consumption equivalents.

Some losers, minor gains, stands in sharp contrast to the harsh effects that trade has on these labor market outcomes—wages falling by two percentage points over the first five years of the shock. There are several reasons for the disconnect between wage outcomes and welfare which is unique to our paper. First is the fact that households are actively using margins of adjustment to mitigate the effects of the shock. Import exposed households increase participation, they build up savings all to smooth out the foreseen, future negative income shock from China. A second reason for the disconnect works through the stochastic nature of comparative advan-

Table 5: Welfare

		Welfare (Baseline)	Δ Real Wages
Initial Exposure	Import Exposed	0.09 [0.05]	-2.55 [0.05]
	Non-Traded	0.71 [0.68]	0.18 [0.70]
	Export Exposed	2.12 [0.25]	5.45 [0.25]
Average		1.02	1.10

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

tage. Due to fluctuating local and world price shocks, at some point comparative advantage will change favor these households. And it was the evidence of Hanson et al. (2015) which disciplined the extent to which comparative advantage persists or does not. This mechanism interacts with the change in trade costs because there are now larger gains associated with advantageous situations. That is households will eventually experience the gains on the right of Figure 8.

The final reason for the disconnect between welfare and wages relates to “option value” effects emphasized by Artuç et al. (2010). Like in Artuç et al. (2010), because trade increases average real wages and the variance of wages, thus households indirectly benefit as the option of moving increases. Moreover, because the move is to a random island in our model the value of the option is stronger for import-exposed households. This is true even absent an average increase in wages—this statement is about the variance of the wages.

These mechanisms lead to substantial welfare gains for those in non-traded regions, about 0.75 percent increase in welfare yet experience essentially no change in real wages. Due to the stochastic nature of comparative advantage, these households benefit from more from the times when comparative advantage will favor them and, also, buffer them when it does not. It also increases the option value of a move as average wages increase. Quantitatively, however, the disconnect between labor market outcomes and welfare is smaller than the import-exposed case with only a half percentage point difference (in contrast to a nearly three percentage point difference for the import exposed).




Figure 13: Welfare Across Islands and Asset States

Finally, those in the exporting regions gain the most, a bit less than two percent with corresponding real wage gains of nearly four percent. Recall several points discussed above. The the labor market gains are purely through changes in the price index and, hence, a rise in real wages. Second, gains in the labor market do not map one-for-one into welfare. Unlike the other two regions, the welfare gains for the export exposed are *smaller* than the wage gains. The stochastic nature of comparative advantage implies that these gains are transitory and large relative to the average. Moreover, a portion of the benefit is simply saved for possible bad times in the future.

6.4. The role of the Autor et al. (2013) evidence

The evidence of Autor et al. (2013) is central to disciplining the small negative welfare effects from trade seen in Figure 13. We illustrate this point by showing what patterns in the data would suggest larger welfare losses from trade. We proceed by re-calibrating the model to match all the same moments, except target a larger wage elasticity of -8.60 rather than -4.30 as in seen the data.

Table 6: Alternative Calibrations

		Big ADH Cal. (-8.60, 4.74)		Baseline (-4.30, 9.53)	
		Welfare	Δ Log Wages	Welfare	Δ Log Wages
Initial Exposure	Import Exposed	-.06 [0.21]	-3.00 [0.21]	0.19 [0.09]	-2.19 [0.09]
	Non-Traded	0.73 [0.33]	0.00 [0.33]	0.75 [0.68]	0.34 [0.68]
	Export Exposed	1.71 [0.46]	3.50 [0.46]	1.64 [0.25]	3.99 [0.25]
Average		0.91	1.01	0.94	1.06

Note: Welfare values are lifetime consumption equivalents; values in brackets report the share of the population in that category. First two columns are from a calibration targeting a ADH wage elasticity of -8.60.

The top two lines of the Table 6 report the key calibration outcomes from this exercise. We are easily able to match the higher wage elasticity of -8.60. All other moments are essentially the same as reported above. The second line of Table 6 reports the demand elasticity. As argued above this is the key parameter for which the wage elasticity is informative about and is the parameter. Here the demand elasticity is near five—half the value found in our baseline calibration with the value of ten.

Because this calibration results in a smaller elasticity of substitution across goods and the change in the volume of trade is the same, the implication is that there is a stronger pass-through of the trade shock into wages. The right most column reports the change in real wages across different exposure to trade. Here wages fall by more for the import exposed; those in the middle experience literally no change; and those at the top experience smaller wage gains. So this parameterizations essentially results in more losers, less winners in terms of labor earnings. Welfare essentially mimics the same patter, but with the losses from trade being small.

7. Final Thoughts

Researchers with more experience than us have remarked to us in (sometimes in a semi-pejorative way) the following: “We know that the standard incomplete market model gives outcomes which are in, welfare terms, near a complete markets benchmark.” Well that’s the point—take

the evidence of Autor et al. (2013) and give households limited, but some self-insurance mechanisms and the interpretation about the negative effects from trade *completely change*. With that said, our paper is a starting point that suggests where to look if our prior is that trade had negative consequences. One area to explore are different financial market frictions, specifically housing wealth or forms of liquid and non-liquid wealth (e.g., Kaplan and Violante (2014)).

New Keynesian researchers often comment that Keynesian frictions are needed, in particular wage rigidities. To keep things simple, we deliberately focused on a competitive framework. Again, this is a starting point, not the end and we have provided a natural benchmark for more research. Models with a richer labor market structure: within commute zone selection, unemployment, wage rigidities (see, e.g., Beraja et al. (2016))— are important extensions.

Is all well? No more complaining about the effects of trade—not quite. The economy that we analysed is (most likely) not constrained efficient. That is the competitive equilibrium does not correspond with a planner who faces the same constraints. This suggests the government policy has a role to play and specifically as an economy is increasingly exposed to trade. As one example, our parallel work in Lyon and Waugh (2018) explores these issues and finds that a progressive tax system is an important tool to mitigate the negative consequences of trade yet reap the gains from trade. Further exploration of these issues is an important area of future research.

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, S141–S198.
- ACEMOGLU, D. AND P. RESTREPO (2017): "Robots and Jobs: Evidence from US Labor Markets," Tech. rep., National Bureau of Economic Research.
- AIYAGARI, S. R. (1994): "Uninsured idiosyncratic risk and aggregate saving," *The Quarterly Journal of Economics*, 109, 659–684.
- ALMUNIA, M., P. ANTRÀS, D. LOPEZ-RODRIGUEZ, AND E. MORALES (2018): "Venting Out: Exports During a Domestic Slump," Tech. rep., National Bureau of Economic Research.
- ARKOLAKIS, C., A. COSTINOT, AND A. RODRÍGUEZ-CLARE (2012): "New Trade Models, Same Old Gains?" *American Economic Review*, 102, 94–130.
- ARTUÇ, E., S. CHAUDHURI, AND J. MCLAREN (2010): "Trade shocks and labor adjustment: A structural empirical approach," *The American Economic Review*, 100, 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, 2121–2168.
- BERAJA, M., E. HURST, AND J. OSPINA (2016): "The aggregate implications of regional business cycles," Tech. rep., National Bureau of Economic Research.
- BERNANKE, B. (2005): "The global saving glut and the US current account deficit," Tech. rep.
- 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, 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, 625–663.
- DIX-CARNEIRO, R. (2014): "Trade liberalization and labor market dynamics," *Econometrica*, 82, 825–885.
- DIXIT, A. (1987): "Trade and insurance with moral hazard," *Journal of international economics*, 23, 201–220.
- (1989a): "Trade and Insurance with Adverse Selection," *The Review of Economic Studies*, 56, 235–247.

- (1989b): “Trade and Insurance with Imperfectly Observed Outcomes,” *The Quarterly Journal of Economics*, 104, 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, 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’Economie*, 18, 258–272.
- EATON, J. AND S. KORTUM (2002): “Technology, geography, and trade,” *Econometrica*, 70, 1741–1779.
- GALLE, S., A. RODRÍGUEZ-CLARE, AND M. YI (2017): “Slicing the Pie: Quantifying the Aggregate and Distributional Effects of Trade,” Working Paper 23737, National Bureau of Economic Research.
- 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,” Tech. rep., 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, 953–969.
- HUGGETT, M. AND S. OSPINA (2001): “Aggregate precautionary savings: when is the third derivative irrelevant?” *Journal of Monetary Economics*, 48, 373–396.
- JONES, C., V. MIDRIGAN, AND T. PHILIPPON (2018): “Household Leverage and the Recession,” Tech. rep., National Bureau of Economic Research.
- KAMBOUROV, G. (2009): “Labour Market Regulations and the Sectoral Reallocation of Workers: The Case of Trade Reforms,” *Review of Economic Studies*, 76, 1321–1358.
- KAPLAN, G. AND G. L. VIOLANTE (2014): “A model of the consumption response to fiscal stimulus payments,” *Econometrica*, 82, 1199–1239.
- 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, 1477–1507.

- LUCAS, R. E. AND E. C. PRESCOTT (1974): "Equilibrium search and unemployment," *Journal of Economic Theory*, 7, 188–209.
- LYON, S. AND M. WAUGH (2018): "Redistributing the Gains From Trade Through Progressive Taxation," .
- MIAN, A., K. RAO, AND A. SUFI (2013): "Household Balance Sheets, Consumption, and the Economic Slump*." *Quarterly Journal of Economics*, 128.
- MIAN, A. AND A. SUFI (2014): "What explains the 2007–2009 drop in employment?" *Econometrica*, 82, 2197–2223.
- 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," Tech. rep., National Bureau of Economic Research.
- NAKAMURA, E. AND J. STEINSSON (2014): "Fiscal stimulus in a monetary union: Evidence from US regions," *American Economic Review*, 104, 753–92.
- (2018): "Identification in macroeconomics," *Journal of Economic Perspectives*, 32, 59–86.
- NEWBERY, D. M. AND J. E. STIGLITZ (1984): "Pareto inferior trade," *The Review of Economic Studies*, 51, 1–12.
- ROGERSON, R. (1988): "Indivisible labor, lotteries and equilibrium," *Journal of monetary Economics*, 21, 3–16.
- SIMONOVSKA, I. AND M. E. WAUGH (2014): "The Elasticity of Trade: Estimates and Evidence," *Journal of International Economics*, 92, 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.