The London School of Economics and Political Science

Three Essays on Globalization and Economic Development

March 2013

Benjamin Faber

A thesis submitted to the London School of Economics for the degree of Doctor of Philosophy.

Declaration

I certify that the thesis I have presented for examination for the MPhil/PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified in it).

I warrant that this authorization does not, to the best of my belief, infringe the rights of any third party.

I declare that my thesis consists of 37,178 words.

Abstract

How do falling barriers to flows of trade and information both within and across countries affect economic livelihoods in developing countries? The three chapters presented in this PhD thesis aim to contribute to our understanding of this question.

Table of Contents

Acknowledgements

I am grateful to Robin Burgess, Dave Donaldson, Henry Overman, Stephen Redding, Michael Storper, and Daniel Sturm for their advice and encouragement at various stages of the PhD.

I thank my parents, Rosel and Georg, for giving me the confidence to complete this work.

I thank my brother, Dominic, for giving me the ambition to complete this work.

I thank my amazing wife, Vivi, for her patience, presence and confidence at every stage of this process.

Last but not least, many thanks go to the many inspiring friends I have met during the PhD, especially Felix Weinhardt, Rosa Sanchis-Guarner, Horacio Larreguy, Michael Best, and Francisco Costa to name just a few.

Finally, financial support from the UK Economic and Social Research Council is very gratefully acknowledged.

Chapter 1

Trade Integration, Market Size, and Industrialization: Evidence from China's National Trunk Highway System*

Benjamin Faber†

Abstract

This paper presents evidence in favor of the hypothesis that trade integration between large and small markets can reinforce the concentration of economic activity. Krugman's (1980) home market effect provided a microfoundation for the proposition that market size is a determinant of industrialization. The same channel also has important, yet so far untested implications for the consequences of falling trade costs between asymmetric markets. I exploit China's National Trunk Highway System as a large scale natural experiment to test for the home market channel of trade integration. The network was designed to connect provincial capitals and cities with an urban population above 500,000. As a side effect, a large number of small peripheral counties were connected to large metropolitan city regions. To guide estimation, I derive qualitative and quantitative predictions from a tractable general equilibrium trade model. To test these predictions, I construct least cost path spanning tree networks as instruments for route placements on the way between targeted city nodes. The results suggest that network connections had negative growth effects among peripheral counties due to reduced industrial output growth. Counterfactual estimations of the calibrated structural model suggest that the home market channel can both qualitatively and quantitatively account for the observed effects.

Keywords: Market integration; transport infrastructure; home market effect
JEL Classification: F12; F15; O18; R12

*I am grateful to Lawrence Crissman at the ACASIAN Data Center at Griffith University for his support on the GIS data. I thank the MIT Economics Department for their hospitality when working on this project. I am grateful to Robin Burgess, Dave Donaldson, Esther Duflo, Guy Michaels, Henry Overman, Nancy Qian, Steve Pischke, Steve Redding, Frederic Robert-Nicoud, Thomas Sampson, and Daniel Sturm for helpful comments. This research was financially supported by the UK Economic and Social Research Council.
†London School of Economics and Centre for Economic Performance; Email: b.s.faber@lse.ac.uk

1 Introduction

A large share of world trade takes place between regions within countries.[1] In this context, transport infrastructure investments have been a prominent policy tool that directly affects the degree of within-country trade integration.[2] These policies frequently combine national efficiency with regional equity objectives under the presumption that integration promotes both growth as well as the spreading of economic activity to peripheral regions.[3]

Trade theory has suggested otherwise. In the presence of increasing returns to scale in production, the home market effect in Krugman (1980) and Helpman and Krugman (1985) provided a microfoundation for the idea that market size is a determinant of industrialization. The same microeconomic channel also suggests that falling trade costs between *ex ante* asymmetric regions can reinforce the concentration of economic activity. Despite the fact that transport networks almost inevitably connect both metropolitan and peripheral regions, our existing empirical knowledge about the role of market size in trade integration is very limited.

This paper exploits China's National Trunk Highway System (NTHS) as a large scale natural experiment to contribute to our understanding of this question. The first contribution is related to a growing body of empirical evaluations of transport infrastructure.[4] It has been an implicit assumption of this literature that relative market size has no implications for the consequences of falling trade costs. A common basis for this assumption is the neoclassical tradition of constant returns to scale in production. This paper relaxes this assumption and exploits the NTHS as a source of plausibly exogenous variation in trade cost shocks across a large number of *ex ante* asymmetric regions to learn about the role of market size in trade integration.

A second contribution is to present a novel empirical test of one of the central microeconomic channels of new trade theory. Existing empirical literature on the home market effect has mainly followed the original comparative static in Krugman (1980) and Helpman

[1] This follows from the gravity literature in international trade. For evidence using microdata on US plant transactions see Hilberry and Hummels (2008).

[2] Transport infrastructure has been the second most important spending category in World Bank lending over 2001-06, of which 73% were spent on highways and roads (www.worldbank.org).

[3] From the World Bank's Transport Business Strategy 2008-2012 (2008, pp. 03): "*One of the best ways to promote rural development is to ensure good accessibility to growing and competitive urban markets.*" For a discussion of the combination of efficiency with regional equity objectives in the context of European Structural Fund spending on transport see Puga (2002).

[4] See discussion of related literature further below in this section for references.

and Krugman (1985) in terms of a testable cross-sectional relationship between consumption and production shares across countries. This paper, on the other hand, exploits a large scale infrastructure policy to test for the home market channel in the context of falling trade costs.

The findings of the paper present evidence in favor of the home market channel of integration. The main results can be summarized as follows. NTHS network connections had a significant negative effect on GDP growth among peripheral counties on the way between the targeted metropolitan network nodes relative to non-connected peripheral counties due to reduced industrial output growth. These effects vary significantly across pre-existing county characteristics as predicted by theory. Counterfactual estimations of the calibrated structural model suggest that the home market channel can both qualitatively and quantitatively account for the observed effects. Finally, counterfactual policy simulations suggest significant aggregate welfare gains from trade, and provide the insight that the observed increase in concentration of nominal production does not necessarily imply parallel welfare distributional consequences due to significant price index reductions among connected peripheral counties.

To guide estimation, I adapt the canonical model in Helpman and Krugman (1985) to an empirical setting with multiple and *ex ante* asymmetric regions.[5] In addition to the tractability and closed form solutions that this choice facilitates, it assures that the predictions and their empirical evaluation are based on the identical set of economic forces that underlie the original microfoundation of the home market effect.

In a first step, I derive a series of qualitative predictions on both average connection effects as well as their heterogeneity with respect to pre-existing county characteristics from a tractable three region policy scenario. In the second step, I derive quantitative predictions in the full multi-region setting based on simulations of the calibrated structural model subject to observed inter-county trade cost changes from Chinese transport network data.

To empirically test these predictions one requires exogenous variation in trade cost changes across a sufficient number of *ex ante* asymmetric regions. In this context, China's NTHS provides a wide range of pre-existing regional size asymmetries in combination with the explicit targeting of a subset of metropolitan network nodes. Its stated objectives were to

[5] Krugman (1980) and Helpman and Krugman (1985) are based on identical microfoundations of the home market effect but impose alternative conditions to pin down equilibrium outcomes. The former imposes regional trade balance conditions, whereas the latter introduces a freely tradeable numeraire sector. The choice of the Helpman-Krugman model is due to its maintained tractability in a multi-asymmetric region setting. See also Behrens *et al.* (2009) for a discussion of this choice in a multi-region setting.

connect provincial capitals, cities with an urban registered population above 500,000, and border crossings on a single expressway network (World Bank, 2007b, see Figure 1). While the targeted metropolitan city centers represented less than 1.5% of China's land area and 14% of its population, they accounted for 50% of China's non-agricultural production before the bulk of the network was built in 1997. The average pre-existing market size difference between non-targeted peripheral counties and targeted metropolitan city centers was 1:24 in terms of county GDP.

Despite the stated objectives, the assumption of random route placements on the way between targeted city nodes would be strong. Both the NTHS planning process and descriptive statistics suggest that planners targeted politically important and economically prosperous counties on the way between targeted destinations. To address these concerns I propose an instrumental variable strategy based on hypothetical least cost path spanning tree networks. These correspond to the question which routes planners would have been likely to build if the sole policy objective had been to connect all targeted destinations subject to construction cost minimization. I use remote sensing data on land cover and elevation in combination with Dijkstra's (1959) optimal route algorithm to compute least costly construction paths between targeted destinations. I then feed these bilateral cost parameters into Kruskal's (1956) minimum spanning tree algorithm to identify the subset of routes that connect all targeted nodes to minimize network construction costs. I also construct a more extensive, but less precise straight line spanning tree network that is subject to a different trade-off between route precision and the number of captured bilateral connections.[6]

The baseline identifying assumption is that county location along an all-China least cost spanning tree network affects changes of county level economic outcomes only through NTHS highway connections, conditional on province fixed effects, distance to the nearest targeted node, and controls for pre-existing political and economic characteristics. To assess the validity of the exclusion restriction I report how baseline IV point estimates are affected by the inclusion of county controls, and test the model's predictions on the heterogeneity of the connection effects. In addition, the relatively recent nature of the NTHS also allows me to test for network connection effects on identical county samples both before and after the

[6] The second network is related to Banerjee *et al.* (2009) who propose straight line connections between Treaty Ports and historical cities as instruments for railway lines. The combination of least cost path routes, such as straight lines, with a spanning tree algorithm allows me to instrument for the choice of bilateral connections in addition to route placements on any given connection.

network was built as a placebo falsification test.

In conclusion, the paper provides two novel insights. First, the paper presents evidence in favor of the hypothesis that market integration can reinforce the concentration of economic activity. Second, the paper presents a novel empirical test of the home market effect using a very different setting and source of variation compared to existing literature discussed below. These findings are in support of theoretically motivated policy discussions in the trade literature (Fujita *et al.*, 1999; Baldwin *et al.,* 2003), and serve to emphasize the importance of potentially unintended general equilibrium consequences when evaluating and planning large scale transport infrastructure policies.

This research is related to existing empirical literature on the home market effect in international trade (Davis and Weinstein, 1996; 1999; 2003; Head and Ries, 2001; Hanson and Xiang, 2004; Brulhart and Trionfetti, 2009). This literature has mainly followed the original comparative static result of the home market effect in Krugman (1980) and Helpman and Krugman (1985), that an increase in relative market size leads to an increase in the share of industrial production[7], and adapted it to a multi-sector industrial setting in order to exploit variation in markets sizes and production across countries and industries.[8] Instead of relying on cross-sectional variation in consumption and production shares, this paper tests for the home market channel in the context of falling trade costs between *ex ante* asymmetric regions.

The paper is also related to a growing empirical literature on the evaluation of transport infrastructure. Recent contributions have studied the economic effects on suburbanization (Baum-Snow, 2007), skill premia in local labor markets (Michaels, 2008), long term GDP effects (Banerjee *et al.*, 2009), gains from trade (Donaldson, 2010), urban form (Baum-Snow *et al.*, 2012), and city growth (Duranton and Turner, 2012).[9] This paper's contribution is to

[7] There is debate about the robustness of the home market effect. Davis (1998) finds that it disappears when all sectors incur transport costs. Krugman and Venables (1999) show that the home market effect exists as long as some homogenous goods have low transport costs or some differentiated goods have zero fixed costs.

[8] A notable exception is Head and Ries (2001) who also follow the original prediction, but in addition report results on how changes in US-Canada trade barriers affect production shares depending on pre-existing relative market sizes across industries. They do not find evidence in support of the home market channel of integration. Rather than using cross-sector variation within manufacturing between two integrating countries, the present paper exploits variation in trade cost shocks across a large number of asymmetric regions within a country. Also note that not all existing studies rely on a cross-country setting. Davis and Weinstein (1999) use variation in consumption and production shares across Japanese regions.

[9] In a related paper Banerjee *et al.* (2009) use straight line connections between Treaty Ports and historical cities in China to study long term average effects of railway connections in the late 19th and early 20th century on contemporary economic outcomes. This paper's focus is on the role of market size in contemporary trade integration in an emerging market environment subject to road network additions.

study a large scale infrastructure policy to test increasing returns to scale trade theory and learn about the role of market size in trade integration.

The remainder of the paper is structured as follows. Section 2 presents the model and derives a set of qualitative predictions to guide estimation. Section 3 describes the policy background and data. Section 4 presents the empirical strategy. Section 5 reports estimation results. Section 6 presents the quantitative analysis and counterfactual estimations. Section 7 concludes.

2 The Model

The model is based on Helpman and Krugman (1985) and adapted to a setting with multiple and *ex ante* asymmetric regions. In addition, I introduce capital as an input to industrial production and allow this factor to be mobile across regions as in Martin and Rogers (1995). As discussed in more detail below, this serves to adapt the original cross-country model without factor mobility to a within country setting with partial factor mobility without altering the original set of microeconomic forces at play.

The economy is populated by a continuum of agents who are distributed over R regions. There are two sectors of production, labelled agriculture (A) and industry (M), and two factors of production labelled labor (L) and capital (K). The former is assumed to be immobile across regions, while the latter is mobile. Mobile stocks of capital are owned by workers, and returns to capital are repatriated across regions.

2.1 Preferences

The representative consumer in each region has two-tier preferences, where the upper tier is a Cobb-Douglas nest of consumption of agriculture (which will be the numeraire good) and a composite of industrial varieties. Industrial goods enter as a constant elasticity of substitution (CES) sub-utility function defined over a continuum of industrial varieties $i (i = 1, 2, ...N)$. Consumer utility in region $j (j = 1, 2, ...R)$ is given by:

$$U_j = C_{Mj}^{\mu} C_{Aj}^{1-\mu} \quad C_{Mj} = \left(\int_{i=0}^{N} c_{ij}^{1-1/\sigma} di \right)^{\frac{1}{(1-1/\sigma)}} \quad 0 < \mu < 1 < \sigma \tag{1}$$

C_{Mj} and C_{Aj} are consumption of industry and agriculture in region j respectively, c_{ij} is consumption of manufacturing variety i in region j, μ is the expenditure share on industry, and σ is the elasticity of substitution between varieties. Standard utility maximization yields a constant division of expenditure between sectors and CES demand for an industrial variety i in region j:

$$c_{ij} = \frac{p_{ij}^{-\sigma}}{\int_{i=0}^{N} p_{ij}^{1-\sigma} di} \mu Y_j \tag{2}$$

Y_j is total regional factor income of labor (L_j) and capital (K_j), with wage rate w_j and capital return π_j:

$$Y_j = w_j L_j + \pi_j K_j \tag{3}$$

2.2 Technology

The numeraire agricultural sector requires a_A units of labor to make one unit of A. It is subject to perfect competition, constant returns to scale and faces no trade costs. Marginal cost pricing implies that $p_{Aj} = a_A w_j$ and costless trade equalizes prices and wages across regions so that $p_{Aj} = p_A$ and $w_j = w$ as long as some positive fraction of A is produced in every region.[10]

The industrial sector M is subject to increasing returns, Dixit-Stiglitz monopolistic competition and iceberg trade costs. Each firm of a continuum of industrial producers requires one fixed cost unit of capital K, and a_M units of L to produce a unit of M. This implies a cost function $\pi + w a_M x$, where x is firm level output. It is costless to ship industrial goods within a region, but $\tau_{jk} - 1$ units of the good are used up in transportation between two regions j and k. It is assumed that $\tau_{jk} = \tau_{kj}$. It proves convenient to define $\phi_{jk} = \tau_{jk}^{1-\sigma}$ as the "freeness" of trade ranging from 0 (prohibitive costs) to 1 (costless trade).

Dixit-Stiglitz monopolistic competition and demand in (2) imply that mill pricing is optimal for industrial firms, so that the price ratio of a variety in an export region k over its local market price in j is τ_{jk}. For a variety i produced in region j but also sold in another region k this is:

$$p_{ij} = \frac{w a_M}{1 - 1/\sigma}, \quad p_{ik} = \tau_{jk} \frac{w a_M}{1 - 1/\sigma} \tag{4}$$

[10] The Online Appendix derives the formal condition of incomplete specialization across regions.

2.3 Equilibrium

Because the marginal cost of industrial firms depends on the immobile factor whose price is pinned down by costless trade in the numeraire sector, industrial f.o.b. prices are equalized across regions and consumer prices differ only by transport costs. As capital enters as fixed cost component in industrial production, this also implies that capital returns are equal to the operating profit of a typical variety. Under Dixit-Stiglitz competition, this is equal to the value of sales divided by σ. Normalizing the price of agriculture to be the numeraire and choosing units of A such that $p_A = a_A = w = 1$, we can use demand in (2) and mill pricing in (4) to solve for the equilibrium returns to the mobile factor:

$$\pi_j = \left(\sum_k \frac{\phi_{jk} S_{Yk}}{\sum_m \phi_{mk} S_{Nm}} \right) \frac{\mu}{\sigma} \frac{Y}{K} \tag{5}$$

S_Y represents regional shares of total expenditure, and S_N are regional shares of the mass of total industrial varieties. Y and K stand for total expenditure and the total capital endowment across all regions respectively. Given repatriation of capital returns to immobile owners, regional expenditure shares are a deterministic function over regional shares of capital owners and labor endowments, S_K and S_L respectively:

$$S_{Yj} = \left(1 - \frac{\mu}{\sigma}\right) S_{Lj} + \frac{\mu}{\sigma} S_{Kj} \tag{6}$$

Because capital is freely mobile across regions, there are two possible types of equilibria: core-periphery outcomes where S_N can be 0 or 1, and interior location equilibria. Given all regions maintain some positive fraction of industrial activity, capital returns are equalized so that the long run equilibrium location condition is given by $\pi_j = \pi$ for $0 < S_{Nj} < 1$. The profit equation (5) coupled with inter-regional profit equalization yield a system of R equations that can be solved for an $R * 1$ vector of regional industrial production shares as a function of an $R * R$ bilateral trade cost matrix and an $R * 1$ vector of regional expenditure shares that are in turn determined by regional endowments as stated in (6).

2.4 Discussion

Equilibrium profits in (5) are a positive function of access to consumer expenditure, and decreasing in access to competing industrial producers. The interplay between these two channels, where the former enters as agglomeration force and the latter as dispersion force, gives rise to the home market effect in the original two region setting of Helpman and Krugman (1985) as soon as the regional symmetry in expenditure shares is broken.

Considering profits in (5) in the original two region setting, it is readily shown that the direct corollary comparative static result of the home market effect is that falling trade costs between two *ex ante* asymmetric markets lead to an increase of industrial concentration in the larger market.[11] The intuition behind this result is that falling trade costs attenuate both the agglomeration and the dispersion force, but at differential rates. On one hand, lower trade costs decrease the relative disadvantage of higher product market competition in the larger market because the relative increase in competition is stronger for the smaller region.

On the other hand, lower trade costs also decrease the market access advantage of the larger region because the relative increase in market access is stronger for the smaller region. The microfoundation of the home market channel of trade integration is that the decrease in regional differences of market competition outweighs the decrease in market access differentials. Falling trade costs attenuate the dispersion force at a faster rate than the agglomeration force.

The introduction of capital mobility to the Helpman-Krugman model extends this effect to total production rather than just industrial production, without altering the original microeconomic channel at work. Because capital returns are repatriated to immobile owners, capital mobility does not introduce additional self-reinforcing agglomeration forces as in the new economic geography literature (Fujita *et al.*, 1999). It does, however, allow for the possibility that integration leads to reinforced concentration of aggregate economic activity, as regional GDP changes in parallel to industrial production when a factor used in the production of industry is mobile. This relationship is less than proportional because the immobile factor of production remains productive within the region.

[11] Denote a larger "core" region by superscript C and a smaller peripheral region by superscript P. The profit equation for the periphery becomes: $\pi^P = \left(\frac{S_Y^P}{S_N^P + \phi(1-S_N^P)} + \phi\frac{1-S_Y^P}{\phi S_N^P + 1 - S_N^P}\right)\frac{\mu}{\sigma}\frac{Y}{K}$. Given an isomorphic expression for the core region and equalization of capital returns, we can solve for: $S_N^P = \frac{1}{2} + \left(\frac{1+\phi}{1-\phi}\right)\left(S_Y^P - \frac{1}{2}\right)$. The derivative over trade freeness is $\frac{\partial S_N^P}{\partial \phi} = \frac{\phi+1}{(\phi-1)^2}\left(S_Y^P - \frac{1}{2}\right) - \frac{\phi-1}{(\phi-1)^2}\left(S_Y^P - \frac{1}{2}\right)$. This is positive for regions with $S_Y > \frac{1}{2}$, zero for symmetric regions, and negative for regions with $S_Y < \frac{1}{2}$.

Finally, capital mobility enters in combination with immobile labor in the model. In addition to the tractability and closed form solutions that this choice facilitates, there are two deeper reasons. First labor mobility would reinforce the centripetal force of integration already present in the model by adding an agglomeration force of "circular causation" that is not present in the original microfoundation of the home market effect. Second, in the Chinese setting there is a compelling empirical case to be made for a channel of integration that does not rely on perfect labor mobility. The household registration or "hukou" system represents a well documented restriction on migration flows (e.g. Au and Henderson, 2006).

2.5 Qualitative predictions from a three-region policy scenario

This subsection derives a series of closed form qualitative predictions from a tractable policy scenario to motivate the empirical analysis. It is evident from (5) that predictions in a full multi-region setting require the parameterization of the $R * R$ matrix of initial bilateral trade costs, their changes, as well as the vector of initial expenditure shares. This approach will be taken in Section 6 to derive quantitative predictions from the calibrated structural model. Proofs of the following propositions are rendered to the Online Appendix.

Empirical estimations will be based on the comparison of changes of economic outcomes among peripheral county regions that were connected to new NTHS routes relative to non-connected peripheral counties. Given that in general equilibrium it would be a strong assumption that non-connected regions are not affected at all by the network, the most basic policy scenario thus requires at least three regions. Consider two initially identical peripheral regions and one larger metropolitan core region, denoted by superscripts $P1$, $P2$ and C respectively, that are identical in terms of tastes, technology, and initial bilateral trade costs. Geometrically, one can think of this scenario as three regions located on the endpoints of an equilateral triangle. The profit equation in the first peripheral region becomes:

$$\pi^{P1} = \left(\frac{S_Y^{P1}}{S_N^{P1} + (1 - S_N^{P1})\phi} + \phi \frac{S_Y^{C}}{S_N^{C} + (1 - S_N^{C})\phi} + \phi \frac{S_Y^{P2}}{S_N^{P2} + (1 - S_N^{P2})\phi} \right) \frac{\mu}{\sigma} \frac{Y}{K} \tag{7}$$

Profits in the core region are isomorphic, and profits in the second peripheral region are given by $\pi^{P2} = 1 - \pi^{P1} - \pi^{C}$. Initial peripheral symmetry implies that $S_Y^{P1} = S_Y^{P2}$, and ϕ is the identical bilateral trade freeness between all three regions at an initial period. We now introduce asymmetric trade integration in the most convenient way. Let $\alpha\phi$ denote the bilateral trade freeness between peripheral region 1 and the core region after a negative bilateral trade cost shock, while ϕ is the unchanged initial trade freeness between all regions. Initially, $\alpha = 1$, while after the trade cost shock takes effect, α is in the range $1 < \alpha < \frac{1}{\phi}$.[12]

Using peripheral symmetry ($S_Y^{P2} = S_Y^{P1}$), introducing the asymmetric trade cost shock ($\phi^{P1} = \alpha\phi$), and solving for the equilibrium difference of industrial activity between connected and non-connected peripheral regions subject to profit equalization $\pi^{P1} = \pi^{P2} = \pi^C$, we get:

$$S_N^{P1} - S_N^{P2} = \left(\left(\frac{1}{2} - \frac{3}{2} S_Y^C \right) \frac{(\alpha - 1)\phi}{1 - \alpha\phi} + 1 \right) \frac{1 + \alpha\phi - 2\phi^2}{(1-\phi)(1 + \alpha\phi - 2\phi)} - \frac{3\phi}{1 + \alpha\phi - 2\phi} - 1 \quad (8)$$

This provides a closed form solution for peripheral differences in industrial activity as a function of relative market sizes, initial levels of trade costs, and the degree of asymmetric trade integration. At the initial $\alpha = 1$ position, perfect symmetry between peripheral regions leads to $S_N^{P1} - S_N^{P2} = 0$. The question is what happens to industrial production in the connected peripheral county relative to the non-connected one after the trade cost shock materializes. The derivative of interest is $\frac{\partial(S_N^{P1} - S_N^{P2})}{\partial\alpha}$. As can be seen from (8), the sign of this derivative in principle depends on the extent of the pre-existing core-periphery gradient summarized in S_Y^C, the level of pre-existing trade integration ϕ, as well as the extent of asymmetric trade integration captured by α.

It is clear from (8) that for any given scenario of core-periphery integration $1 < \alpha < \frac{1}{\phi}$ and initial trade costs ϕ, the difference in industrial production shares between the integrating and the non-integrating periphery becomes more negative as the core-periphery size asymmetry (summarized by S_Y^C) increases. Using this insight, one can solve for the necessary degree of the core-periphery gradient at which $\frac{\partial(S_N^{P1} - S_N^{P2})}{\partial\alpha} < 0$ holds for any combination of initial trade costs and trade cost shock asymmetry. As long as the metropolitan region is at least twice the size of an individual peripheral region, $\frac{\partial(S_N^{P1} - S_N^{P2})}{\partial\alpha} < 0$ holds across any configuration of ϕ and α.[13] The Online Appendix provides proofs of this result and the following propositions.

[12] This is equivalent to modeling a proportional change of the iceberg trade cost (τ) between the integrating peripheral and the core region. Let $\phi^{P1} = (\widetilde{\alpha}\tau)^{1-\sigma}$, where $\frac{1}{\tau} < \widetilde{\alpha} < 1$, then $\alpha = \widetilde{\alpha}^{(1-\sigma)}$.

[13] Existing theoretical literature on preferential trade integration in a setting with increasing returns to scale production has found increased industrial output in the trading block relative to the non-integrating region (Puga and Venables, 1997). The focus of the present analysis is on the consequences of preferential integration when regions within the trading block are asymmetric. An important feature is the tractability of the present model which allows me to consider a wide range of regional size asymmetries in the context of China. This is in contrast to simulations of infinitessimal deviations from perfect symmetry used in new

Proposition 1 *Falling trade costs between a sufficiently uneven core-periphery pair of regions lead to a reduction of industrial production in the integrating periphery relative to a non-integrating peripheral control region.*

Descriptive statistics in Table 1 give a clear indication that the model's size asymmetry threshold is exceeded when comparing non-targeted peripheral counties to the targeted metropolitan city regions.[14] Following from the previous discussion, the model makes additional predictions about county level changes to overall GDP and agricultural output. Aggregate GDP moves in parallel to industrial output, but less than proportional because labor formerly used by industry remains productive in the region. Proposition 1 thus holds for aggregate production, but we expect a lower point estimate on the elasticity of peripheral GDP to trade cost reductions compared to industrial activity. The reallocation of labor to the agricultural numeraire sector implies that falling trade costs have the opposite effect on agricultural output growth.

Proposition 2 *The negative effect of integration holds, but to a lesser extent, for total regional production, and is reversed in sign for agricultural production.*

In addition to the predictions on the average effects of integration among peripheral counties formalized in Propositions 1 and 2, the richness of the empirical setting also allows to test how the home market channel should affect peripheral counties differently. From (8), the first cross-derivative prediction is that the home market effect should be more pronounced among peripheral counties whose initial level of trade costs vis-a-vis the core region is lower: $\frac{\partial^2(S_N^{P1}-S_N^{P2})}{\partial\alpha\partial\phi} < 0$. For a given trade cost reduction, the marginal effect on industrial and aggregate production should be more negative at higher initial levels of ϕ.

Proposition 3 *The negative effects of integration on industrial and total production are more pronounced among peripheral regions with initially lower trade costs to the larger core region.*

This interaction effect is related to what the trade literature has referred to as home market magnification (Baldwin *et al.*, 2003). The intuition is that falling trade costs attenuate the peripheral location advantage of less market crowding at a faster rate than the

economic geography models.

[14]Beyond the three-region case, all predictions will also be confirmed in the simulation results of the calibrated full multi-region model in Section 6.

metropolitan market access advantage, so that at lower initial trade costs between core and periphery a given trade cost reduction will require a larger relocation of industrial production to equalize the rate of capital return.[15]

The second cross-derivative prediction is that, holding initial trade freeness constant, the home market effect is stronger among peripheral counties whose size differential vis-a-vis the core is more pronounced: $\frac{\partial^2(S_N^{P1}-S_N^{P2})}{\partial\alpha\partial S_Y^C}<0$.

Proposition 4 *The negative effects of integration on industrial and total production are more pronounced among peripheral regions with an initially stronger market size differential to the core region.*

The prediction that the home market channel should operate more strongly among smaller peripheral regions is also intuitive. Falling trade costs weaken the dispersion force at a faster rate than the agglomeration force, so that for a larger core-periphery size gradient, and thus higher initial levels of agglomeration and dispersion forces, a given trade cost reduction requires more industrial concentration in the core to equalize profits. The Online Appendix provides a graphical illustration of Propositions 1-4 for a set of parameter combinations.

3 Background And Data

3.1 China's National Trunk Highway System

In 1992, the Chinese State Council approved the construction of the "7-5" network, consisting of seven horizontal and five vertical axes, under the National Trunk Highway Development Program (Asian Infrastructure Monthly, 1995; World Bank, 2007b) (see Figure 1). The NTHS was constructed at an estimated cost of US$ 120 billion over a 15-year period until the end of 2007, spanning approximately 35,000 km of high speed four-lane highways (Li and Shum, 2001; Asian Development Bank, 2007; World Bank, 2007a).

Its stated objectives were to connect all provincial capitals and cities with an urban registered population above 500,000 on a single expressway network, and to construct routes between targeted centers and the border in border provinces as part of the Asian Highway Network. NTHS routes are four-lane high speed limited access toll ways. The speed limit

[15] Notice that this result is not driven by the choice of modelling trade integration in terms of a proportional reduction in τ. All predictions hold when considering absolute changes in τ instead, but the former are more convenient for deriving closed form solutions.

is 120 km/h, and a common minimum speed limit is 70km/h. Road quality, congestion, and driving speed of the modern expressways are in clear contrast to pre-existing national highways (speed limit 100 km/h) and provincial highways (speed limit 80 km/h) that are also subject to road tolls.

The network was originally earmarked for completion by 2020, but was completed ahead of schedule by the end of 2007. Planners at the Chinese Ministry of Communications divide the construction into a "kick-off" phase between 1992-1997, and "rapid development" between 1998-2007 (World Bank, 2007a). The reason behind the acceleration of construction efforts in 1998 is that highway construction became part of the government's stimulus spending after the Asian financial crisis (Asian Development Bank, 2007).

To finance the great majority of NTHS routes, the central government encouraged province and county level governments to raise funds by borrowing against future toll revenues. Roughly 70% was financed from province and county level debt, and 10-15% was contributed by the central government. Private sector participation was also encouraged with up to 5% of financing stemming from domestic and foreign investors (Asian Development Bank, 2007).

Construction was undertaken almost entirely by Chinese state owned enterprises, part of which were assigned directly to particular localities, part of which were participating in contract auctions.[16] Given the progress in the construction of the NTHS ahead of plan, the State Council approved an even more ambitious follow-up blue print for highway construction in 2004. The so called "7-9-18" system has the stated objective to connect all cities with an urban registered population of more than 200,000. It is scheduled to be completed by 2020.

3.2 Data

Geo-referenced administrative boundary data for the year 1999 was obtained from the ACASIAN Data Center at Griffith University in Brisbane, Australia. These data provide a county-level geographical information system (GIS) dividing the surface of China into 2341 county level administrative units, 349 prefectures, and 33 provinces. Chinese administrative units at the county level are subdivided into county level cities (shi), counties (xian), and urban wards of prefecture level cities (shixiaqu).

County level socioeconomic records are taken from Provincial Statistical Yearbooks for

[16] Until the "Measures on Tenders and Bids for Contracts for Construction Projects" came into effect in May 2003, competitive bidding, was recommended but not mandatory (World Bank, 2007a).

the years 1990, 1997 and 2006, as well as the 1990 Chinese population census. The statistical yearbook records for 1997 and 2006 were obtained from the University of Michigan's China Data Center, and the 1990 census data as well as statistical yearbook data for 1990 were obtained from the China in Time and Space (CITAS) project at the University of Washington. The Provincial Statistical Yearbook series report county level GDP broken up into agriculture, industry, and services gross value added, as well as local government revenues and registered county populations. The 1990 Population Census provides county level data on population, education, and employment shares by sector.

These sources result in a database of 1748 historically consistent geo-referenced county units that have non-missing reporting values in the Provincial Statistical Yearbooks of 1997 and 2006 (75% of Chinese administrative units). Close to the entirety of this county sample (1706 of 1748) also report socioeconomic records in the 1990 Population Census, and 1238 of the 1748 report local government revenues in the CITAS Provincial Statistical Yearbooks for the year 1990.[17] Table 1 presents a set of descriptive statistics, and the Online Appendix describes the data sources and processing in more detail.

Geo-referenced NTHS highway routes as well as Chinese transport network data were obtained from the ACASIAN Data Center. NTHS highway routes were digitized on the basis of a collection of high resolution road atlas sources published between 1998 and 2007. These atlas sources made it possible to classify NTHS segments into three categories that coincide with the main construction phases described by the Ministry of Communications: opened to traffic before mid-1997 (10% of NTHS), opened to traffic between mid-1997 and end of 2003 (81% of NTHS), and opened to traffic after 2003 (9% of NTHS).[18] A list of the atlas publications as well as a more detailed description of the data processing and NTHS classifications is given in the Online Appendix. Finally, land cover and elevation data that are used in the construction of least cost path highway routes were obtained from the US Geological Survey Digital Chart of the World project, and complemented by higher resolution Chinese hydrology data from the ACASIAN data center.

[17] Only a fraction of the reporting counties in 1997 and 2006 report production data in the Provincial Statistical Yearbooks for 1990.

[18] The available series of atlas sources did not allow to date the opening to traffic of each segment of the 35,000 km NTHS road network. See the Online Appendix for a listing of the atlas publications.

4 Empirical Method

I use the collection of data described in the previous section to estimate the effects of NTHS network connections among peripheral counties between 1992-2003 on changes of economic outcomes between 1997-2006. Given that 89% of the reported NTHS connections until the end of 2003 were completed during the phase of "rapid development" between 1998-2003, the main source of variation used in the estimations stems from network connections during this five year period 1998-2003 and their effects on changes of economic outcomes over the nine year period 1997-06. The baseline estimation strategy is a difference in differences specification of the form:

$$\ln(y_{ip}^{2006}) - \ln(y_{ip}^{1997}) = \alpha + \gamma_p + \beta Connect_{ip} + \eta X_{ip} + \epsilon_{ip} \ , \tag{9}$$

where y_{ip} is an outcome of interest of county i in province p, γ_p is a province fixed effect, $Connect_{ip}$ indicates whether i was connected to the NTHS between 1992-2003, and X_{ip} is a vector of control variables. I classify highway connections using GIS with a dummy indicator that takes the value of one if any part of county i is within a 10km distance of a NTHS highway that was opened to traffic before the end of 2003. Alternatively, I run specification (9) with a continuous treatment variable, $lnDistHwy_{ip}$, which stands for the logarithm of great circle distance to the nearest NTHS highway segment opened to traffic before the end of 2003, measured from the center of each county unit.

The error term ϵ_{ip} could be correlated across counties that were connected to a similar part of the network during a similar period between 1992-2003. I therefore cluster standard errors at the level of 33 Chinese provinces. Alternatively, I follow Conley (1999) and allow for spatial dependence to be a declining function over bilateral county distances without imposing parametric assumptions. Finally, due to the fact that the explicitly targeted network nodes are China's largest city regions that encompass multiple county level units, I exclude county observations within a 50 km commuting radius around the targeted city centers.[19]

4.1 Least cost spanning tree networks

Estimating specification (9) by OLS would imply the assumption that county connections between nodal cities were randomly assigned within provinces. Given the policy setting of

[19] See Garske et al. (2011) for a study of commuting patterns and distances in China.

the NTHS, this assumption would be strong. The NTHS was planned in 1992 to establish the backbone of a modernized road transport system for China. Province and county governments borrowed against future expressway toll revenues to finance its construction. This background raises the concern that planners targeted politically important and economically prosperous regions on the way between the network's targeted destinations. This concern is supported by descriptive statistics presented in Table 1. Peripheral counties connected to the network by 2006 were on average larger, richer, more urbanized, and more industrialized than non-connected peripheral counties before the bulk of the network had been built in 1997.[20]

To address these concerns, I construct two hypothetical minimum spanning tree highway networks as instruments for actual route placements (see Figures 2 and 3). I refer to the first as least cost path spanning tree network, and to the second as Euclidean spanning tree network. Both instruments correspond to the question of which routes central planners would have been likely to construct if the sole policy objective had been to connect all targeted destinations on a single network in a least costly manner. The least cost path network yields more precise route predictions between any given bilateral connection on an all-China minimum spanning tree due to its use of land cover and elevation data, while the Euclidean network covers a larger set of the actually built bilateral network routes.

The least cost path network depicted in Figure 2 is constructed in a two step procedure. The first step is to compute least cost highway construction paths between all possible targeted destination pairs on the basis of remote sensing data on land cover and elevation. To this end, I adopt a simple construction cost function from the transport engineering literature that assigns higher construction costs to land parcels with steeper slope gradients and land cover classified as water, wetlands, or built structures (Jha *et al.*, 2001; Jong and Schonfeld, 2003).[21] I use the remote sensing data to create a construction cost surface covering the PR China in a rectangular grid of cost pixels (see Online Appendix for details and illustrations). I then implement Dijkstra's optimal route algorithm to construct least cost highway construction paths between all possible bilateral city connections as well as provincial capitals and the border in border provinces. In the second step, I extract the estimated aggregate construction cost of each bilateral connection and feed them into Kruskal's minimum span-

[20] Reported differences are statistically significant at the 1% level.

[21] As discussed further below, I will also include these geographical characteristics used in the construction of the instrument as direct county level controls to address the concern that these might affect changes in economic outcomes directly.

ning tree algorithm. This algorithm yields the minimum number of least cost connections (i.e. number of targeted nodes minus one) to connect all targeted destinations on a single continuous network to minimize aggregate construction costs.

To construct the Euclidean spanning tree network depicted in Figure 3, the first step is to compute great circle distances between all possible bilateral connections of the network, which is done by applying the Haversine formula to bilateral coordinate pairs as well as provincial capitals and the border in border provinces. I then compute Kruskal's algorithm to identify the minimum number of edges that connect all targeted destinations subject to the minimization of total network distance. To compensate for the loss of route precision, I account for the fact that Chinese planners constructed many more than the minimum number of spanning tree connections. I therefore re-run Kruskal's algorithm within separate geographical subdivisions after dividing China into either North-Center-South or East-Center-West geographical divisions. These six additional spanning tree computations add nine bilateral routes in addition to the single all-China minimum spanning tree. The Online Appendix provides further details and additional illustrations of these computations.

4.2 Additional county controls and identifying assumption

The minimization of a network construction cost objective function from which the instruments in Figures 2 and 3 are derived is aimed to address the concern of non-random highway placements on the way between targeted destinations. However, the exclusion restriction could be violated if locations along least cost road construction paths between major economic centers in China are correlated with economic county characteristics due to history and sorting. Furthermore, the instrument is likely to be mechanically correlated with distance to the nearest targeted metropolis. I therefore estimate regressions before and after including a set of additional county controls that could be correlated with the instrument while also affecting the change of economic outcomes between 1997-2006.

Counties closer to targeted city centers are mechanically more likely to lie on a least cost spanning tree path than counties situated farther away. Concerns about the exclusion restriction arise if distances to the major cities of China are correlated with economic county characteristics that also affect growth trajectories. I include the log distance between counties and the nearest targeted metropolitan city center to address this concern.

Conditional on county distance to the targeted centers, location on least cost road construction paths between major economic centers in China could be correlated to political and economic county characteristics due to historical trade routes. To address such concerns, I include a set of observable controls for pre-existing county level political status and economic conditions. The two political controls are dummy variables indicating whether the county seat was a prefectural capital or has city as opposed to township status in 1990. The concern is that higher administrative status might be historically concentrated along least cost path routes between important economic centers.

Concerning pre-existing economic conditions, I use data from the 1990 Census at the county level which allows me to compute the share of agricultural employment in total county employment, the logarithm of county level urban registered population, as well as the share of above compulsory schooling attainment in total county population above 20 years of age in 1990.[22] These controls are aimed to address concerns that counties along least cost connections between major cities differ in terms of both their economic composition (shares of skilled labor and sectoral specialization), as well as in their mass of economic activity (urban populations).

The baseline identifying assumption is that county location along an all-China least cost spanning tree network affects changes in county level economic outcomes only through NTHS highway connections, conditional on province fixed effects, distance to the nearest targeted city region, administrative status and county-level economic conditions in 1990. I address a series of robustness checks after reporting baseline estimation results.

5 Estimation Results

5.1 Average network connection effects

This subsection reports the empirical results and robustness checks on Propositions 1 and 2 of the theory section. Table 2 presents the first stage results for the least cost path and the Euclidean network instruments as well as their combined first stage results. First stages are run for binary NTHS connection indicators as well as the log distance to the nearest NTHS segment. Both the least cost path and the straight line networks are strongly significant

[22] Categories beyond the compulsory 9-year curriculum are senior middle school, secondary technical school, technical college, junior college and university.

within province predictors of actual NTHS placements conditional on log distance to the nearest targeted node and the full set of pre-existing political and economic county characteristics. County controls are related to NTHS exposure mostly as expected. NTHS route connections are more likely for counties with lower distances to the targeted metropolitan city centers, larger pre-existing urban populations, and city status.[23]

Both instruments remain statistically significant when included simultaneously, confirming that the two spanning tree networks capture slightly different sources of the increased likelihood of route placements. While the least cost path instrument is a more precise predictor of placements on any given bilateral connection, the Euclidean tree instrument captures a higher proportion of the actually built network connections.

Table 3 presents OLS and IV results when regressing log changes of county level outcomes on the binary network treatment variable before and after including the full set of 1990 county controls. The Online Appendix reports results after replacing the binary network treatment with the log county distance to the nearest NTHS segment. The results discussed in this section are confirmed (in opposite sign as expected) in these continuous NTHS connection specifications.

The instrumental variable estimates of the NTHS connection effect in Table 3 are negative and statistically significant for industrial output growth, non-agricultural output growth, local government revenue growth, as well as total GDP growth. Two important patterns emerge. The first is that the OLS point estimates are less negative than the IV estimates. The second is that the inclusion of additional county controls for pre-existing political status and economic conditions leads to more negative point estimates of the NTHS connection effect.

The first of these is in line with the discussed concern that planners targeted places with higher expected returns to infrastructure investments and/or higher expected traffic demand, which is also apparent in the descriptive statistics of Table 1. The second attests the conditionality of the exclusion restriction on controlling for pre-existing political and economic county differences. The results in Table 3 suggest that county location along least costly

[23] When individually included, higher pre-existing shares of agricultural employment decrease the likelihood of route placements, and higher shares of educated population increase it. These correlations are no longer significant when both controls are added simultaneously as reported in Table 2. Finally, the identifier for prefecture level capital status in 1990 enters in opposite sign than expected (decreased likelihood). This is due to the simultaneous inclusion of the city identifier, so that the coefficient is driven by approximately 1% of relatively remote prefecture level capitals that are not also classified as cities in 1990.

road construction paths between major cities in China is at least partly correlated with pre-existing county characteristics, such as the size of urban populations, the share of educated labor, and the degree of industrialization. As noted in the previous subsection, these correlations could be driven by settlements and sorting along historical trade routes. However, the finding that conditioning on pre-existing county differences leads to more negative connection treatment effects on industrial output growth, total output growth, as well as local government revenue growth suggests that these characteristics are positively associated with economic growth, rather than negatively.

Nevertheless, the sensitivity of the IV point estimates to the inclusion of county controls in principle raises the concern that the estimated effects remain biased in either direction due to omitted unobserved differences that are correlated with the instrument. A related concern is that counties along an all-China spanning tree network had different pre-existing growth trends before the highway network came into effect. To address such concerns, I make use of the fact that the majority of the reporting county sample in 1997 and 2006 also reported local government revenues in the Provincial Statistical Yearbooks for the year 1990. If the exclusion restriction is satisfied conditional on the included county controls, then we should expect to find no significant relationship between NTHS treatments and local government revenue growth prior to the network, when estimated on the identical county sample for both periods.

Table 4 presents OLS and IV results for both instruments in both periods 1990-1997 and 1997-2006. For completeness, the table also reports results for the continuous NTHS exposure variable measured by the log county distance to the nearest NTHS route for both periods. The county sample is smaller than in the previous regressions firstly because not all reporting counties in 1997-2006 have non-missing entries for local government revenue in 1990, and secondly because these regressions exclude counties that were connected to NTHS routes built between 1992-1997 (10% of the NTHS).

The connection indicator enters negatively and statistically significantly only for the NTHS period, and the log distance to the nearest NTHS segment enters positively and statistically significantly only for the NTHS period. Furthermore, these outcomes are not driven by differences in the size of the standard errors of the point estimates across the two different periods, but by changes in the point estimates themselves. These results provide a reassuring robustness check of the exclusion restriction conditional on the included set of pre-existing political and economic county controls. A series of additional robustness checks will be discussed at the end of this subsection, and the following subsection tests the model's predictions on the heterogeneity of the NTHS connection effects.

According to second stage IV results using both spanning tree instruments and the full set of county controls as reported in the final column of Table 3, NTHS connections on the way between targeted destinations have on average reduced GDP growth by about 18 percent over a nine year period between 1997-2006 compared to non-connected peripheral counties. Local government revenue growth is reduced by approximately 23 percent. These adverse growth effects appear to be mainly driven by a decline in industrial output growth of approximately 26 percent over the nine year period.[24]

Results in Table 3 for agricultural GDP growth are close to zero and not statistically significant. This finding is in contrast to the model's prediction about the reallocation of labor from industry to the agricultural sector. This result could be due to labeling both more and less industrialized activities as agricultural in county level economic accounts. An economic explanation of this finding could be related to factor market rigidities or adjustment costs and the frequent empirical finding that the inter-sectoral reallocation of resources following trade shocks fails to be confirmed in the data (Goldberg and Pavcnik, 2007).

The final result from Table 3 is that NTHS connections had no effect on county population growth. The point estimates are close to zero and statistically insignificant conditional on controls. This result is consistent with the above discussion of Chinese migration controls and suggests that the stated output growth effects are not driven by significant differences in population growth across counties. The following subsection on the interaction effects of NTHS connections also addresses a further robustness check concerning the possibility of unreported outmigration that could remain uncaptured when estimating the NTHS effect on registered county populations.

In addition to the results reported in this section, the Online Appendix includes a series of additional robustness checks concerning the average NTHS connection effects on production and government revenue growth. These address concerns about i) controlling for potential direct effects of the land cover and elevation characteristics used in the construction of the least

[24]The cited estimates correspond to point estimates as shown in the regression tables after converting log points back to percentage changes and rounding.

cost path instrument, ii) adjusting standard errors for spatial dependence following Conley (1999) instead of province level clustering, iii) controlling for construction activity already underway in 1997, iv) excluding mountainous regions due to least cost path endogeneity concerns, v) excluding the Golmud-Lasa railway route completed over the same period, and vi) controlling for proximity to historical trade routes such as the Silk Road. The results discussed in this section are robust in their magnitude and statistical significance in these additional specifications.

The Online Appendix also provides a discussion and estimation results concerning the proportion and observable characterization of the complier group of counties that drive the estimated local average treatment effects. Descriptive statistics and the pattern of coefficient estimates discussed above suggest that planners targeted economically prosperous counties on the way between targeted city regions. The additional estimations address the concern that least cost spanning tree location might have affected actual highway placements only for a subset of remote and economically stagnant counties on the way between targeted nodes, so that the estimated local average NTHS connection effects might systematically differ from population average effects. The results presented in the Online Appendix provide evidence against this concern, showing that the predictive power of the instruments does not significantly vary across observable pre-existing county characteristics. The reasons behind this finding are linked to the nature of the spanning tree prediction errors compared to actual NTHS route placements which are further explored in reference to a set of cartographic illustrations.

5.2 Interaction effects and alternative channels

This subsection reports the empirical results and robustness checks on Propositions 4 and 5 of the theory section. To test the model's additional predictions on the heterogeneity of the network's effects among non-targeted peripheral counties, I estimate specification (9) after including interaction terms between the NTHS connection indicator and the county log distance to the nearest targeted metropolitan node, as well as an indicator for counties with above mean 1990 total employment sizes. The log distance to the nearest metropolitan node is aimed to capture the initial degree of trade freeness between core and periphery (Proposition 3). The categorization of large counties in terms of 1990 employment sizes is aimed to capture the pre-existing degree of the core-periphery size differential (Proposition 4). Alternatively, the binary NTHS connection indicators in these interactions are replaced by the continuous treatment variable $lnDistHwy_{ip}$ as in the previous subsection.

Table 5 reports second stage least squares results for industrial output growth as well as aggregate GDP growth as dependent variables after instrumenting for NTHS connection and its interaction terms with the least cost path network instrument.[25] The first column for each dependent variable reproduces the average treatment estimate conditional on the full set of county controls as reported in Table 3. The second columns then introduce the additional two interaction terms. In the binary network connection specifications both interaction terms enter positively and statistically significantly for both industrial and aggregate output growth. When using the log distance to the nearest NTHS segment as continuous treatment variable instead, the interaction terms enter negatively and statistically significant for both dependent variables.

A potential concern with these two stage least squares results is that the first stage F-statistics significantly drop after instrumenting for NTHS treatments as well as its interactions. The Online Appendix reports an additional set of results that compare two stage least squares (2SLS) estimates using both instruments to estimations by limited information maximum likelihood (LIML). The fact that the variation of first stage F-statistics between the different specifications across Columns 2-4[26] in Table 5 have little effect on the point estimates of the two interaction terms of interest, and the fact that the reported LIML point estimates are slightly higher than the 2SLS estimates reported in the Online Appendix indicate that weak instrument bias is unlikely to be a concern.[27]

The reported results provide empirical evidence for significant heterogeneity in the NTHS connection effects among peripheral counties that confirms the additional cross-derivative predictions of the model. In particular, while the majority of both large and small peripheral counties are estimated to have experienced negative GDP growth effects, a subset of large peripheral counties that are also subject to large initial trade costs with respect to the targeted metropolitan city centers are estimated to have experienced positive growth effects due to

[25] In the cross-derivative regressions the least cost path spanning tree instrument performes better in terms of first stage predictive power than the Euclidean instrument or both instruments combined. The Online Appendix reports these specifications using both instruments instead.

[26] The final column is not estimated on the identical county sample.

[27] See for example Angrist and Pischke (2008, Section 4.6) for a discussion 2SLS and LIML estimates in the context of weak instrument concerns.

NTHS connections.[28]

In the following, I consider the interaction effects of peripheral NTHS connections in the light of alternative channels. While the reported negative average effects on not just sectoral but aggregate output growth are *a priori* difficult to reconcile in absence of the home market channel, one might nevertheless be concerned that the observed average effects and their interactions are partly driven by alternative microeconomic channels.

A first remaining concern could be that county size and/or distance to the nearest city center are correlated with relative production costs across counties linked to neoclassical channels of trade. In particular, relative cost (dis-)advantages in industrial production could be correlated with these characteristics. Columns 3-5 in Table 5, report results after including a series of additional interaction terms with respect to pre-existing political and economic county characteristics. Reported estimates now control for the heterogeneity of the NTHS connection effect on county level industrial or aggregate output growth with respect to 1990 differences in the share of skilled population, the share of agricultural employment, county political status as a city or a prefecture level capital, and local government revenue per capita. These additional interaction terms are aimed to capture differences in relative factor prices or technologies across counties that could be correlated to market size or distance to the nearest targeted city regions.

The finding that the estimates on the interactions between NTHS treatments and metropolitan proximity and county size are confirmed in statistical significance is reassuring. Furthermore, the point estimates of these coefficients are hardly affected by the inclusion of the additional interaction terms across Columns 3-5. This provides support for the interpretation that the estimated reduced form effects are driven by the home market channel, rather than capturing alternative relative cost determinants of integration.

A second remaining concern could be that NTHS connections have reduced the costs of unregistered outmigration to larger metropolitan markets. Despite the fact that we find no significant migration response from registered population records, one could suggest that unregistered outmigration happens on a large scale and is significantly affected by NTHS network connections. The negative growth effects in this setting could then be driven by an unregistered reduction in employment which reduces both industrial production and GDP.

[28] The median and mean distances to the nearest trageted city node were 168 km and 203 km respectively.

To assess this explanation against the proposed theoretical channel, I consider the interaction effects of NTHS connections that would arise under such a channel.

The concern would be that unregistered migrants move to equalize real wages across counties net of the costs of migration, metropolitan centers have higher real wages, and NTHS connections tip the balance for a larger number of peripheral citizens to outmigrate compared to non-connected counties. Similar to the alternative channels considered above, a testable distinction compared to the home market channel is that interaction effects should be driven by regional factor returns, and not by market size or the initial trade cost position. In particular, one would expect the interaction point estimates with respect to county size and metropolitan distance to be sensitive to the inclusion of the additional interaction with respect to the log of local government revenue per capita reported in 1990 which serves as a proxy for regional income per capita differences.[29] The results reported in Table 5 prove robust to the concern of large scale unregistered outmigration responses.

6 Structural Estimation And Results

The first objective of this section is to assess to what extent the Helpman-Krugman microfoundation of the home market effect can quantitatively account for the observed reduced form effects. To this end, I calibrate the model to fit the county level distribution of industrial activity in 1997, and simulate the network's effects subject to observed reductions in bilateral county trade costs from Chinese transport network data. I then document simulation results across a wide range of parameter combinations to assess the robustness of the qualitative predictions from the stylized three-region model in the full multi region setting, and characterize the parameter space that best accounts for the observed reduced form effects. The second objective is to use the best fitting parameterization of the structural model in counterfactual policy estimations to learn about the network's effect on the concentration of economic activity in China as well as its welfare implications.

[29] Only a fraction of counties reporting in 1997-2006 reported GDP in the Provincial Statistical Yearbooks for 1990.

6.1 Calibration and simulation

6.1.1 Calibration to 1997 distribution of industrial production

The profit equation in (5) coupled with capital return equalization and the restriction that regions retain some positive fraction of production in both sectors yield a system of R equations that can be solved for an $R * 1$ vector of regional industrial production shares as a function of an $R * R$ bilateral trade cost matrix and an $R * 1$ vector of initial expenditure shares which are in turn determined by endowment shares as presented in (6).

Given 1997 levels of bilateral trade freeness parameters, $\tau^{1-\sigma} = \phi$, and the observed 1997 county level distribution of industrial output, I can calibrate unobserved county level expenditure shares to fit the distribution of observed industrial activity before the majority of the NTHS network were built.[30] One advantage of this calibration approach is that the calibrated regional expenditure shares will capture not just domestic market access, but also county level differentials in cross-border market access as the observed distribution of industrial activity in 1997 reflect the sum of effective market access.

To be able to compute the 1997 expenditure share vector, one needs to compute the pre-existing matrix of bilateral trade costs. I follow the standard approach from the gravity literature in international trade and model initial levels of transport costs as a function of bilateral distance between two counties j and k, $\tau_{jk} = D_{jk}^{\delta}$. This approach requires two parameter choices in the calibration exercise. The first parameter is the elasticity of transport costs with respect to distance (δ). The second parameter is the elasticity of substitution between varieties (σ). The choices of these two parameters jointly determine the distance elasticity of trade between bilateral county pairs: $\delta(1-\sigma)$. As described below, I consider a wide range of parameter combinations.

The two parameter choices allow me to compute the $R * R$ matrix of bilateral trade freeness parameters on the basis of great circle distances between all bilateral county pairs. Combined with the observed initial distribution of county level industrial activity, this allows me to calibrate the model to fit the pre-existing county level distribution of industrial activity and compute the vector of pre-existing expenditure shares.

Counties with 1997 statistical yearbook observations on industrial output yield an all-

[30] Regional expenditure shares are the sum of immobile labor income and returns to immobile owners of mobile capital.

China estimation sample of 1679 regions. In line with the empirical analysis presented above, each of the 54 targeted metropolitan city regions have been aggregated to a single region encompassing a 50 km commuting buffer around the city center. The 1679 estimation sample thus contains 54 targeted city regions, and 1625 non-targeted peripheral regions.

6.1.2 Simulation of NTHS effects subject to observed trade cost reductions

Given initial bilateral trade costs and the initial county vectors S_Y and S_N, I can simulate county level economic changes due to the NTHS system subject to observed bilateral trade cost changes. To estimate the $R * R$ matrix of bilateral trade cost changes due to the addition of new NTHS routes, I use road network data provided by the ACASIAN data center to create a geographical information system of China's primary and secondary road network for mid-1997 (before the bulk of expressways were built) and at the end of 2003 (after the bulk of the expressway network was built).[31] This allows me to apply a minimum cost routing algorithm on both network datasets to extract bilateral transport cost measures between all possible bilateral county pairs before and after the network was built.

I model the trade cost reduction effects of new NTHS route additions in terms of bilateral transport cost percentage changes that I can compute from the GIS network database and the optimal routing algorithm. In theory, new route additions will affect the iceberg trade costs τ_{jk} in the trade freeness parameter $\phi_{jk} = \tau_{jk}^{1-\sigma}$. Iceberg trade costs imply that $(\tau_{jk} - 1)$ units of industrial varieties are used up in transport between any given bilateral pair, so that $(\tau_{ij} - 1)$ is the ad valorem trade cost between counties. I use the estimated percentage reduction in bilateral transport costs from the network analysis to obtain an estimate of the change in $(\tau_{jk} - 1)$ between any bilateral county pair.[32]

The network routing analysis to compute bilateral percentage transport cost changes requires assumptions about the relative per kilometer cost of transport across different road types in China. In absence of empirical estimates of these parameters across Chinese provin-

[31] The Online Appendix provides an illustration of the road network used in the computation. The present analysis abstracts from railroad transport, air transport and inland waterways mainly to avoid the need to parameterize relative per km transport costs across modes, and as a changing function over distance for rail, air, and waterways. According to national transport statistics provided by the Chinese National Bureau of Statistics for the year 2000, 77% of Chinese domestic freight was accounted for by road transport (www.stats.gov.cn/english). Since waterways and railways are used for higher weight-to-value ratios, this percentage is likely a lower bound estimate of the relative importance of road transport in domestic trade.

[32] Denoting total bilateral transport costs as TC_{jk}, ad valorem bilateral transport costs for a given industrial variety between bilateral county pair j and k are $(\tau_{jk} - 1) = \frac{TC_{jk}}{p}$, so that $\frac{d(\tau_{jk}-1)}{\tau_{jk}-1} = \frac{dTC_{jk}}{TC_{jk}}$.

cial highways, national highways, and the new NTHS expressways, I take the relative reported speed limits as my baseline estimation. This choice reflects the empirical finding in, for example, Combes and Lafourcade (2005) that time varying costs of road transport, such as drivers' wages, account for a substantial part of total transport costs. Given speed limits of 120 km/h, 100 km/h, and 80 km/h for NTHS expressways, national highways, and provincial highways respectively, the baseline per km cost ratios when choosing national highways as the reference point are approximately 0.85 for NTHS routes, 1 for national highways, and 1.25 for provincial highways. To address the empirical uncertainty of this baseline estimate, each simulation will be run across a band of five equally spaced alternative parameterizations that are centered around the baseline estimate as further discussed below.

These relative per kilometer trade cost parameters are attached to network arcs of different road types in the GIS road network dataset of China both before and after the additional NTHS expressway connections were added over the period 1997-2003. The minimum cost routing algorithm produces bilateral network transportation costs subject to the original network in 1997 (with 10% of the NTHS completed) and subject to the network at the end of 2003 (with 81% of the NTHS completed). The resulting matrix of bilateral percentage cost reductions are then used to compute the $R * R$ matrix of bilateral trade freeness parameters for the post-NTHS network. Given the calibrated S_Y vector and the post-NTHS trade cost matrix, equation (5) yields a system of R equations in R unknown post-NTHS industrial production share outcomes. The solution to this S_N vector of simulated post-NTHS county industrial production shares then yields results for simulated industrial output growth at the county level.

6.1.3 Parameter ranges and number of simulations

The choice parameters for calibrations and simulations are the elasticity of trade costs to distance (δ), the elasticity of substitution between industrial varieties (σ), which together determine the trade cost elasticity of trade ($\delta(1-\sigma)$). In addition, I estimate simulated results for each parameter combination across a range of five parameterizations for the relative per km cost of NTHS routes relative to existing national highways that is centered around the baseline estimate of 0.85. These range from 0.75-0.95 and are equally spaced in deviations of 0.05.

I consider a wide range of parameter combinations with respect to δ and σ, allowing δ to vary between 0.2-1.6 in intervals of 0.02 and allowing σ to vary between 2-10 in intervals of 0.2. Each of these 71*41=2911 simulations is computed across the five alternative parameterizations of relative NTHS transport cost savings. Each of these single simulations yields a vector of simulated county level changes of industrial activity across the 1679 regions.

6.2 Quantitative assessment of the home market channel

The first question of this subsection is how robust the qualitative predictions from the three region policy scenario are in the full multi-region setting with observed bilateral trade cost changes. Figure A.1 depicts regression coefficients for simulated county industrial output growth on the binary NTHS connection treatment dummy estimated for the identical peripheral county sample for which reduced form estimation results are reported in Table 3. The point estimates are derived from the identical regression specification including the full set of 1990 county level controls and province fixed effects. The OLS point estimate of each separate vector of county level simulation results is plotted against the trade elasticity parameter implied by the combinations of δ and σ.[33] The horizontal line indicates the preferred instrumental variable estimate displayed in the final column of Table 3.

The first point to note is that the prediction of a negative NTHS connection effect among peripheral counties (Proposition 1) is confirmed in the full multi-region setting subject to observed trade cost changes. Not surprisingly, for prohibitively high initial cross-county trade costs, NTHS induced trade cost reductions have zero effect on industrial output growth differences between connected and non-connected peripheral counties. Once the initial level of trade costs in the system drops below prohibitive, the home market channel is at work and the core-periphery prediction of the model becomes stronger the lower the initial level of trade costs.

Figure A.2 shifts attention from the predicted average connection effect among peripheral counties to its predicted heterogeneity with respect to pre-existing county characteristics (Propositions 3 and 4). The y-axis displays the difference of simulated average connection effects on industrial output growth between county groups with above and below mean distance to the nearest targeted node, and with above and below mean employment size in 1990. Connection effects are estimated from four separate regressions of simulated outcomes on the

[33] All displayed results correspond to the baseline parameterization of the relative transport cost savings on NTHS routes compared to existing national highways.

NTHS connection dummy and province fixed effects. The x-axis displays the range of trade cost elasticities that yield results in the proximity of the reduced form average connection effect displayed by the horizontal line in Figure A.1. In confirmation of the cross-derivative predictions of the three-region model and the reduced form results presented in Table 5, the NTHS effect is predicted to be more negative among smaller peripheral counties and lower initial trade costs to targeted metropolitan cities.

The second question of this subsection concerns the quantitative performance of the model's home market channel in accounting for the estimated reduced form treatment effects from the IV estimations. To this end, I implement a two-stage parameter grid search across all simulation results. In the first stage, I identify all parameter combinations for which the estimated average treatment effect is statistically significant at the 10% level and within half a percentage point of the estimated reduced form treatment effect from the identical specification and the identical county sample reported in the final column of Table 3. In the second step, I further limit the selection to simulations with statistically significant (10% level) and positive interaction effects with respect to distance to the nearest targeted node and 1990 employment size, and compute the sum of squared deviations of simulated interaction effects from those reported in Column 2 of Table 5. The two-stage procedure places higher weight on matching the average observed connection effect of the network before further selecting on the basis of its heterogeneity across counties.

These narrow search criteria yield three best fitting parameter combinations that are reported in Table 6. The first point to notice is that the best fitting parameter combinations are clearly within the range of commonly found empirical estimates for both the elasticity of substitution and the implied trade elasticity to distance. For the best fitting parameter combination, these respective values are 2.4 and -1.6. For example Redding and Sturm (2008) have obtained the identical parameter estimate of the trade elasticity to distance also in the context of land based transport for German regions. Secondly, the best fitting parameterizations are based on the same baseline estimate of relative transport costs of new NTHS expressways relative to existing national highways. Given that the estimate is based on relative travel times, this is in line with existing empirics on the importance of time costs in road transport.[34]

[34] See for example Combes and Lafourcade (2005).

Finally, the best fitting parameter space depicts a trade off between the two choice parameters that together determine the trade elasticity over distance. Figure A.3 displays the relationship between the trade cost elasticity to distance and the elasticity of substitution when conditioning on the parameter space that yields statistically significant and close to identical average connection effects compared to the IV point estimates in the final column of Table 3. Similar to the findings in Redding and Sturm (2008), the apparent trade off between the two parameters that jointly determine the elasticity of trade to distance suggests that the latter is the binding parameter that governs the model's prediction on the intensity of the effects of a given trade cost shock.

6.3 Counterfactual policy estimations

The final part of the paper uses the best fitting parameterization of the calibrated structural model for the purpose of counterfactual policy estimations. Table 7 presents regression results of simulated industrial output growth on an identifier variable for targeted metropolitan city regions before and after including province fixed effects. While the identification strategy of the preceding empirical analysis was based on variation of highway connections among non-targeted peripheral counties, the structural model allows to analyze the network's effects across the all-China county sample. The model suggests that the NTHS network had a significantly negative average growth effect among all peripheral counties in China, whereas a small but statistically significant positive effect is the case for targeted metropolitan centers.

Table 8 presents estimation results on changes of industrial concentration as measured by a Herfindahl index for the all-China county sample as well as for within peripheral or within metropolitan county groups. Following from the results discussed up to this point, the NTHS is found to have increased the concentration of industrial activity in China. Interestingly, the model suggests that industrial concentration has increased most within the peripheral county group. The first reason is that peripheral regions constitute a more homogenous group prior to the network so that differential output growth rates across counties have a stronger effect on industrial concentration within this group.[35] The second reason is that, as indicated by the estimation results on Proposition 5, initially larger peripheral counties were affected more

[35] In contrast, even strongly reduced peripheral output growth has a minor effect on a national index of Chinese industrial concentration because of the small pre-existing relative mass of industry that peripheral counties represent relative to the metropolitan regions.

positively by new highway connections within the Chinese periphery.

Finally, the structural model can help to shift attention from nominal output growth to county level price index consequences. In the Helpman-Krugman model gains from trade arise purely from increased regional consumer access to industrial varieties captured by the ideal price index.[36] Given preferences in (1), the ideal price index for region j is given by:

$$P_j = \left(\sum_k \phi_{jk} S_{Nk} \right)^{\frac{-\mu}{\sigma-1}} \tag{10}$$

where μ is the expenditure share of industry in consumption. Applying the best fitting parameter combination identified in Table 6 to parameterize σ and the matrix of initial bilateral ϕ_{jk}'s, and including an observed household expenditure share on non-agricultural merchandise of $\mu = 0.5$ from the Chinese input-output table in 2000, I can estimate county level price indices according to (10) on the basis of observed industrial output shares (S_{Nk}) in 1997, as well as for simulated industrial output shares after the NTHS was built.

Table 9 presents OLS regression results for simulated price index changes. The estimated population weighted county level average price index reduction due to the construction of the NTHS network is -1.1% across all counties in China. Evaluated at nominal Chinese GDP in the year 2000, the model thus predicts substantial NTHS induced gains from trade of 10% in annual terms of the total officially stated network construction costs (120 billion US dollars).

Columns 3 and 4 then shed light on the welfare distributional predictions of the parameterized model. The first point to notice is that all county groups are predicted to have benefited from the NTHS in terms of local price index reductions for industrial varieties. This finding is a result of two counteracting forces among peripheral counties. On average, peripheral counties are predicted to have experienced reductions of industrial output growth which adversely affects the local price index (to increase). On the other hand, peripheral counties have gained access to the mass of metropolitan industrial varieties that is many times larger than their individual shares of production. The predicted net result of these two forces is that all counties have benefited from greater access to industrial varieties.

The second point is that the group of connected peripheral counties who have experienced the largest adverse production effects is also the group with the highest predicted average price

[36] As discussed in the theory section, nominal wages and capital returns are invariant to trade cost changes. Repatriation of returns to mobile capital and costless reallocation of labor to the freely traded numeraire sector imply that local price index changes represent a sufficient welfare statistic in the model.

index reductions. The important insight of the structural model is that the observed increase in the concentration of industrial and aggregate economic activity does not necessarily imply parallel welfare distributional effects across regions. The reason is that small peripheral production centers obtain cheaper access to the bulk of the industrial production mass located in the core production regions, so that the effect on local price indices is likely to be strongest for the connected periphery.

In contrast to the network's effects on county level economic activity, the absence of data on local price indices and household incomes in Provincial Statistical Yearbooks prevents the comparison of the model's predictions to reduced form empirical evidence on the NTHS welfare effects. In this context, it is important to underline that the model's welfare predictions abstract from other potential channels of gains from trade, such as comparative advantages or productivity effects on one side, and on the other side abstract from potential impediments to the estimated gains such as factor reallocation costs or decreasing returns to scale in agriculture. While the counterfactual estimations of the model provide important additional insights, the absence of supporting reduced form evidence should serve as a cautionary note against firm conclusions on the welfare effects.

7 Conclusion

Increasing returns to scale trade theory makes important predictions about the consequences of trade integration between regions with uneven market sizes. The home market effect in Krugman (1980) and Helpman and Krugman (1985) provided a microfoundation for the proposition that market size is a determinant of industrialization. The same channel also suggests that falling trade costs between asymmetric regions can reinforce the concentration of economic activity. This paper exploits China's NTHS as a source of plausibly exogenous variation in trade cost shocks across a large number of *ex ante* asymmetric regions to test for the role of market size in trade integration.

The paper presents evidence in favor of the home market channel of trade integration. The presented findings provide two novel insights. First, the paper presents evidence in favor of the hypothesis that market integration can reinforce the concentration of economic activity. Second, the paper presents a novel empirical test of the home market effect using a very different setting and source of variation compared to existing literature. These findings are in support of theoretically motivated policy discussions in the trade literature (Fujita *et al.*, 1999; Baldwin *et al.,* 2003), and serve to emphasize the importance of potentially unintended general equilibrium consequences when evaluating and planning large scale transport infrastructure policies.

The counterfactual estimations of the calibrated structural model also allow to shed light on a set of economic consequences that are outside the scope of reduced form estimation. In particular, the model suggests that the network led to significant aggregate welfare gains from trade, and illustrates an important distinction between the nominal and real distributional consequences across regions. Because connected peripheral regions gain access to the mass of industrial varieties produced in the metropolitan regions, the estimated consumer price index reductions suggest that the increased concentration in nominal production does not foretell parallel welfare distributional effects. The latter insight serves to emphasize caution when drawing welfare conclusions from observed changes to the regional distribution of nominal production.

References

ADB (2007). A retrospective analysis of the road sector 1997-2005. *Asian Development Bank Operations Evaluation Department Report.*

Angrist, J. & Pischke, J. (2008). *Mostly harmless econometrics: An empiricist's companion.* Princeton Univ Pr.

Au, C. & Henderson, J. (2006). How migration restrictions limit agglomeration and productivity in china. *Journal of Development Economics*, 80(2), 350–388.

Baldwin, R., Forslid, R., Martin, P., Ottaviano, G., & Robert-Nicoud, F. (2003). *Economic geography and public policy.* Princeton Univ Pr.

Banerjee, A., Duflo, E., & Qian, N. (2009). On the road: Access to transportation infrastructure and economic growth in china. *Unpublished manuscript, Massachusetts Institute of Technology, Department of Economics.*

Baum-Snow, N. (2007). Did highways cause suburbanization? *Quaterly Journal of Economics*, 122(2), 775.

Baum-Snow, N., Brandt, L., Henderson, J., Turner, M., & Zhang, Q. (2012). Roads, railroads and decentralization of chinese cities. *mimeo, Brown University.*

Behrens, K., Lamorgese, A., Ottaviano, G., & Tabuchi, T. (2009). Beyond the home market effect: Market size and specialization in a multi-country world. *Journal of International Economics*, 79(2), 259–265.

Brülhart, M. & Trionfetti, F. (2009). A test of trade theories when expenditure is home biased. *European Economic Review*, 53(7), 830–845.

Combes, P. & Lafourcade, M. (2005). Transport costs: measures, determinants, and regional policy implications for france. *Journal of Economic Geography*, 5(3), 319–349.

Davis, D. (1998). The home market, trade, and industrial structure. *American Economic Review*, 88(5), 1264–76.

Davis, D. & Weinstein, D. (1996). Does economic geography matter for international specialization? *National Bureau of Economic Research Working Paper*.

Davis, D. & Weinstein, D. (1999). Economic geography and regional production structure: an empirical investigation. *European economic review*, 43(2), 379–407.

Davis, D. & Weinstein, D. (2003). Market access, economic geography and comparative advantage: an empirical test. *Journal of International Economics*, 59(1), 1–23.

Dijkstra, E. (1959). A note on two problems in connexion with graphs. *Numerische mathematik*, 1(1), 269–271.

Donaldson, D. (2010). Railroads of the raj: Estimating the impact of transportation infrastructure. *mimeo, London School of Economics*.

Douglas, D. (1994). Least-cost path in gis using an accumulated cost surface and slopelines. *Cartographica The International Journal for Geographic Information and Geovisualization*, 31(3), 37–51.

Duranton, G. & Turner, M. (2008). Urban growth and transport. *University of Toronto Working Paper*.

Fujita, M., Krugman, P., & Venables (1999). *The spatial economy: cities, regions and international trade*. Wiley.

Garske, T., Yu, H., Peng, Z., Ye, M., Zhou, H., Cheng, X., Wu, J., & Ferguson, N. (2011). Travel patterns in china. *PloS one*, 6(2).

Goldberg, P. & Pavcnik, N. (2007). Distributional effects of globalization in developing countries. *Journal of Economic Literature*, 45, 39–82.

Hanson, G. & Xiang, C. (2004). The home-market effect and bilateral trade patterns. *American Economic Review*, (pp. 1108–1129).

Head, K. & Ries, J. (2001). Increasing returns versus national product differentiation as an explanation for the pattern of us-canada trade. *American Economic Review*, (pp. 858–876).

Helpman, E. & Krugman, P. (1985). Market structure and international trade. *Cambridge (Mass.)*.

Hering, L. & Poncet, S. (2009). The impact of economic geography on wages: Disentangling the channels of influence. *China Economic Review*, 20(1), 1–14.

Hillberry, R. & Hummels, D. (2008). Trade responses to geographic frictions: A decomposition using micro-data. *European Economic Review*, 52(3), 527–550.

Jha, M., McCall, C., & Schonfeld, P. (2001). Using gis, genetic algorithms, and visualization in highway development. *Computer-Aided Civil and Infrastructure Engineering*, 16(6), 399–414.

Jong, J. & Schonfeld, P. (2003). An evolutionary model for simultaneously optimizing three-dimensional highway alignments. *Transportation Research Part B: Methodological*, 37(2), 107–128.

Krugman, P. (1980). Scale economies, product differentiation, and the pattern of trade. *The American Economic Review*, 70(5), 950–959.

Krugman, P. & Venables, A. (1999). How robust is the home market effect? *mimeo, London School of Economics*.

Kruskal, J. (1956). On the shortest spanning subtree of a graph and the traveling salesman problem. *Proceedings of the American Mathematical society*, 7(1), 48–50.

Li, S. & Shum, Y. (2001). Impacts of the national trunk highway system on accessibility in china. *Journal of Transport Geography*, 9(1), 39–48.

Michaels, G. (2008). The effect of trade on the demand for skill: evidence from the interstate highway system. *The Review of Economics and Statistics*, 90(4), 683–701.

Puga, D. (2002). European regional policies in light of recent location theories. *Journal of Economic Geography*, 2(4), 373.

Puga, D. & Venables, A. (1997). Preferential trading arrangements and industrial location. *Journal of International Economics*, 43(3-4), 347–368.

Redding, S. J. & Sturm, D. M. (2008). The costs of remoteness: Evidence from german division and reunification. *American Economic Review*, 98(5), 1766–97.

Unknown (1995). China paves the way for expressway network. *Asian Infrastructure Monthly*, June.

WB (2007a). China study tour by new delhi transport office team. *World Bank Internal Report*.

WB (2007b). Domestic trade impacts of the expansion of the national expressway network in china. *World Bank EASTR Working Paper No. 14*.

WB (2008). Safe, clean, and affordable: Transport for development. *World Bank Transport Business Strategy 2008-2012*.

Figures (for inclusion in text)

Figure 1: China's National Trunk Highway System

The figure shows Chinese county boundaries in 1999 in combination with the 54 targeted city nodes and the completed expressway routes of the National Trunk Highway System (NTHS) in the year 2007.

Figure 2: Least cost path spanning tree network

The network in red colour depicts the completed NTHS network in 2007. The network in black colour depicts the least cost path spanning tree network. The black routes are the result of a combination of least cost path and minimum spanning tree algorithms. In the first step Dijkstra's (1959) optimal route algorithm is applied to land cover and elevation data in order to construct least costly paths between each bilateral pair of the targeted centers. In the second step, these bilateral cost parameters are fed into Kruskal's (1956) minimum spanning tree algorithm to identify the minimum number of bilateral routes that connect all targeted cities on a single continuous network of the PR China to minimize total route construction costs. Border connections are least costly paths between provincial capitals to the border in border provinces. A detailed description of these computations and additional maps can be found in the Online Appendix.

Figure 3: Euclidean spanning tree network

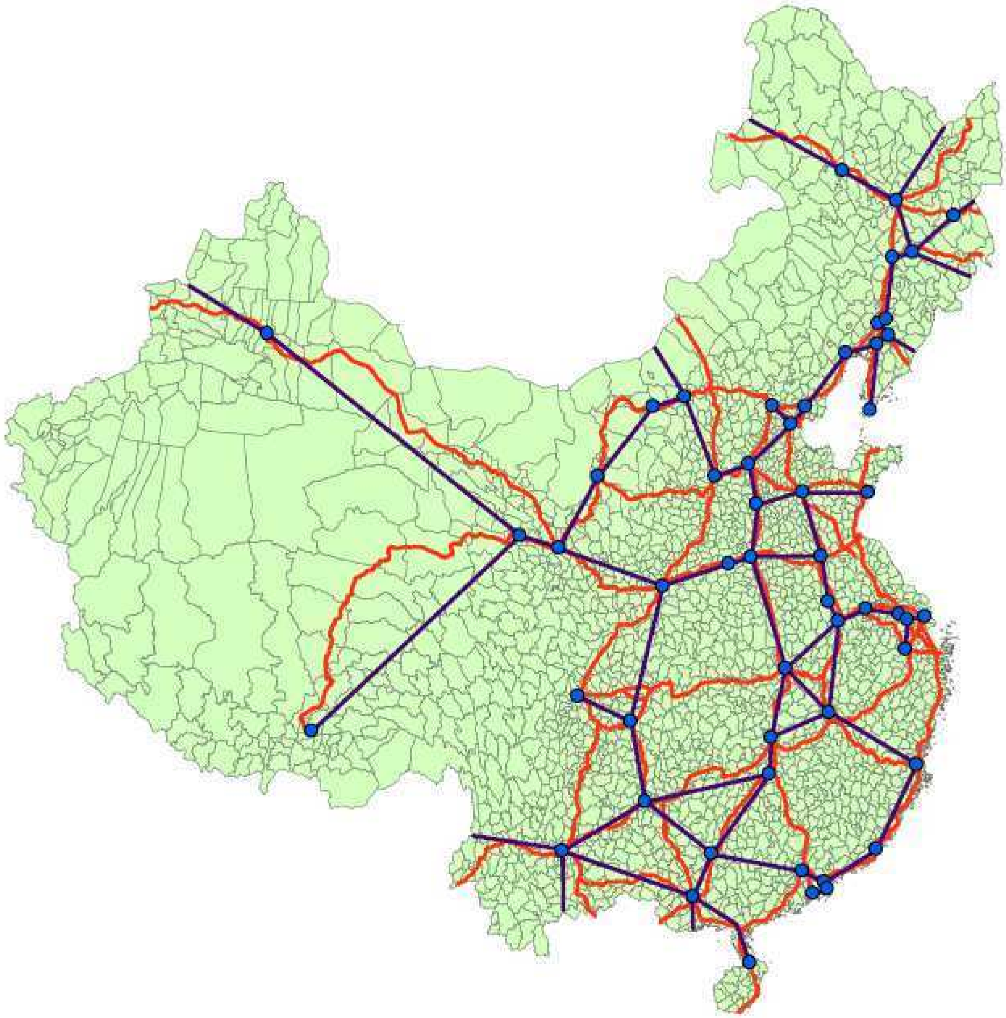

The network in red colour depicts the completed NTHS network in 2007. The network in purple colour depicts the Euclidean spanning tree network. The routes are the result of applying Kruskal's (1956) minimum spanning tree algorithm to bilateral Euclidean distances between targeted nodes. This algorithm is first run for the all-China network, and then repeated within North-Center-South and East-Center-West divisions of China. These repetitions add 9 routes to the original 53 bilateral connections. Connections between provincial capitals of border provinces and the border are minimum Euclidean distance paths. A detailed description of these computations and additional maps can be found in the Online Appendix.

Tables (for inclusion in text)

Table 1: County level descriptive statistics for 1997

	Targeted City Centers	Connected Periphery	Non-connected Periphery	National Share of Targeted City Centers
Population (10,000)	233.24	56.96	38.48	0.14
Urban population (10,000)	179.69	10.77	5.83	0.44
GDP (100 Million Yuan)	517.86	32.58	15.09	0.42
GDP per capita (Yuan)	21435.06	5142.16	3637.09	-
Local government revenue (100 Million Yuan)	38.23	1.23	0.57	0.60
Industrial gross value added (100 Million Yuan)	194.61	14.93	5.58	0.39
Non-Agricultural gross value added (100 Million Yuan)	505.75	24.42	9.74	0.50
Agricultural output share	0.04	0.34	0.42	-
Land area (km^2)	1543.09	3057.47	4513.40	0.013
Number of counties	54	424	943	54

The first three columns present mean 1997 levels, and the fourth column presents national shares by county groups. Targeted city centers refer to the central city county units (shixiaqu) of targeted metropolitan regions. Peripheral counties are counties outside a 50 km commuting buffer around the targeted city centers.

Table 2: First stage regressions

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent Variable:	Connect	Connect	Connect	lnDistHwy	lnDistHwy	lnDistHwy
Least Cost Path IV	0.323***		0.254***	0.317***		0.245***
	(0.0574)		(0.0635)	(0.0645)		(0.0635)
Euclidean IV		0.243***	0.144**		0.280***	0.193***
		(0.0529)	(0.0560)		(0.0599)	(0.0657)
lnDistNode	-0.130***	-0.127***	-0.104***	0.588***	0.635***	0.426***
	(0.0376)	(0.0295)	(0.0323)	(0.130)	(0.112)	(0.122)
Prefect Capital	-0.124*	-0.129*	-0.120*	0.437**	0.429*	0.413*
	(0.0648)	(0.0736)	(0.0658)	(0.209)	(0.229)	(0.215)
City Status	0.0891**	0.0929**	0.0847**	-0.297***	-0.296***	-0.270***
	(0.0403)	(0.0437)	(0.0399)	(0.0946)	(0.103)	(0.0951)
lnUrbPop90	0.106***	0.115***	0.107***	-0.228***	-0.244***	-0.227***
	(0.0225)	(0.0217)	(0.0209)	(0.0691)	(0.0640)	(0.0636)
Educ90	-0.273	-0.303	-0.302	-1.671	-1.747	-1.626
	(0.598)	(0.656)	(0.601)	(1.697)	(1.804)	(1.666)
AgShare90	-0.170	-0.216	-0.167	0.0238	-0.00173	-0.0160
	(0.182)	(0.189)	(0.179)	(0.537)	(0.555)	(0.533)
Constant	-0.212	-0.314	-0.388	2.321**	2.627**	2.695**
	(0.335)	(0.299)	(0.293)	(1.103)	(1.049)	(0.992)
Obs	1342	1342	1342	1342	1342	1342
R^2	0.222	0.204	0.233	0.401	0.394	0.414
First stage F-Stat	31.61	21.07	20.31	24.09	21.82	15.00

All regressions include province fixed effects. Columns 1-3 report results for binary NTHS connection indicators among peripheral counties. Columns 4-6 report results for the log distance to the nearest NTHS segment among peripheral counties. lnDistNode is log county distance to the nearest targeted city node. Prefect Capital and City Status are binary indicators for respective county status in 1990. lnUrbPop90 is log 1990 county urban population. Educ90 is the 1990 county share of above compulsory schooling in 20+ population. AgShare90 is the 1990 county share of agricultural employment. Standard errors are clustered at the province level and stated in parentheses below point estimates. ***1%, **5%, and *10% significance levels.

Table 3: Testing Propositions 1 and 2 on average effects of peripheral network connections

Dependent Variables		(1) OLS No Controls	(2) OLS With Controls	(3) LCP IV No Controls	(4) LCP IV With Controls	(5) Euclid IV No Controls	(6) Euclid IV With Controls	(7) Both IVs No Controls	(8) Both IVs With Controls
	Connect	-0.0529	-0.0356	-0.284**	-0.304**	-0.246*	-0.287*	-0.272***	-0.297***
Change ln(IndGVA)		(0.0418)	(0.0499)	(0.118)	(0.145)	(0.148)	(0.154)	(0.0965)	(0.108)
1997-2006	Obs	1302	1280	1302	1280	1302	1280	1302	1280
	R^2	0.242	0.255						
	Connect	-0.0411	-0.0266	-0.243**	-0.252**	-0.270**	-0.296**	-0.251***	-0.268***
Change ln(NonAgGVA)		(0.0335)	(0.0375)	(0.0983)	(0.117)	(0.122)	(0.131)	(0.0877)	(0.0969)
1997-2006	Obs	1285	1262	1285	1262	1285	1262	1285	1262
	R^2	0.270	0.284						
	Connect	-0.0497*	-0.0914***	-0.0542	-0.223*	-0.175	-0.315**	-0.0926	-0.257***
Change ln(GovRevenue)		(0.0285)	(0.0295)	(0.109)	(0.120)	(0.117)	(0.132)	(0.0893)	(0.0996)
1997-2006	Obs	1290	1285	1290	1285	1290	1285	1290	1285
	R^2	0.275	0.334						
	Connect	-0.00204	-0.0144	-0.106	-0.177*	-0.178	-0.254**	-0.127	-0.203**
Change ln(GDP)		(0.0245)	(0.0276)	(0.0830)	(0.0942)	(0.112)	(0.116)	(0.0824)	(0.0886)
1997-2006	Obs	1297	1272	1297	1272	1297	1272	1297	1272
	R^2	0.228	0.264						
	Connect	-0.00344	-0.00790	0.000194	-0.0252	-0.0305	-0.0597	-0.00865	-0.0371
Change ln(AgGVA)		(0.0210)	(0.0220)	(0.0631)	(0.0789)	(0.0672)	(0.0728)	(0.0545)	(0.0630)
1997-2006	Obs	1335	1313	1335	1313	1335	1313	1335	1313
	R^2	0.202	0.208						
	Connect	0.00488	-0.00217	0.0395**	0.0264	0.0183	0.0104	0.0333*	0.0207
Change ln(Population)		(0.00456)	(0.00568)	(0.0188)	(0.0234)	(0.0242)	(0.0262)	(0.0183)	(0.0215)
1997-2006	Obs	1337	1314	1337	1314	1337	1314	1337	1314
	R^2	0.234	0.271						

All regressions include province fixed effects. LCP IV stands for the least cost path spanning tree instrument. Euclid IV stands for the straight line spanning tree instrument. "No controls" columns refer to regressions on NTHS treatment and log county distance to the nearest targeted city node. "With controls" indicates a full set of 1990 county controls (city status and prefecture capital dummies, log urban population, share of agricultural employment, and share of above compulsory school attainment in 20+ population). The dependent variables in order as listed are county level industry gross value added, manufacturing plus services gross value added, local government revenue, total GDP, agricultural gross value added, and population. Standard errors are clustered at the province level and stated in parentheses below point estimates. ***1%, **5%, and *10% significance levels.

Table 4: Placebo falsification test before and after the network was built

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dependent Variable:	OLS	OLS	LCP IV	LCP IV	Euclid IV	Euclid IV	Both IVs	Both IVs
Change ln(LocGovRev)	1990-97	1997-06	1990-97	1997-06	1990-97	1997-06	1990-97	1997-06
Panel A: Binary								
Connect	0.0154	-0.0848**	0.0143	-0.151	0.117	-0.282**	0.0563	-0.204***
	(0.0410)	(0.0360)	(0.0853)	(0.0974)	(0.107)	(0.129)	(0.0647)	(0.0467)
Obs	894	894	894	894	894	894	894	894
R^2	0.274	0.339						
First stage F-Stat			19.635	19.635	19.091	19.091	14.930	14.930
Panel B: log Distance								
ln(DistHwy)	-0.0114	0.0160	-0.0409	0.0854*	-0.00442	0.185**	-0.0274	0.122***
	(0.0142)	(0.0190)	(0.0350)	(0.0470)	(0.0573)	(0.0783)	(0.0329)	(0.0430)
Obs	894	894	894	894	894	894	894	894
R^2	0.275	0.336						
First stage F-Stat			18.696	18.696	17.306	17.306	11.259	11.259

All regressions include province fixed effects. LCP IV stands for the least cost path spanning tree instrument. Euclid IV stands for the straight line spanning tree instrument. Regressions include the full set of controls. Panel A presents results for binary NTHS connection indicators, and Panel B presents results for log distance to the nearest NTHS segment. Standard errors are clustered at the province level and stated in parentheses below point estimates. ***1%, **5%, and *10% significance levels.

Table 5: Testing Propositions 4 and 5 on interaction effects of peripheral network connections

Dependent Variable:	Change ln(IndGVA) 1997-2006					Change ln(GDP) 1997-2006				
	(1)	(2)	(3)	(4)	(5)	(1)	(2)	(3)	(4)	(5)
Panel A: Binary										
Connect	-0.304**	-4.281***	-4.236***	-3.147	-6.001**	-0.177*	-3.571***	-3.483***	-3.218**	-4.597**
	(0.145)	(1.569)	(1.620)	(2.146)	(2.575)	(0.0942)	(1.011)	(1.023)	(1.533)	(1.848)
Connect*ln(DistNode)		0.748***	0.740***	0.784***	0.891*		0.636***	0.626***	0.649***	0.759**
		(0.270)	(0.277)	(0.267)	(0.472)		(0.172)	(0.172)	(0.170)	(0.312)
Connect*Emp90Dum		0.450*	0.473*	0.468*	0.689		0.404**	0.485***	0.412**	0.642**
		(0.255)	(0.253)	(0.262)	(0.444)		(0.196)	(0.184)	(0.193)	(0.315)
Obs	1280	1280	1280	1280	1020	1272	1272	1272	1272	1024
First stage F-Stat	29.966	3.462	1.765	2.601	1.317	27.972	4.724	1.717	2.517	1.553
Panel B: log Distance										
lnDistHwy	0.0954	1.465***	1.495***	1.067	1.689***	0.0639	1.105***	1.109***	0.861*	1.238***
	(0.0674)	(0.455)	(0.463)	(0.669)	(0.562)	(0.0434)	(0.318)	(0.324)	(0.507)	(0.437)
lnDistHwy *ln(DistNode)		-0.236***	-0.239***	-0.245***	-0.239***		-0.181***	-0.181***	-0.180***	-0.184***
		(0.0748)	(0.0762)	(0.0737)	(0.0821)		(0.0494)	(0.0501)	(0.0502)	(0.0619)
lnDistHwy*Emp90Dum		-0.266***	-0.250***	-0.250***	-0.258**		-0.192***	-0.188***	-0.184***	-0.192**
		(0.0823)	(0.0823)	(0.0821)	(0.111)		(0.0693)	(0.0673)	(0.0675)	(0.0833)
Obs	1280	1280	1280	1280	1020	1272	1272	1272	1272	1024
First stage F-Stat	22.367	4.649	2.720	2.286	2.004	21.698	4.842	2.876	2.549	2.355

All regressions include province fixed effects. Reported results are two-stage least squares estimates using the least cost path spanning tree to instrument for NTHS treatments as well as their reported interaction terms. lnDistNode is log county distance to the nearest targeted city node. Emp90Dum is a dummy for counties with above mean levels of county employment in 1990. All regressions include a full set of county controls. Columns 1 and 2 do not include controls for additional interaction terms. Column 3 reports results after including additional interaction terms between NTHS treatments and a dummy indicator for city status in 1990, as well as a dummy indicator for prefecture level capital status in 1990. Column 4 reports results after including NTHS interactions with 1990 county shares of above compulsory schooling in 20+ population, as well as 1990 county shares of agricultural employment. Column 5 reports results after including an interaction term with log 1990 county government revenue per capita. Standard errors are clustered at the province level and stated in parentheses below point estimates. ***1%, **5%, and *10% significance levels.

Table 6: Best fitting parameter combinations

Rank	Distance Parameter (δ)	Elasticity of Substitution (σ)	Trade Cost Elasticity (Φ)	Simulated Average Treatment Effect	Sum of Squared Deviations from Two Estimated Interaction Terms
1	1.14	2.4	-1.60	-0.293***	0.156
2	1.04	2.6	-1.66	-0.301***	0.197
3	0.96	2.8	-1.73	-0.298***	0.223

The table presents three parameter combinations that yield statistically significant simulated average connection effects within +/-0.5 percentage points of the estimated instrumental variable coefficient in the final column of Table 3. In addition, the presented configurations yield statistically significant (10% level) cross-derivate effects with respect to county size and distance to the nearest targeted center in line with the effects estimated by two stage least squares in Table 5. The ranking refers to the sum of squared deviations over the two estimated interaction coefficients presented in Table 5. ***1%, **5%, and *10% significance levels.

Table 7: Simulation results for the all-China county sample

Dependent Variable: Simulated Change ln(IndOutput)	(1)	(2)
Peripheral Counties (Constant)	-0.313***	
	(0.0761)	
Targeted Metropolitan City Regions	0.319***	0.312***
	(0.0712)	(0.0873)
Province Fixed Effects	No	Yes
Obs	1679	1679
R^2	0.01	0.18

The table presents results from an OLS regression of simulated changes in log county industrial output on an identifier that takes the value of one for targeted metropolitan city regions, and zero otherwise. Simulated outcomes are based on the best fitting parameter configuration presented in Table 6. Column 2 includes province fixed effects. Standard errors are clustered at the province level and stated in parentheses below point estimates. ***1%, **5%, and *10% significance levels.

Table 8: Simulated changes to china's concentration of industrial activity

	All-China Counties	Within Peripheral County Group	Within Targeted City Region Group
Observed Herfindahl 1997	0.0129	0.0036	0.0381
Simulated Herfindahl After NTHS	0.0134	0.0039	0.0383
Percentage Change	3.9	8.3	0.5
Observations	1679	1625	54

Table 9: Simulated effects on local price indices

Dependent Variable: Simulated Change ln(Price Index)	(1)	(2)	(3)	(4)
All-China Average	-0.0106***	-0.0107***		
	(0.000918)	(0.00103)		
Non-Connected Periphery (Constant)			-0.00242***	
			(0.000295)	
Connected Periphery			-0.0254***	-0.0258***
			(0.00260)	(0.00273)
Targeted Metropolitan City Regions			-0.00243***	-0.00410***
			(0.000528)	(0.000874)
Province fixed effects	No	No	No	Yes
1997 Population Weighted	No	Yes	No	No
Obs	1679	1671	1679	1679
R^2	0.000	0.000	0.208	0.234

The table presents results from OLS regressions of simulated changes in the log county ideal price index on a constant (in Columns1 and 2), and on two group identifiers as indicated in Columns 2 and 3 where non-connected peripheral counties constitute the third group in the reference category. Simulated outcomes are based on the best fitting parameter configuration presented in Table 6. Standard errors are clustered at the province level and stated in parentheses below point estimates. ***1%, **5%, and *10% significance levels.

Appendix

Figure A.1: Simulated average treatment effects across the parameter space

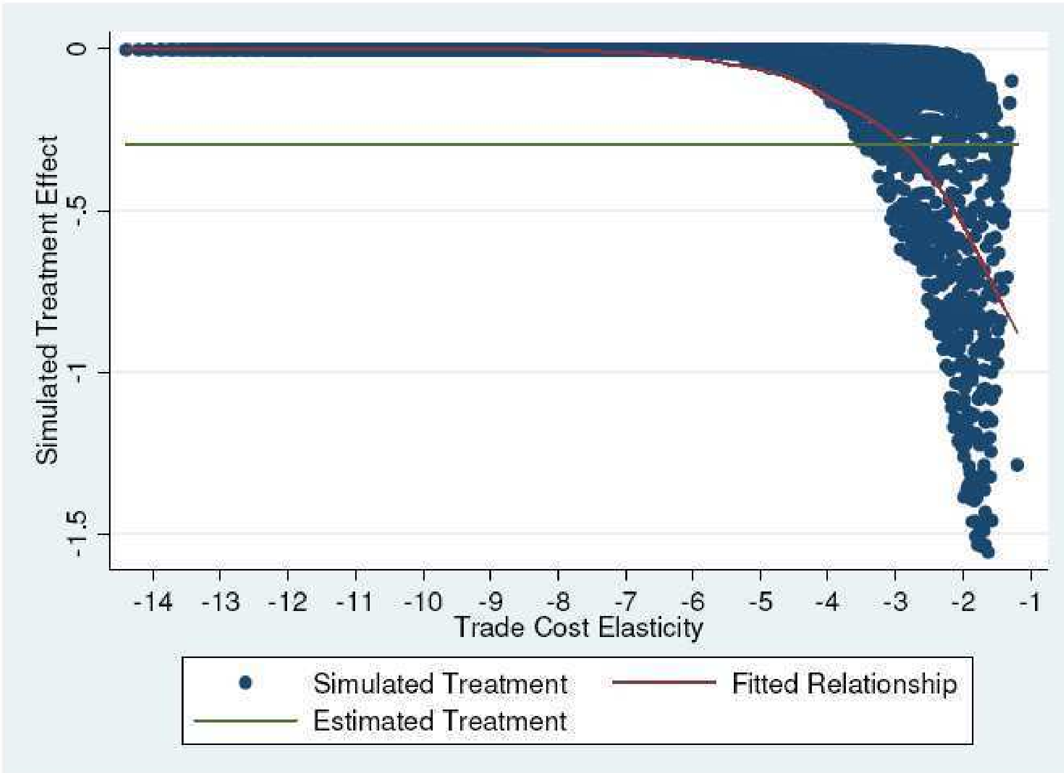

The y-axis displays regression coefficients of county level simulated industrial growth on binary NTHS connection treatments, province fixed effects, and a full set of 1990 county controls for the peripheral county sample. Each dot represents a separate vector of simulated county growth rates. The horizontal line indicates the estimated instrumental variable NTHS effect for the identical peripheral county sample and specification as presented in the final column of Table 3. Trade cost elasticity is the elasticity of trade flows to distance ($\delta*(1-\sigma)$). Depicted results are based on the baseline estimate of NTHS transportation costs relative to existing roads as described in the text. The graph suppresses extreme simulation results at lower levels of trade costs at which a significant fraction of peripheral counties are at the corner solution with zero remaining industry. The fitted line refers to a locally weighted linear least squares regression (bandwidth 0.8).

Figure A.2: Simulated treatment effects for different county groups

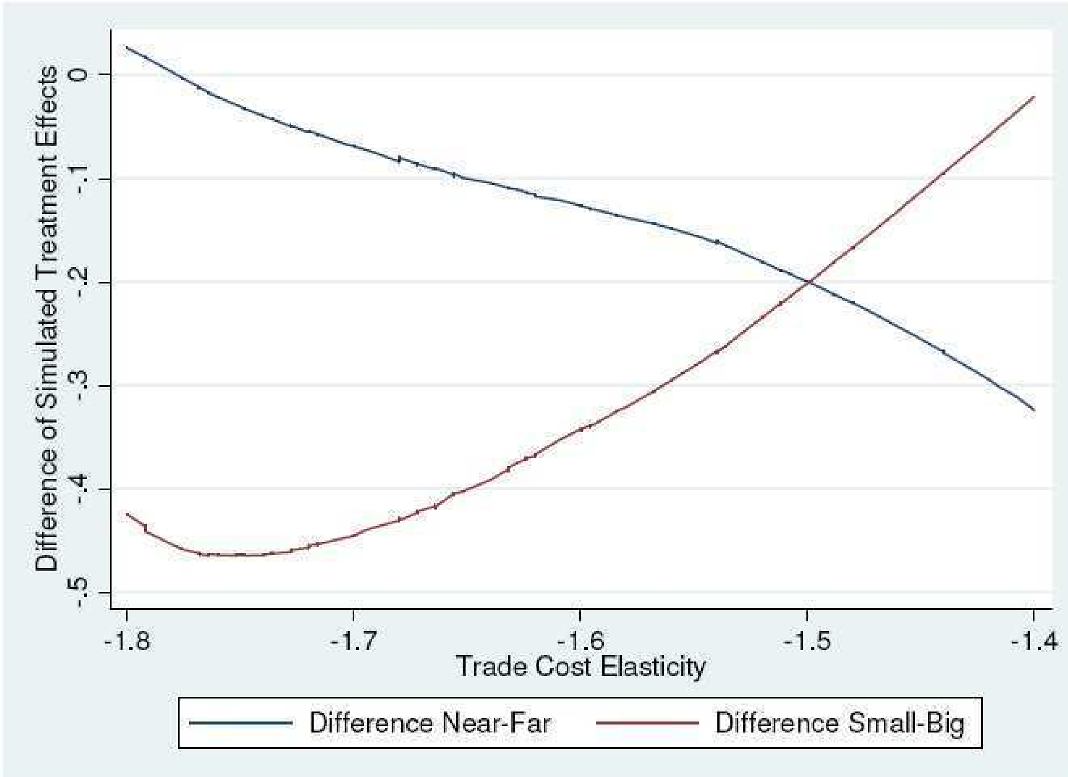

The y-axis displays the difference of regression coefficients of simulated county industrial growth on binary NTHS connection treatments and province fixed effects between different groups of peripheral counties. Regression coefficients are estimated separately for sub-samples of above mean ("Far") and below mean distance ("Near") to the nearest targeted metropolitan region, as well as above mean ("Big") and below mean ("Small") 1990 total employment. The fitted lines refer to locally weighted linear least squares regressions (bandwidth 0.8). Trade cost elasticity refers to the elasticity of trade flows to distance ($\delta^*(1-\sigma)$).

Figure A.3: Parameter configurations within 10% of estimated average treatment effect

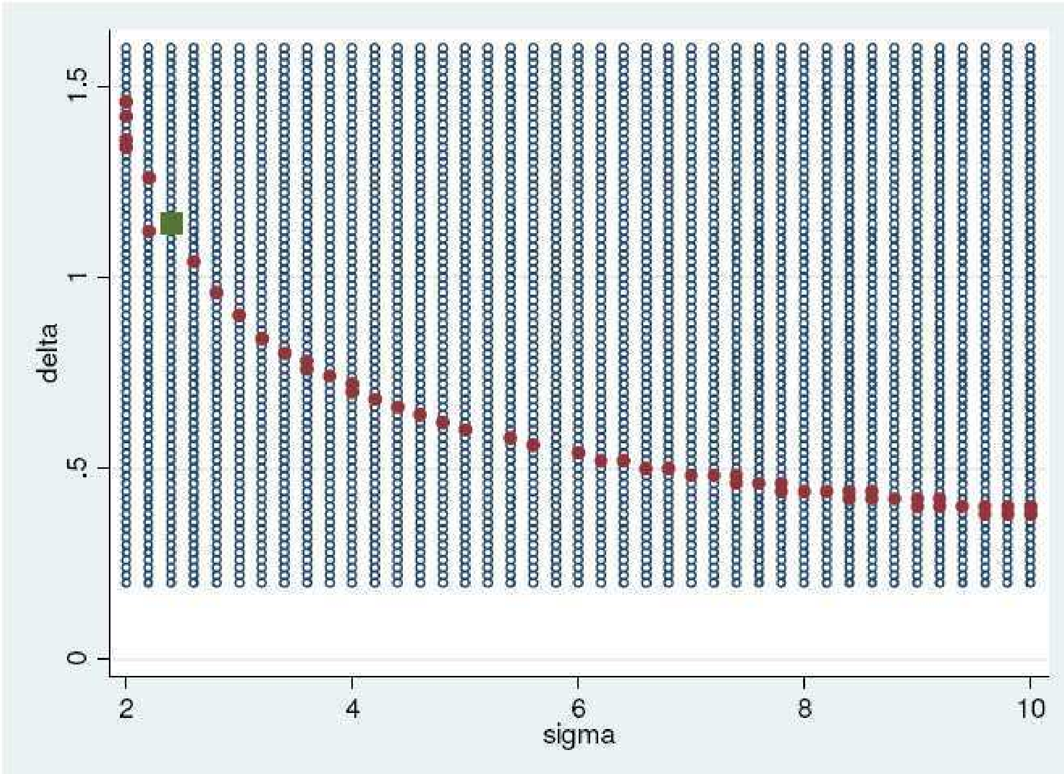

The graph depicts the parameter space where each circle represents one vector of simulated county outcomes subject to the baseline parameterization of NTHS transport costs relative to existing roads described in the text. The solid circles indicate parameter combinations of the distance parameter (δ) and the elasticity of substitution (σ) for which the simulated average treatment effect among peripheral counties is statistically significant at the 10% level and within three percentage points of the estimated 2nd stage IV estimate presented in the final column of Table 3 (within 10% of point estimate). The solid square indicates the best fitting parameter combination identified by the grid search and presented in Table 6.

Online Appendix - Trade Integration, Market Size, and Industrialization

Benjamin Faber*

This appendix proceeds in six sections. Appendix 1 provides proofs. Appendix 2 presents estimation results concerning the proportion and characterization of complier counties that drive the local average treatment effects estimated in the paper. Appendix 3 presents additional estimation and robustness results for both average NTHS connection effects and their interactions. Appendix 4 describes the datasets and construction of variables. Appendix 5 describes the construction of the least cost path and Euclidean spanning tree networks. Appendix 6 describes the transport network routing analysis used to estimate bilateral trade cost changes in the simulations of the structural model.

Appendix 1: Proofs

Appendix 1.1: Non-Specialization

This section solves for the formal condition under which the numeraire sector is produced in all regions. For this to hold, any $J-1$ subset of regions must be unable to satisfy total demand for that sector in equilibrium. We start from the inequality that the total labor endowment in any region j must exceed the regional labor employed by industrial production in the region:

$$S_{Lj}L > S_{Nj}Na_M x \tag{1.1.1}$$

Where the total mass of industrial varieties, N, is equal to the world endowment of capital K. Using the zero profit condition to characterize the equilibrium firm output scale, $x = \pi\frac{\sigma-1}{a_M}$, and substituting into 1.1.1 we get:

$$S_{Lj}L > S_{Nj}K\pi_j(\sigma-1) \tag{1.1.2}$$

Equilibrium returns to capital are given by world expenditure on capital, $\frac{\mu}{\sigma}Y$, divided by the capital stock, so that $\pi = \pi^* = \frac{\mu}{\sigma}\frac{Y}{K}$. Substituting into 1.1.2, expressing world expenditure as $Y = \frac{L}{1-\frac{\mu}{\sigma}}$, and rearranging, we get:

$$S_{Lj} > S_{Nj}\frac{\sigma-1}{\sigma/\mu-1} \tag{1.1.3}$$

The incomplete specialization condition in 1.1.3 becomes weaker as the expenditure share of the numeraire sector $(1-\mu)$ increases.

*London School of Economics and Centre for Economic Performance; Email: b.s.faber@lse.ac.uk

Appendix 1.2: Proposition 1

Proof. This section solves the three region model for the minimum necessary degree of core-periphery size asymmetry (summarized in S_Y^C) under which equation (8) in the paper is negative for all possible combinations of pre-existing trade costs and asymmetric trade integration across equilibria in which no region loses all of its industrial base. We start by solving for:

$$S_N^{P1} - S_N^{P2} = \left(\left(\frac{1}{2} - \frac{3}{2}S_Y^C\right)\frac{(\alpha-1)\phi}{1-\alpha\phi} + 1\right)\frac{1+\alpha\phi-2\phi^2}{(1-\phi)(1+\alpha\phi-2\phi)} - \frac{3\phi}{1+\alpha\phi-2\phi} - 1 < 0 \tag{1.2.1}$$

Which yields:

$$S_Y^C > \left(\left(\left(\frac{3\phi}{1-2\phi+\alpha\phi} + 1\right)\frac{(1-\phi)(1-2\phi+\alpha\phi)}{1+\alpha\phi-2\phi^2} - 1\right)\frac{1-\alpha\phi}{(\alpha-1)\phi} - \frac{1}{2}\right) * \left(-\frac{2}{3}\right) \tag{1.2.2}$$

We now solve for the maximum degree of core-periphery size asymmetry for which all regions retain some industrial activity at the initial equilibrium prior to the trade cost shock as a function of the initial trade freeness. We start with:

$$S_N^C = S_Y^C \frac{1+\phi-2\phi^2}{(1-\phi)^2} - \frac{\phi}{1-\phi} < 1 \tag{1.2.3}$$

Which yields:

$$S_Y^C < \frac{1-\phi}{1+\phi-2\phi^2} \tag{1.2.4}$$

We can now ask what is the minimum necessary core-periphery size gradient at which $S_N^{P1} - S_N^{P2}$ remains negative for all combinations of initial trade costs and asymmetric integration given that in equilibrium all regions retain some industrial activity. The first step is to notice that 1.2.2 is decreasing in the size of the preferential trade shock of the integrating peripheral region (α) in the range $1 < \alpha < \frac{1}{\phi}$. Setting α to an arbitrarily small differential above unity, we can then solve for the minimum necessary size of the metropolitan region by setting equal 1.2.2 and 1.2.4. This yields $S_Y^C > 0.5$. ∎

This result can be graphically illustrated as follows:

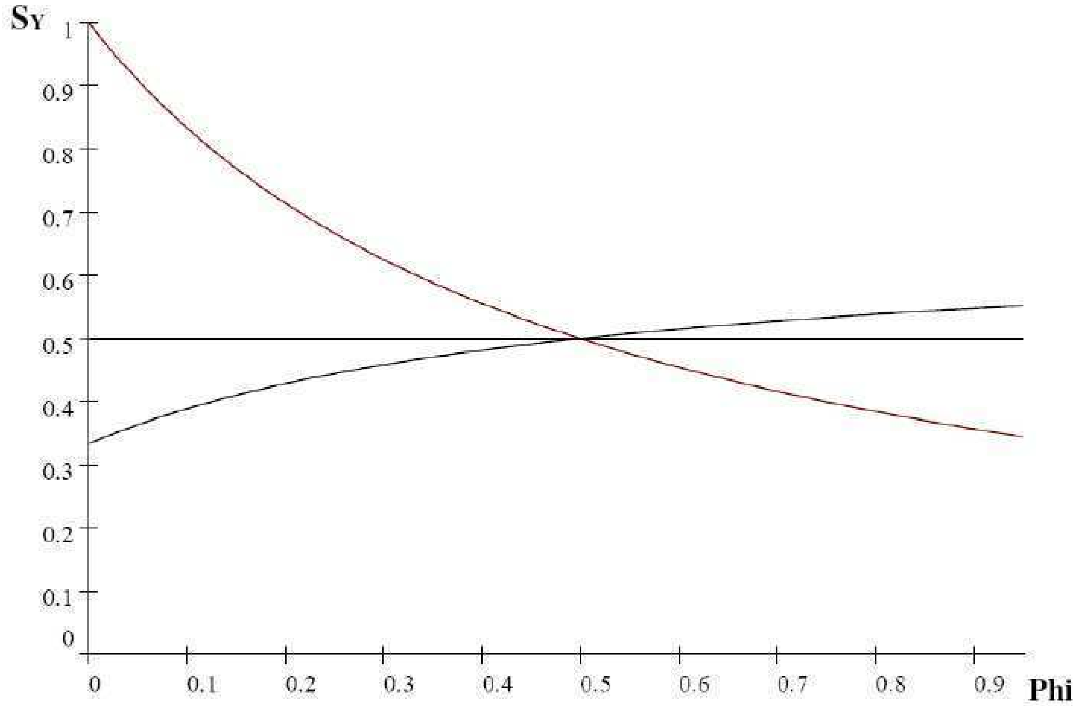

The x-axis displays the degree of initial trade freeness between regions. The y-axis displays the share of total expenditure located in the metropolitan region. The upward sloping function is given in equation 1.2.2 subject to an arbitrarily small preferential trade cost shock of α >1. The downward sloping function is given in equation 1.2.4.

Appendix 1.3: Proposition 2 (GDP)

Proof. The previous section proves that $\frac{\partial (S_N^{P1} - S_N^{P2})}{\partial \alpha} < 0$ holds as long as $S_Y^C > 0.5$. This section proves that changes in regional total output differences between the two peripheral regions move in parallel to changes in the difference of industrial output, but less than proportional. We start by expressing the difference of nominal industrial production between the two peripheral regions as a function of the difference in the mass of manufacturing firms:

$$GDP_{Ind}^{P1} - GDP_{Ind}^{P2} = \left(S_N^{P1} - S_N^{P2}\right) Nxp \tag{1.3.1}$$

Using equilibrium outcomes to firm scale ($x = \pi \frac{\sigma - 1}{a_M}$), factory gate prices of industrial varieties ($p = a_M \frac{\sigma}{1 - \sigma}$), and capital returns ($\pi = \frac{\mu}{\sigma} \frac{Y}{K}$), and substituting $N = K$ we get:

$$GDP_{Ind}^{P1} - GDP_{Ind}^{P2} = \left(S_N^{P1} - S_N^{P2}\right) \mu Y \tag{1.3.2}$$

We now relate this to changes in the difference of total regional production. Total regional production is the sum of factor returns that are employed in the region, so that:

$$GDP^{P1} - GDP^{P2} = \left(S_N^{P1} - S_N^{P2}\right) N\pi + \left(S_L^{P1} - S_L^{P2}\right) L \tag{1.3.3}$$

Using equilibrium outcomes of capital returns, and substituting the mass of industrial varieties for the world capital stock as above, we get:

$$GDP^{P1} - GDP^{P2} = \left(S_N^{P1} - S_N^{P2}\right) \frac{\mu}{\sigma} Y + \left(S_L^{P1} - S_L^{P2}\right) L \tag{1.3.4}$$

Labor immobility implies that the second term is constant over time, so that given $0 < \mu < 1 < \sigma$, and the sufficient core-periphery gradient derived in the previous section ($S_Y^C > 0.5$), it must hold true that:

$$\frac{\partial(GDP_{Ind}^{P1} - GDP_{Ind}^{P2})}{\partial\alpha} = \frac{\partial(S_N^{P1} - S_N^{P2})}{\partial\alpha}\mu Y < \frac{\partial(GDP_{Ind}^{P1} - GDP_{Ind}^{P2})}{\partial\alpha} = \frac{\partial(S_N^{P1} - S_N^{P2})}{\partial\alpha}\frac{\mu}{\sigma}Y < 0 \quad (1.3.5)$$

■

Appendix 1.4: Proposition 2 (Agriculture)

Proof. This section proves that changes in regional agricultural output differences between the two peripheral regions are affected by asymmetric trade integration in the opposite way compared to changes in total output or industrial production differences. We start by expressing the difference of nominal regional agricultural production as a function of industrial activity shares and labor endowments using the assumption of full employment:

$$GDP_{Ag}^{P1} - GDP_{Ag}^{P2} = \left(S_L^{P1} - S_L^{P2}\right)L - \left(S_N^{P1} - S_N^{P2}\right)Nxa_M \quad (1.4.1)$$

Using equilibrium outcomes to firm scale ($x = \pi\frac{\sigma - 1}{a_M}$) and capital returns ($\pi = \frac{\mu}{\sigma}\frac{Y}{K}$), and substituting $N = K$ we get:

$$GDP_{Ag}^{P1} - GDP_{Ag}^{P2} = \left(S_L^{P1} - S_L^{P2}\right)L - \left(S_N^{P1} - S_N^{P2}\right)\mu\frac{\sigma - 1}{\sigma}Y \quad (1.4.2)$$

So that given labor immobility, $0 < \mu < 1 < \sigma$, and the sufficient core-periphery gradient derived in the previous section ($S_Y^C > 0.5$), it must hold true that:

$$\frac{\partial(GDP_{Ag}^{P1} - GDP_{Ag}^{P2})}{\partial\alpha} = (-1)\frac{\partial(S_N^{P1} - S_N^{P2})}{\partial\alpha}\mu\frac{\sigma - 1}{\sigma}Y > 0 \quad (1.4.3)$$

■

Appendix 1.5: Proposition 3

Proof. This section provides the proof for the cross-derivative prediction formalized in Proposition 3.

$$S_N^{P1} - S_N^{P2} = \overbrace{\left(\left(\frac{1}{2} - \frac{3}{2}S_Y^C\right)\frac{(\alpha - 1)\phi}{1 - \alpha\phi} + 1\right)}^{\text{"}A\text{"}} * \overbrace{\frac{1 + \alpha\phi - 2\phi^2}{(1 - \phi)(1 + \alpha\phi - 2\phi)}}^{\text{"}B\text{"}} - \overbrace{\frac{3\phi}{1 + \alpha\phi - 2\phi}}^{\text{"}C\text{"}} - 1 \quad (1.5.1)$$

$$\frac{\partial^2(S_N^{P1} - S_N^{P2})}{\partial\alpha\partial\phi} = \frac{\partial^2(A)}{\partial\alpha\partial\phi} * B + \frac{\partial(B)}{\partial\phi}\frac{\partial(A)}{\partial\alpha} + \frac{\partial^2(B)}{\partial\alpha\partial\phi} * A + \frac{\partial(A)}{\partial\phi}\frac{\partial(B)}{\partial\alpha} - \frac{\partial^2(C)}{\partial\alpha\partial\phi} \quad (1.5.2)$$

Solving for these terms, we find that conditional on $0 < \mu < 1 < \sigma$, as well as the core-periphery gradient conditions identified in Appendix 1.2 ($S_Y^C > 0.5$ and $\phi < 0.5$), the following holds: $\frac{\partial^2(A)}{\partial\alpha\partial\phi} < 0$, $B > 0$, $\frac{\partial(B)}{\partial\phi} > 0$, $\frac{\partial(A)}{\partial\alpha} < 0$, $\frac{\partial^2(B)}{\partial\alpha\partial\phi} < 0$, $A < 0$, $\frac{\partial(A)}{\partial\phi} < 0$, $\frac{\partial(B)}{\partial\alpha} < 0$, and $\frac{\partial^2(C)}{\partial\alpha\partial\phi} < 0$. Finally, applying some more algebra, it holds true that the sum of the absolute values of the first two additive terms in 1.5.2 exceeds the sum of the absolute values of the remaining terms, so that $\frac{\partial^2(S_N^{P1} - S_N^{P2})}{\partial\alpha\partial\phi} < 0$. ■

Appendix 1.6: Proposition 4

Proof. This section provides the proof for the cross-derivative prediction formalized in Proposition 4. Following notation from the previous section, the interaction effect with respect to the core-periphery size gradient is:

$$\frac{\partial^2(S_N^{P1} - S_N^{P2})}{\partial\alpha\partial S_Y^C} = \frac{\partial^2(A)}{\partial\alpha\partial S_Y^C} * B \tag{1.6.1}$$

Solving for the first term, we get:

$$\frac{\partial^2(S_N^{P1} - S_N^{P2})}{\partial\alpha\partial S_Y^C} = -\frac{3}{2}\frac{\phi(1-\phi)}{(\alpha\phi - 1)^2} * B < 0 \tag{1.6.2}$$

■

Appendix 1.7: Graphical illustration of Propositions 1-4

Average and cross-derivative predictions from a three-region scenario

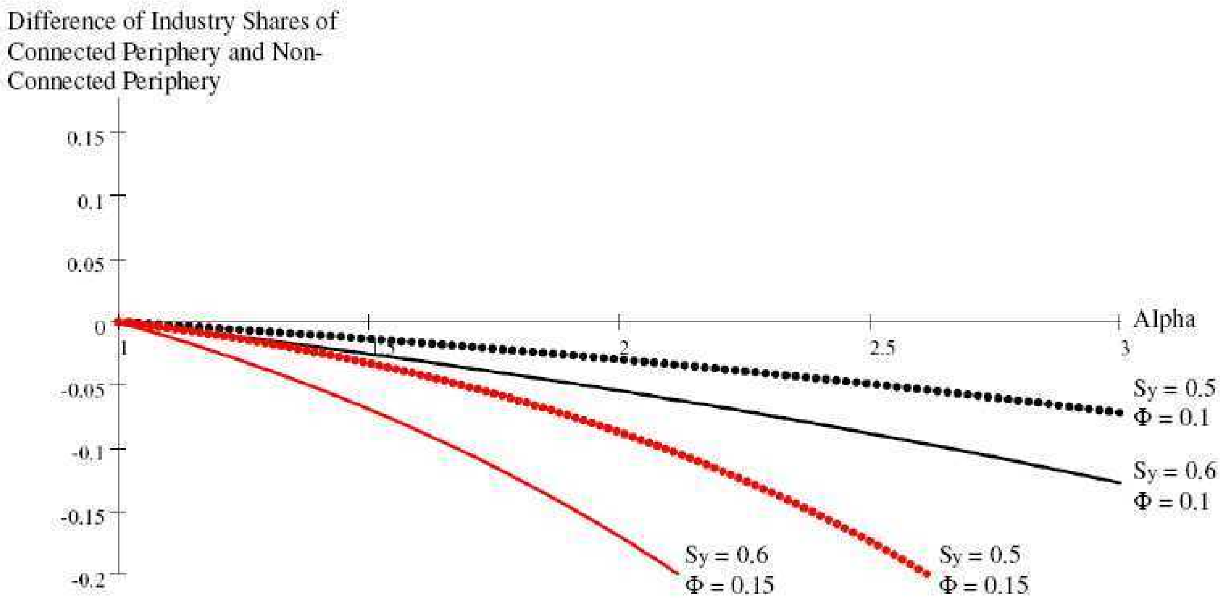

The x-axis displays the degree to which the policy treatment lowers the trade cost of the connected peripheral region to the core region relative to the non-connected peripheral region. The axis starts at the initially identical trade freeness vis-à-vis the metropolitan core (α=1). The y-axis displays the difference of industrial production shares between the connected and the non-connected peripheral regions. Sy is the share of total expenditure located in the metropolitan region, and Φ is the initial trade freeness parameter between all regions.

Appendix 2: Local Average Connection Effects

The instrumental variable estimates presented in the paper represent the local average treatment effect (LATE) of network connections among peripheral counties whose treatment status is affected by location along the all-China least cost spanning tree network. The evaluation literature refers to the latter category as "compliers", as opposed to "always taker" counties that were connected despite their location away from the spanning tree paths.[1]

Descriptive statistics and the pattern of coefficient estimates discussed in the paper suggest that planners targeted economically prosperous counties on the way between targeted city regions. In this empirical context, the concern addressed in this additional set of estimations is that least cost spanning tree location might have affected actual highway placements only for a subset of remote and economically stagnant counties on the way between targeted nodes, so that the estimated local average NTHS connection effects might systematically differ from population average effects.

While it is not possible to identify the complier status of individual counties in the county sample, it is possible to estimate the proportion of compliers among the treated counties as well as their observable characteristics (Abadie, 2003; Angrist and Pischke, 2008). The proportion of compliers among all actually treated NTHS counties is given by:

$$\begin{aligned} P(C_{1i} > C_{0i} | C_i = 1) &= \frac{P(C_i = 1 | C_i > C_{0i}) P(C_{1i} > C_{0i})}{P(C_i = 1)} \\ &= \frac{P(z_i = 1)(E(C_i | z_i = 1) - E(C_i | z_i = 0))}{P(C_i = 1)} \end{aligned} \tag{2.1}$$

where C_i is actual NTHS connection status of county i, C_{1i} and C_{0i} are the connection status in cases where the instrument predicts treatment or not, P and E are probability and expectation operators, and z_i is the treatment value of county i as predicted by the instrument. The second equality makes use of the two facts that the total size of the complier group is given by the Wald first stage, and that by independence $P(z_i = 1 | C_{1i} > C_{0i}) = P(z_i = 1)$. The proportion of compliers among treated counties can then be expressed as the product of the first stage estimate and the proportion of predicted treatments, divided by the proportion of actually treated counties. As presented in the first Column of the Table 2.1, this proportion is estimated to be 22% for both least cost path as well as the Euclidean spanning tree instruments.

The critical question is to what extent these complying counties could be systematically different from the rest of the treated counties. In the case of binary treatments and binary dependent characteristics we know that the relative likelihood of compliers falling into the binary observable category is given by the ratio of the first stage Wald estimated for a particular subgroup over the full sample first stage estimate. To this end, Columns 2-6 of Table 2.1 report first stage point estimates in stated order for counties with above mean 1997 levels of population, urban population, the share of urban population, GDP, and GDP per capita.

If the concern was true that the estimated local average treatment effects are unrepresentative of the population average effects, then we would expect the first stage predictive power of the instruments to differ significantly across observable pre-existing county characteristics. As discussed above, this would constitute evidence of significantly different likelihoods of observables among always takers as opposed to complying counties. In particular, in the present setting one would be concerned that the instrumental variable connection predictors would have a lower estimated effect on actual NTHS route placements among the large, urbanized, and rich county groups represented in Columns 2-6.

The reported results provide evidence against this concern. In particular, the first stage

[1] An implicit assumption is that there are no "defiers" in this context, as location along advantageous construction routes does not cause counties not to be on the network.

point estimates do not significantly differ from the full sample first stage estimate when estimated for different subsamples of counties as indicated across the columns. Figure 2.1 then takes a closer cartographic inspection of actual as opposed to predicted route placements to offer two plausible explanations for the absence of clear observable differences between compliers and always takers.

The figure depicts two snapshots at the county level of the PR China in which counties are color coded according to their nominal levels of GDP in 1997. Both cases compare actual NTHS route placements to predictions from the least cost path spanning tree instrument. Case A illustrates the first point. It is evident that the least cost path algorithm is subject to prediction errors for both bigger urban and smaller rural counties, even in cases where no obvious incentive for deviations from the least cost path is evident. This has to do with the fact that entire bilateral route segments might have been differently picked by the instrument as opposed to NTHS routes, and most importantly, because planners built many more bilateral routes than the minimum number of edges that the spanning tree algorithm picks.

Case B illustrates the second point. When planners do seem to have deviated from the least cost path for an obvious reason (e.g. connection of a prefecture level capital city on the way as indicated in the figure), then this deviation for one "important" county leads to prediction errors for several "unimportant" counties on the way to this target. Both of these features that are evident from the GIS snapshots tend to work against any systematic correlation between the predictive power of the instrument in the first stage and the potential heterogeneity of the highway effect.

Table 2.1: Estimated proportion of compliers and relative likelihoods of observable characteristics

	(1) Full Sample	(2) Pop 97	(3) Urban Pop 97	(4) %Urban Pop 97	(5) GDP 97	(6) GDP Cap 97
Panel A: LCP IV						
Connect 1st Stage	0.418***	0.383***	0.432***	0.494***	0.399***	0.433***
Point Estimate	(0.0601)	(0.0821)	(0.0704)	(0.0599)	(0.0873)	(0.0869)
F-Statistic p-value [Coef=0.418]		0.677	0.839	0.214	0.832	0.864
Obs	1367	650	662	633	673	664
Estimated Proportion of Compliers Among Treated Counties	0.222					
Panel B: Euclid IV						
Connect 1st Stage Point Estimate	0.314***	0.354***	0.375***	0.328***	0.365***	0.337***
	(0.0492)	(0.0690)	(0.0822)	(0.0776)	(0.0784)	(0.0712)
F-Statistic p-value [Coef=0.314]		0.567	0.462	0.860	0.521	0.750
Obs	1367	650	662	633	673	664
Estimated Proportion of Compliers Among Treated Counties	0.221					

Columns present first stage point estimates for regressions of binary NTHS connections on spanning tree connections and province fixed effects across different county samples. All regressions include province fixed effects. LCP IV stands for the least cost path spanning tree instrument. Euclid IV stands for the straight line spanning tree instrument. The first column presents the full sample first stage estimate. The following columns (in stated order) present this estimate for counties with above median 1997 levels of population, urban population, shares of urban population, GDP, and GDP per capita. Standard errors are clustered at the province level and stated in parentheses below point estimates. ***1%, **5%, and *10% significance levels.

Figure 2.1: Cartographic inspection of the instrument

Case A | Case B

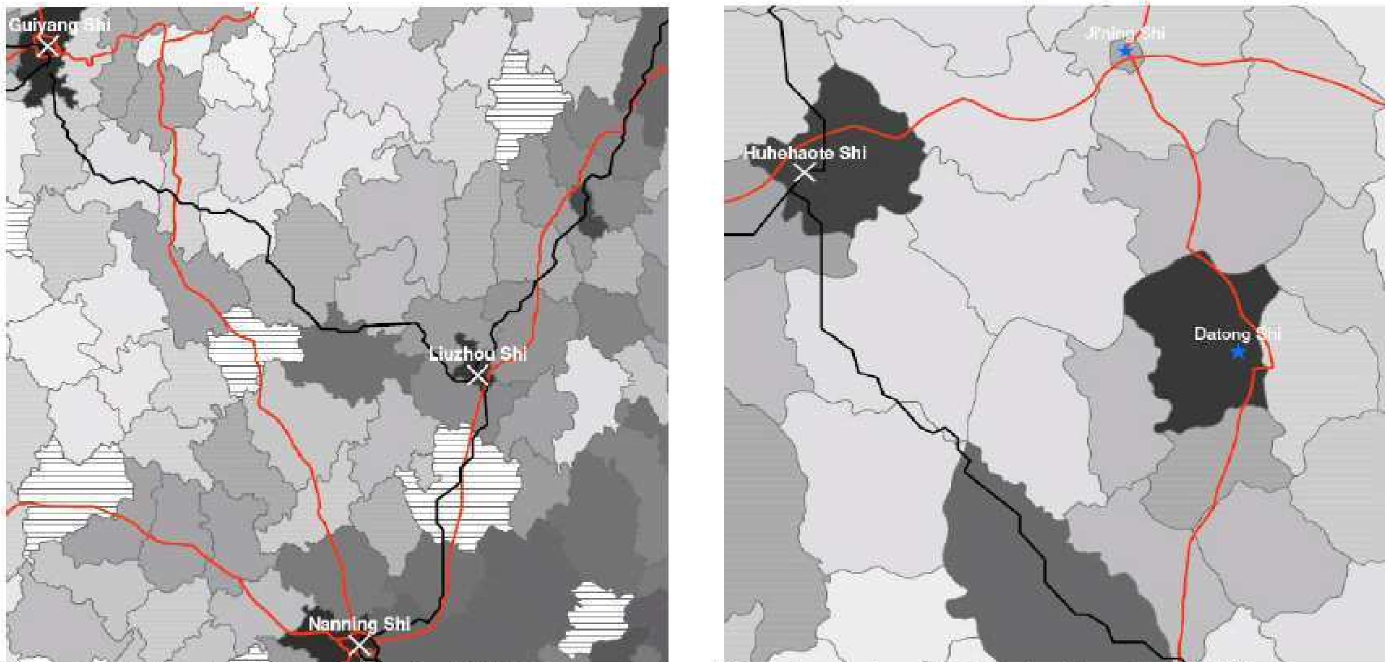

Case A: The network in red colour depicts actual NTHS expressway routes. The network in black colour depicts the least cost path spanning tree network. Crosses indicate targeted metropolitan nodes. Counties are color coded according to their nominal levels of GDP in 1997, where darker colours represent higher values. Striped areas indicate missing 1997 GDP data.

Case B: The network in red colour depicts actual NTHS expressway routes. The network in black colour depicts the least cost path spanning tree network. Crosses indicate targeted metropolitan nodes and stars indicate prefecture capitals. Counties are color coded according to their nominal levels of GDP in 1997, where darker colours represent higher values. Striped areas indicate missing 1997 GDP data.

Appendix 3: Additional Estimation and Robustness Results

This section presents additional estimation and robustness results. Table 3.1 presents OLS and instrumental variable results on average NTHS connection effects after replacing the binary connection identifier by the log distance of a county's center to the nearest expressway segment. The two main patterns discussed in the paper with respect to binary treatment indicators hold in opposite sign for the continuous highway exposure specifications. The first is that OLS coefficients on the effect of log distance to the nearest NTHS route are less positive than IV point estimates. The second is that the IV estimates become more positive after including additional county controls.

While slightly less statistically significant across the full range of specifications, the specifications using log distance to the nearest NTHS segment as continuous exposure to NTHS connections confirm the results discussed in the paper. For the preferred specification with the full set of county controls and using both spanning tree instruments as reported in the final column of the table, distance to the nearest NTHS segment has a positive and statistically significant effect on industrial output growth, non-agricultural output growth, local government revenue growth, as well as aggregate GDP growth.

Table 3.2 reports additional results of the estimations on interaction effects comparing two stage least squares (2SLS) estimates with limited information maximum likelihood (LIML) estimates when using both spanning tree networks to instrument for NTHS connections and its interaction terms with respect to pre-existing county characteristics. The concern addressed in these estimations is that the drop in the first stage F-statistics among specifications with interaction effects could lead to weak instrument bias.

Panel A of the table reports 2SLS results, and Panel B reports LIML results for identical specifications. As was the case for the just identified specifications in the paper, the fact that the point estimates on the interaction terms remain stable across the different specifications across Columns 2-4 with varying levels of first stage F-statistics offers a first reassurance against the concern of weak instrument bias. Secondly, the table indicates that LIML coefficient estimates on the main NTHS effect and its interaction terms are slightly higher across both dependent variables (industrial output growth and GDP growth), as well as across all reported specifications. Given that the LIML estimator has been shown to be less affected by weak instrument bias, this finding provides a second robustness check against this concern.[2]

Finally, Table 3.3 presents estimation results for a series of additional robustness specifications concerning the main average NTHS connection effects on county growth discussed in the paper. The first row of results reproduces the baseline estimates of the NTHS connection effect for industrial output growth, non-agricultural output growth, GDP growth, and local government revenue growth for the preferred specification with both instruments and the full set of pre-existing county controls.

The second row of results presents estimates after adjusting for spatial dependence following Conley (1999). Instead of addressing the concern of spatial autocorrelation of the error term by clustering at the level of 33 provinces, the results are estimated via GMM allowing for spatial dependence as a general decreasing function over bilateral county distance without imposing parametric assumptions on this function. In particular, I estimate each specification allowing for spatial dependence to play a role within a range of bilateral county distances equal to the diameter corresponding to the average province area in China

[2]See for example Angrist and Pischke (2008, Section 4.6) for a discussion 2SLS and LIML estimates in the context of weak instrument concerns.

(600 km). The result is that the standard errors are hardly affected compared to the baseline specification with clustering at the province level, and the change of the standard errors is not systematically higher or lower. Standard errors slightly increase for industrial output growth and local government revenue growth, and slightly decrease for non-agricultural output and aggregate GDP growth.

The third row of results addresses the concern that the geographical characteristics used in the construction of the least cost path instrument could directly affect county growth and thereby lead to a violation of the exclusion restriction. To address this concern, the table reports results after including the average terrain slope gradient, the percentage of water coverage, the percentage of wetlands coverage, and the percentage of developed land coverage as additional county controls. The NTHS connection effect estimates are unaffected by the inclusion of these additional controls, and show a very slight increase.

The fourth row of results addresses the concern that location along least costly paths might be subject to stronger endogeneity concerns in mountainous provinces where valleys provide natural advantages for settlements and economic development. The presented results are estimated after excluding the mountain provinces of Gansu, Qinghai, Sichuan, Tibet, and Xinjiang. The exclusion of these regions also address the concern that due to the mountainous terrain a new long distance railway route was built following closely the route of the NTHS between Golmud and Lasa over the same period. The NTHS connection effects are confirmed in sign and statistical significance for all dependent variables when estimated on the restricted county sample.

The fifth row of results addresses the concern that least costly route locations between the major city regions of China are likely to be correlated with historical trade routes. To this end, I obtained georeferenced routes for the Northern and the Southern routes of the trans-Asian Silk road from the Old World Trade Routes (OWTRAD) Project.[3] The Southern routes of the Silk Road are sometimes referred to as the Tea Horse Road instead. Figure 3.1 provides an illustration of the NTHS network and the Silk Road routes. Reported estimation results include the log distance to the nearest Silk Road segment as an additional county control. The baseline NTHS coefficients are hardly affected by the inclusion of this additional control, indicating that the baseline controls for pre-existing political and economic characteristics have effectively captured county proximity to historical trade routes.

The final row of results addresses the concern that the initial period output levels among NTHS connected peripheral counties might have been inflated by construction activity already underway in 1997. The concern is that the significant negative effects are driven in part by this inflation of the initial levels of economic activity. To address this concern, I include a dummy indicator for road construction underway in 1997 that I collect from the 1998 Atlas source described in the Data Appendix below. The inclusion of this additional county control hardly affects the baseline point estimates of the NTHS connection effects, indicating that road construction activity underway in 1997 did not lead to a spurious negative growth effect among NTHS connected counties.

[3] See www.ciolek.com/owtrad.html.

Table 3.1: Testing average effects using log distance to NTHS instead of binary connections

Dependent Variables		(1) OLS No Controls	(2) OLS With Controls	(3) LCP IV No Controls	(4) LCP IV With Controls	(5) Euclid IV No Controls	(6) Euclid IV With Controls	(7) Both IVs No Controls	(8) Both IVs With Controls
	lnDistHwy	0.000119	-0.00872	0.0987*	0.0954	0.130	0.135	0.112**	0.113*
Change ln(IndGVA)		(0.0157)	(0.0188)	(0.0545)	(0.0674)	(0.0817)	(0.0906)	(0.0533)	(0.0615)
1997-2006	Obs	1302	1280	1302	1280	1302	1280	1302	1280
	R^2	0.241	0.255						
	lnDistHwy	0.000770	-0.00750	0.0929*	0.0939	0.133*	0.145*	0.109**	0.115**
Change ln(NonAgGVA)		(0.0128)	(0.0146)	(0.0487)	(0.0588)	(0.0692)	(0.0781)	(0.0495)	(0.0566)
1997-2006	Obs	1285	1262	1285	1262	1285	1262	1285	1262
	R^2	0.268	0.284						
	lnDistHwy	0.00184	0.0214	0.0670	0.138*	0.149*	0.233***	0.0994*	0.177***
Change ln(GovRevenue)		(0.0156)	(0.0169)	(0.0588)	(0.0724)	(0.0789)	(0.0853)	(0.0559)	(0.0667)
1997-2006	Obs	1290	1285	1290	1285	1290	1285	1290	1285
	R^2	0.274	0.332						
	lnDistHwy	-0.0156	-0.0111	0.0439	0.0639	0.0805	0.113	0.0589	0.0845*
Change ln(GDP)		(0.0103)	(0.0116)	(0.0385)	(0.0434)	(0.0624)	(0.0685)	(0.0433)	(0.0480)
1997-2006	Obs	1297	1272	1297	1272	1297	1272	1297	1272
	R^2	0.230	0.265						
	lnDistHwy	-0.00769	-0.00589	-0.00535	0.00231	0.00556	0.00979	-0.00108	0.00542
Change ln(AgGVA)		(0.00982)	(0.0103)	(0.0322)	(0.0402)	(0.0328)	(0.0381)	(0.0284)	(0.0340)
1997-2006	Obs	1335	1313	1335	1313	1335	1313	1335	1313
	R^2	0.203	0.208						
	lnDistHwy	-0.00320	0.000279	-0.0208***	-0.0172*	-0.0138	-0.0129	-0.0181**	-0.0154
Change ln(Population)		(0.00219)	(0.00245)	(0.00799)	(0.0103)	(0.0112)	(0.0127)	(0.00796)	(0.00941)
1997-2006	Obs	1337	1314	1337	1314	1337	1314	1337	1314
	R^2	0.235	0.271						

All regressions include province fixed effects. LCP IV stands for the least cost path spanning tree instrument. Euclid IV stands for the straight line spanning tree instrument. "No controls" columns refer to regressions on log distance to the nearest NTHS segment and log county distance to the nearest targeted city node. "With controls" indicates a full set of 1990 county controls (city status and prefecture capital dummies, log urban population, share of agricultural employment, and share of above compulsory school attainment in 20+ population). The dependent variables in order as listed are county level industry gross value added, manufacturing plus services gross value added, local government revenue, total GDP, agricultural gross value added, and population. Standard errors are clustered at the province level and stated in parentheses below point estimates. ***1%, **5%, and *10% significance levels.

Table 3.1: Comparing 2SLS and LIML estimations of interaction effects

Dependent Variable:	Change ln(IndGVA) 1997-2006					Change ln(GDP) 1997-2006				
	(1)	(2)	(3)	(4)	(5)	(1)	(2)	(3)	(4)	(5)
Panel A: 2SLS										
Connect	-0.297***	-3.876***	-3.649***	-2.783	-5.856**	-0.203**	-3.496***	-3.389***	-3.004**	-4.729***
	(0.108)	(1.333)	(1.295)	(1.811)	(2.431)	(0.0886)	(0.948)	(0.966)	(1.389)	(1.762)
Connect*ln(DistNode)		0.680***	0.639***	0.744***	0.917**		0.623***	0.605***	0.641***	0.803***
		(0.232)	(0.225)	(0.246)	(0.441)		(0.161)	(0.164)	(0.162)	(0.306)
Connect*Emp90Dum		0.400	0.378	0.467*	0.749*		0.396**	0.389**	0.394**	0.698**
		(0.247)	(0.235)	(0.271)	(0.440)		(0.196)	(0.193)	(0.197)	(0.338)
Obs	1280	1280	1280	1280	1020	1272	1272	1272	1272	1024
First stage F-Stat	18.886	2.047	2.004	1.718	1.204	17.425	2.147	2.149	1.510	1.267
Panel B: LIML										
Connect	-0.297***	-3.914***	-3.706***	-2.872	-6.035**	-0.205**	-3.540***	-3.502***	-3.039**	-4.844***
	(0.108)	(1.353)	(1.327)	(1.878)	(2.556)	(0.0893)	(0.970)	(1.025)	(1.412)	(1.838)
Connect*ln(DistNode)		0.686***	0.649***	0.767***	0.946**		0.631***	0.624***	0.648***	0.822**
		(0.235)	(0.231)	(0.259)	(0.462)		(0.165)	(0.174)	(0.166)	(0.319)
Connect*Emp90Dum		0.406	0.388	0.489*	0.778*		0.404**	0.409**	0.402**	0.718**
		(0.250)	(0.240)	(0.284)	(0.461)		(0.200)	(0.204)	(0.201)	(0.351)
Obs	1280	1280	1280	1280	1020	1272	1272	1272	1272	1024
First stage F-Stat	18.886	2.047	2.004	1.718	1.204	17.425	2.147	2.149	1.510	1.267

All regressions include province fixed effects. Reported results are 2^{nd} stage estimates using the least cost path and the Euclidean spanning tree networks to instrument for NTHS connections as well as their reported interaction terms. lnDistNode is log county distance to the nearest targeted city node. Emp90Dum is a dummy for counties with above mean levels of county employment in 1990. All regressions include a full set of county controls. Columns 1 and 2 do not include controls for additional interaction terms. Column 3 reports results after including additional interaction terms between NTHS treatments and a dummy indicator for city status in 1990, as well as a dummy indicator for prefecture level capital status in 1990. Column 4 reports results after including NTHS interactions with 1990 county shares of above compulsory schooling in 20+ population, as well as 1990 county shares of agricultural employment. Column 5 reports results after including an interaction term with log 1990 county government revenue per capita. Standard errors are clustered at the province level and stated in parentheses below point estimates. ***1%, **5%, and *10% significance levels.

Table 3.3: Additional robustness specifications

Robustness specifications		(1) Change ln(IndGVA) 1997-06	(2) Change ln(NonAgGVA) 1997-06	(3) Change ln(GDP) 1997-06	(4) Change ln(GovRevenue) 1997-06
Baseline estimates	Connect	-0.297***	-0.268***	-0.203**	-0.257***
		(0.108)	(0.0969)	(0.0886)	(0.0996)
	Obs	1280	1262	1272	1285
Conley (1999) standard errors	Connect	-0.297**	-0.268***	-0.203**	-0.257**
		(0.121)	(0.0946)	(0.0800)	(0.100)
	Obs	1280	1262	1272	1285
Control for direct effects of geographical variables used in least cost path construction	Connect	-0.308***	-0.277***	-0.209**	-0.236**
		(0.110)	(0.0988)	(0.0908)	(0.100)
	Obs	1280	1262	1272	1285
Exclude mountain provinces and Golmud-Lasa Railway	Connect	-0.363***	-0.369***	-0.297***	-0.216*
		(0.122)	(0.100)	(0.0940)	(0.123)
	Obs	1043	1032	1040	1039
Control for log distance to historical trade routes	Connect	-0.300***	-0.272***	-0.206**	-0.251**
		(0.106)	(0.0951)	(0.0874)	(0.102)
	Obs	1280	1262	1272	1285
Control for construction underway in 1997	Connect	-0.296***	-0.266***	-0.201**	-0.256**
		(0.111)	(0.0969)	(0.0897)	(0.0996)
	Obs	1280	1262	1272	1285

All regressions include province fixed effects. Reported results are 2nd stage IV estimates using the least cost path and the Euclidean spanning tree networks as instruments for NTHS connections. Regressions include a full set of 1990 county controls. The dependent variables in order of the columns as listed are log changes of county level industrial gross value added, non-agricultural gross value added, total GDP, and local government revenue. Conley (1999) standard errors are estimated via GMM and adjust for spatial dependence without imposing parametric assumptions. I allow spatial dependence to play a role up to a distance range given by the diameter pertaining to the average province area (600 km). Controls for geographical characteristics used in the construction of the least cost path spanning tree instrument are average county slope, and county percentage of wetland water, or developed coverage. Mountainous provinces refer to Gansu, Qinghai, Sichuan, Tibet, and Xinjiang. Historical trade routes are the Northern and Southern routes of the Silk Road (see map presented below). Control for construction in 1997 refers to a dummy variable indicating all counties with reported "under construction" expressway routes in 1997. Except for results using Conley (1999) standard errors, standard errors are clustered at the province level and stated in parentheses below point estimates. ***1%, **5%, and *10% significance levels.

Figure 3.1 The Northern and Southern Routes of the Silk Road

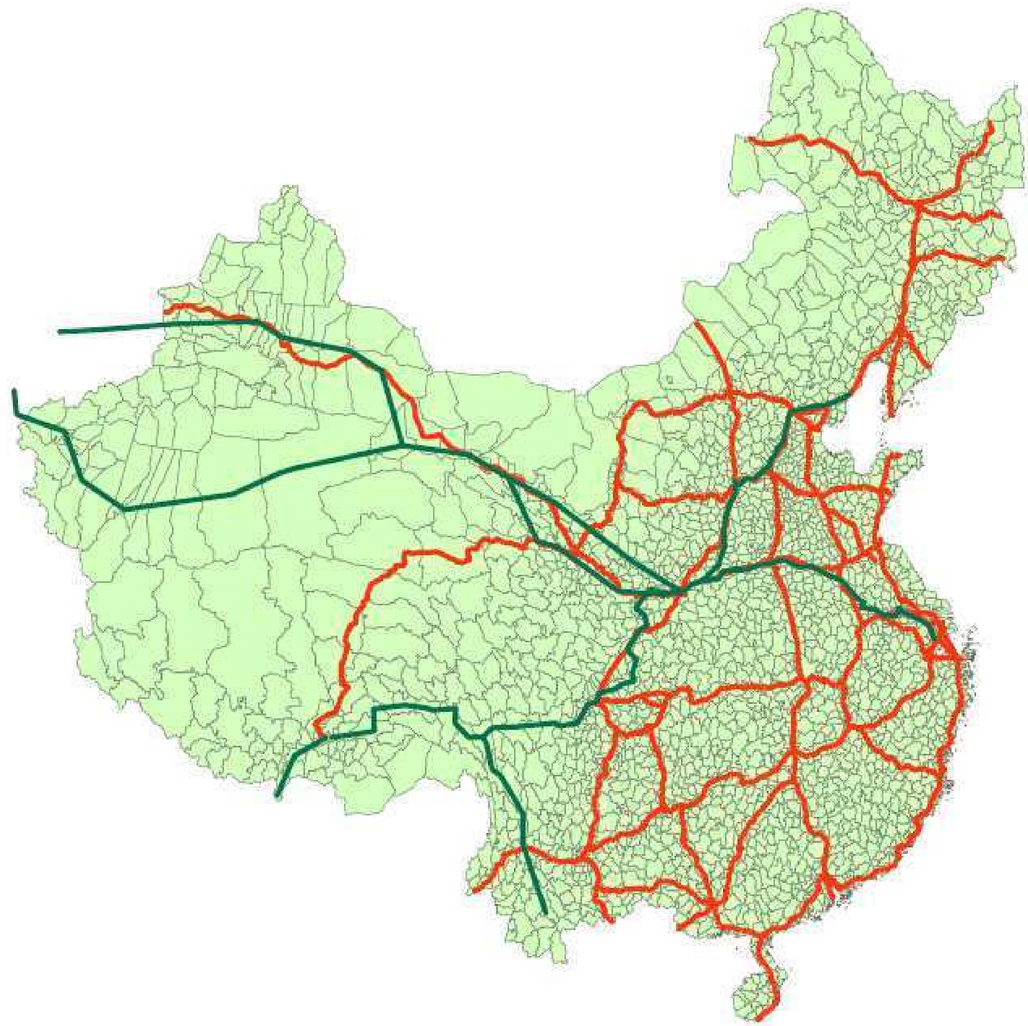

The network in red colour depicts the completed NTHS network in 2007. The green routes represent the Northern and Southern Routes of the Silk Road.

Appendix 4: Data Appendix

GIS Data

Geo-referenced administrative boundary data for the year 1999 was obtained from the ACASIAN Data Center at Griffith University in Brisbane, Australia. These data provide a county-level geographical information system (GIS) dividing the surface of mainland China into 2341 county level administrative units, 349 prefectures, and 33 provinces. Chinese administrative units at the county level are subdivided into county level cities (shi), counties (xian), and urban wards of prefecture level cities (shixiaqu).

Administrative units in China are identified by a system of guo biao codes that allows the matching of records across the GIS and socioeconomic datasets. In addition to guo biao codes, the combination of prefecture and county names were used to double check the consistent matching of administrative units across the datasets. I use reported data on the county area under administration in km^2 from the Provincial Statistical Yearbook series to identify significant boundary changes over time. The historically consistent county sample for estimations on changes 1997-2006 are defined as counties without administrative area changes in excess of 5%. For the placebo falsification test that is estimated on the identical county sample for both the pre- and post-NTHS periods, 1990-1997 and 1997-2007, the same threshold is applied to changes for both periods.

Geo-referenced NTHS highway routes as well as Chinese transport network data were obtained from the ACASIAN Data Center. NTHS highway routes were digitized on the basis of a collection of high resolution road atlas sources published between 1998 and 2007 that is listed below.

(1) *China Newest Public Road Atlas* (1998), Ha Na Bin Map Publishing Company
(2) *China Road Atlas* (2002), Shadong Map Publishing Company
(3) *China Public Road Atlas* (2002), Shandong Map Publishing Company
(4) *China Expressway Atlas* (2003), People's Transport Press
(5) *China Transportation Network Atlas* (2003), Guangdong Map Publishing Company
(6) *China Road Atlas* (2003), Xue Yuan Map Publishing Company
(7) *China Automobile Map* (2003), China World Map Publishing Company
(8) *Chinese People's Road Atlas* (2005), Globe Publishing Company
(9) *China Road Atlas* (2007), Shandong Map Publishing Company

These atlas sources made it possible to classify NTHS segments into three categories: opened to traffic before mid-1997 (10% of NTHS), opened to traffic between mid-1997 and end of 2003 (81% of NTHS), and opened to traffic after the end of 2003 (9% of NTHS). In particular, Source (1) was used to digitize a baseline layer of NTHS routes that were in place by mid year in 1997, and Source (8) was used to digitize a baseline layer of NTHS routes that were in place by the end of 2003. These baseline route maps were then cross-referenced with route information provided in the remaining listed atlas sources. In cases where the remaining atlas sources were at odds with the information of the baseline maps (i.e. routes present in 1997 but not in 2000 or thereafter, or routes present in 2003 but not thereafter), a decision was taken on the basis of the majority of sources (for the 1997 layer), or after tracking down highway openings through press releases on highway opening ceremonies for a small number of cases where Sources (8) and (9) were at odds.

Finally, land cover and elevation data that are used in the construction of least cost path highway routes were obtained from the US Geological Survey Digital Chart of the World (DCW) project, and complemented by higher resolution Chinese hydrology data from the ACASIAN data center. The higher resolution hydrology data from ACASIAN was used to assure that rivers were not interrupted by grid cells coded as mostly covered by land in the lower resolution raster data on land cover obtained from the DCW.

Socio-Economic Data

The Provincial Statistical Yearbook series report production approach county GDP broken up into primary, secondary, and tertiary gross value added. Value added is reported as gross output value less intermediate inputs and value added tax. Traditionally, construction is included together with manufacturing under the secondary industrial sector. The county level data are collected from local establishments under the supervision of the provincial governments, and the Provincial Statistical Yearbooks constitute a separate process of data collection from the national Statistical Yearbook series that is undertaken by the National Bureau of Statistics.

The reported production output data collected by local governments in principle cover the entirety of producing establishments located in the area of the county authority. This is in contrast to central government production statistics that are based on a cut-off of 5 million Yuan annual revenues for so called directly reporting industrial enterprises.[4] The data are collected by teams of local bureaucrats in the form of surveys that are filled out by the establishments located in the jurisdiction of the county.

The Provincial Statistical Yearbooks also provide local government revenues that are reported from the revenue accounts of local authorities. Government revenues mainly consist of industrial and commercial taxes (including value added tax) as well as corporate income taxes (Lin, 2009). The population records contained in the yearbook series refer to locally registered populations under the household registration system.

The CITAS data from the 1990 Population Census provide county level data on population broken up by urban and non-urban, education, and employment shares at the county level. The 1990 Census was the fourth census conducted by the National Bureau of Statistics, and the information therein was recorded on the basis of household questionnaires that were collected locally. Population figures refer to registered county level populations, and agricultural employment shares, as well as above compulsory schooling shares of the population are computed using the county totals and subtotals reported in the CITAS data. The control variable for urban population in 1990 is registered residents in urban wards taken from the CITAS population records.[5]

Appendix 5: Construction of Spanning Trees

This section describes the construction of the least cost path and Euclidean spanning tree networks depicted in Figures A.3 and A.4 in the paper. The stated objectives of the NTHS in 1992 were to connect all provincial capitals and cities with an urban registered population above 500,000 and connect targeted nodes to border segments as part of the Asian Highway Network in border provinces. In the 1990 Chinese Population Census, 54 cities correspond to these criteria.[6] The following computation steps have been executed in ESRI's ArcGIS software.

To construct the least cost path spanning tree network depicted in Figure A.3 of the paper, I adapt a simple construction cost function from the transport engineering literature

[4] This difference is sometimes cited as one of the reasons for discrepancies between the national and the sum of province level economic accounts.

[5] This definition of urban population used in the control variable does not directly correspond with the official Chinese administrative definition. Traditionally two characteristics are used to classify urban residents for Chinese official use. The first is that at least one person of the household holds a "non-agricultural occupation" (industry or services). The second is that the sub-county level administrative unit ("zhen"=ward) is classified urban as opposed to rural.

[6] The records of the 1990 Population Census became available for administrative use in 1991, and constituted the highest quality and most recent information about population registries at the ward level ("zhen") across China at the time of decision making for the NTHS in 1992. According to the Chinese administrative definition, the urban registered population of a central city is the sum of households with urban occupations in the wards ("zhen") of the central city county level units (shixiaqu) of the municipality. 1990 Census population and occupation data at the sub-county ward level was provided by the ACASIAN Data Center from archival records held at Griffith University library.

(Jha *et al.*, 2001; Jong and Schonfeld, 2003).[7]

$$c_i = 1 + Slope_i + 25 * Developed_i + 25 * Water_i + 25 * Wetland_i \tag{5.1}$$

c_i is the cost of crossing a pixel of land i, $Developed_i$ indicates whether the pixel is covered by built structures, $Slope_i$ is i's average slope gradient, $Water_i$ and $Wetland_i$ are dummies indicating whether i is covered by water or wetland. This simple specification of the construction cost function implies that shorter and flatter routes will be preferred, while high costs are assigned to crossing water bodies, wetlands, or built structures. I use the remote sensing data on land cover and elevation described in the previous section to compute c_i for a continuous grid of land parcels covering the PR China. For computational feasibility, I reclassify the original resolutions of the elevation and land cover grids from 30 arc seconds (approximately 0.82x0.82 km^2) to 2x2 km^2 grid cells.[8] This yields an isotropic cost surface grid covering the PR China in a rectangle of approximately 4.7 million 2x2 km grid cells. Figure 5.1 provides a graphical illustration of this construction cost surface.

I then proceed to construct least cost highway paths between all 1431 ($\frac{54*53}{2}$) possible bilateral pairs of targeted city nodes. To achieve this, I follow the accumulative cost minimization procedure pioneered by Douglas (1994). The first step is to compute Dijkstra's (1959) optimal route algorithm to identify the least costly path between each one of the 54 nodes and every cell center of the grid covering the PR China's surface. To calculate the cost of moving from the center of an origin cell to the center of one of the eight directly adjacent cells, there are two types of cost functions subject to which Dijkstra's algorithm is computed:

$$c_{od_1} = \frac{c_o + c_{d_1}}{2} * \pi, \text{ and}$$

$$c_{od_2} = \sqrt{2} * \frac{c_o + c_{d2}}{2} * \pi \tag{5.2}$$

where c_{od_1} is the cost of moving from the origin cell to one of four horizontally or vertically adjacent cells, c_{od_2} is the cost of moving to one of four diagonally adjacent cells, c_o, c_{d_1} and c_{d2} are the assigned construction costs of the respective cells, and π is the km cell resolution.

These computations result in 54 separate cumulative cost rasters, each containing about 4.7 million 2x2 km pixels covering the PR China. Each cell is assigned the accumulative cost associated with Dijkstra's optimal route solution between any origin cell on the raster to one of the 54 nodal destinations. In addition, the computations yield 54 separate directional backlink raster files, assigning a code between 1-8 to each cell on the grid that indicates the moving direction from any cell to one of its eight adjacent cells on the identified least cost path to the particular targeted node of the grid.

The accumulative cost rasters and directional backlink rasters for each of the 54 nodes ultimately enable me to construct 1431 hypothetical least cost highway construction paths between all possible nodal connections. In the second stage, I then extract the aggregate construction cost of each possible bilateral connection in order to compute Kruskal's minimum spanning tree algorithm. This algorithm identifies 53 least cost connections that connect each of the 54 targeted cities on a single network. This yields an all-China spanning tree network.

[7] The choice of the cost factor to be 25 is informed by empirical estimates of the per lane mile construction cost of highways relative to bridges. See for example WSDT (2002).

[8] To assure the continuation of rivers, the reclassified grid cells were classified as covered by water if any of the area was classified as water in the higher resolution grid.

The final step to constructing the network depicted in Figure A.3 in the paper is to apply the least cost path algorithm to find least costly connections between capitals of border provinces and segments of China's border. Least costly paths to any segment of the border within the same compass quadrant (NE, SE, SW, NW) as NTHS routes were constructed without imposing *ex ante* restrictions on the end points located on the border.

To construct the straight line spanning tree network depicted in Figure A.4 of the paper, the first is to compute great circle distances between all possible 1431 bilateral connections of the network, which is done by applying the Haversine formula to bilateral coordinate pairs. I then compute Kruskal's algorithm to identify the minimum number of edges that connect all targeted cities subject to the minimum aggregate distance impedance on the network. To account for the fact that Chinese planners construct many more than the minimum spanning tree connections, I re-run Kruskal's algorithm after dividing China into North-Center-South, as well subject to East-Center-West geographical areas.[9] These two additional estimations add 9 bilateral routes in addition to the 53 connections that resulted from the all-China estimation. The final step is to include minimum great circle distance connections from provincial capitals of border provinces to the nearest border segment within the same compass quadrants as NTHS routes.

[9] I define these geographical areas on the basis of six geographic regions with official administrative recognition in China: East, North, North-East, North-West, South-Central, and South-West.

Figure 5.1: Construction cost raster

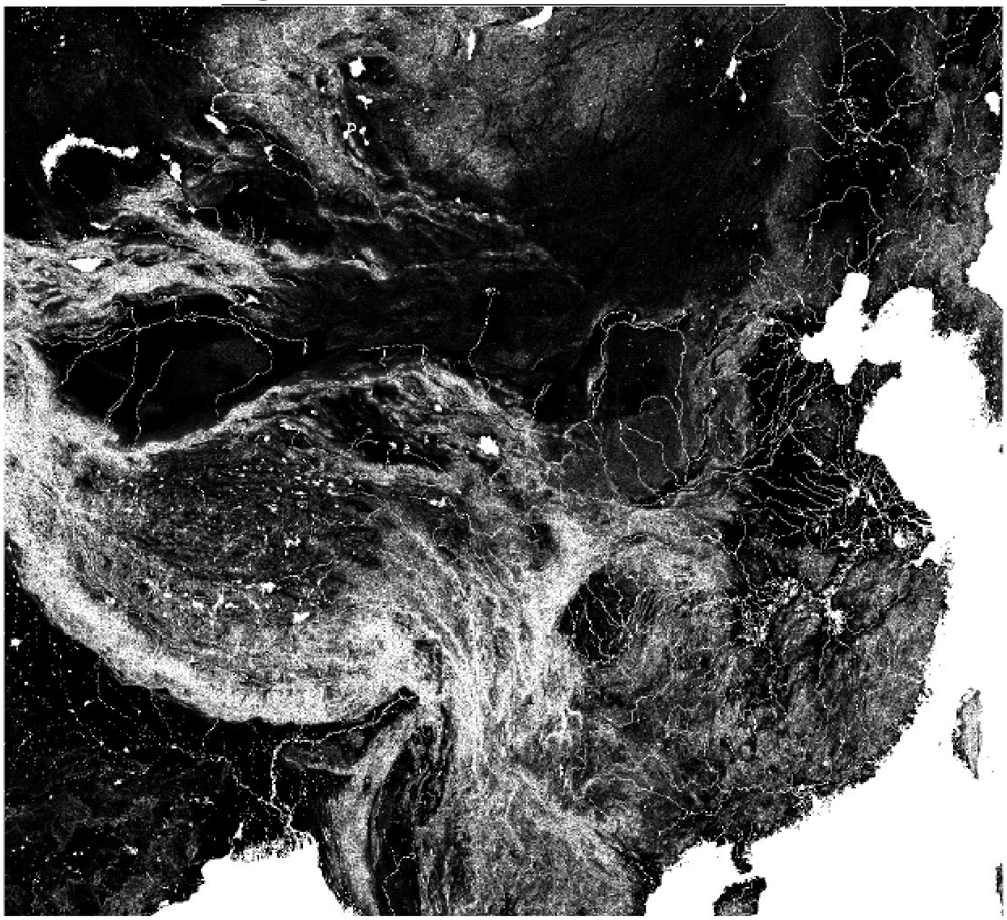

The figure depicts the construction cost raster used as input into the least cost path algorithm. The colour scale ranges from white (very high cost of crossing a parcel of land) to black (very low cost of crossing a square km parcel of land). The cost assignment is based on land gradient (slope) as well as land cover (water, wetlands, and developed land), and described in more detail in the text.

Appendix 6: Network Routing Analysis and Simulations

The final section describes the network routing analysis used to estimate bilateral trade cost changes in the simulations of the structural model. The following computation steps were executed using ESRI's ArcGIS software including the Network Analyst extension. For each parameter assumption about the relative per km cost of NTHS expressways compared to pre-existing national highways that are centered around the baseline estimate of 0.85 (0.75, 0.8, 0.85, 0.90, 0.95), the following optimal network routing analysis is computed twice, once on the basis of the road network in place at the end of 1997, and once after adding the NTHS routes that had been built between mid-1997 and end of 2003. Figure 6.1 depicts the primary and secondary road networks in China for these two comparison periods.

Each of the 1679 county centroids is assigned to the closest segment of the Chinese road network in 1997 or in 2003 as depicted in Figure 6.1. Each segment of the road network system is assigned a km length of the segment as well as a per km cost factor. These cost factors are equal to 1 for national highway segments, 1.25 for provincial highway segments and one of the five aforementioned relative cost factors for NTHS routes existing in 1997 or existing in 2003.

The network routing algorithm then solves the system of 1679 origins and 1679 destinations for the least costly accumulative network cost (sum of products of segment lengths times per km costs) of reaching a given destination from a given origin. The algorithm used to solve this computation problem in the ArcGIS Network Analyst software is Dijkstra's (1959) optimal route algorithm to find the shortest path between an origin and a destination on a weighted graph, where the weights are transport costs assigned to each segment of the network.

This yields a matrix of least costly bilateral network transport costs both for 1997 and 2003, and across the five alternative assumptions about relative per km transportation costs on the network. From these solution matrices, I compute bilateral percentage changes in total bilateral transport costs. These are then used to compute post-NTHS bilateral trade costs matrices where the solution matrix of bilateral percentage transport cost changes is used in combination with the initial calibration of bilateral trade cost levels described in the paper in order to compute the 2003 matrix of bilateral trade freeness parameters $\phi_{ij} = \tau_{ij}^{1-\sigma}$.

Figure 6.1: China's primary and secondary road network before and after the NTHS

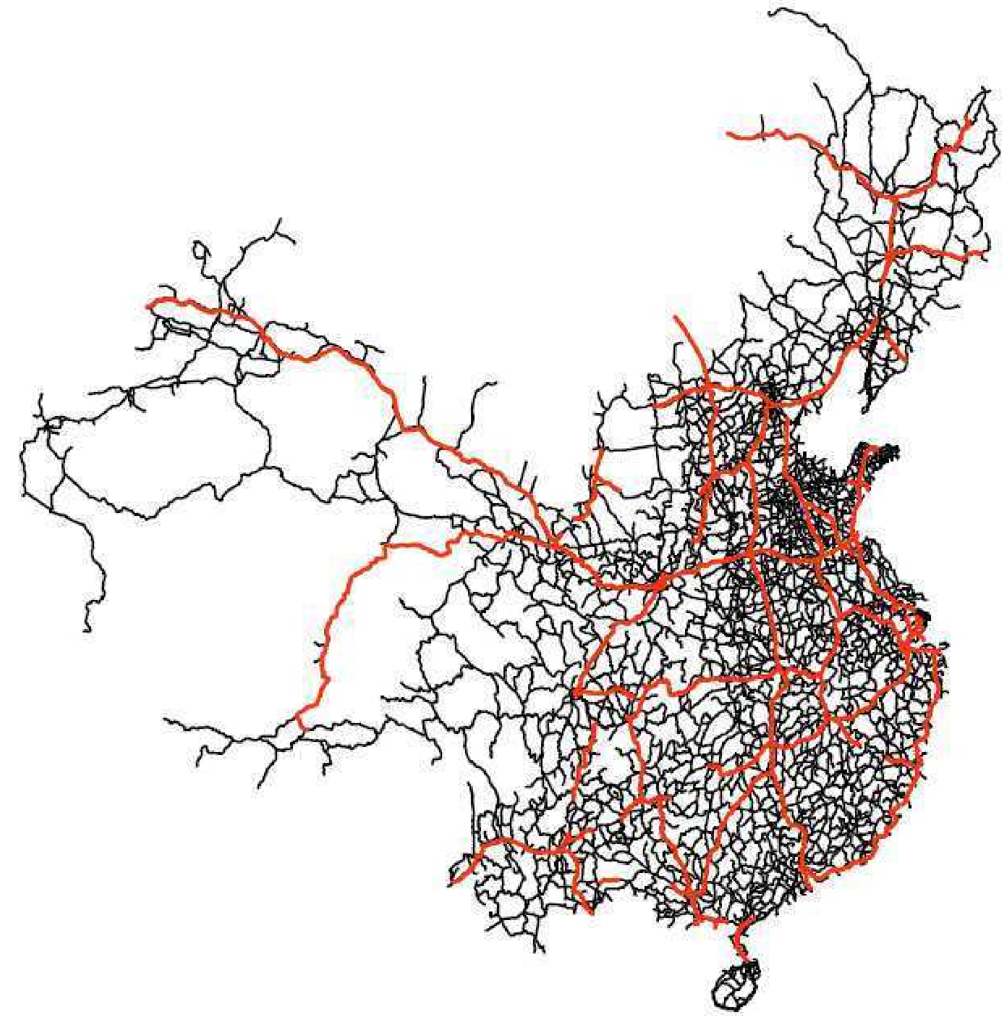

The network in black colour indicates existing national and provincial highways in 1997 as well as a small number of previously built expressways. The network in red colour represents expressway additions that were completed between mid-1997 to end of 2003.

References

Abadie, A. (2003). Semiparametric instrumental variable estimation of treatment response models. *Journal of Econometrics*, 113(2), 231–263.

Angrist, J. & Pischke, J. (2008). *Mostly harmless econometrics: An empiricist's companion.* Princeton Univ Pr.

Lin, S. (2009). *The Rise and Fall of China's Government Revenue.* East Asian Institute, National University of Singapore.

WSDT (2002). *Highway Construction Cost Comparison Survey Final Report.* Washington State Department of Transport.

Chapter 2

Trade Liberalization, the Price of Quality, and Inequality: Evidence from Mexican Store Prices*

Benjamin Faber†

Abstract

This paper considers a novel distributional channel of developing country trade liberalization that operates through differences in cost of living inflation between rich and poor households. Motivated by the observed pattern of vertical differentiation across Mexican households in consumption and plants in production, I consider quality choice as a channel that links differences in consumption baskets between the rich and the poor to differences in plant technologies and, thus, relative price changes. Guided by this theoretical framework, I draw on a new collection of microdata covering Mexican households, plants, and stores to estimate this channel empirically in the context of Mexico's trade liberalization under NAFTA. In particular, I present evidence that cheaper access to US inputs reduces the relative price of higher quality products in Mexican cities. In turn, because richer households consume higher product quality, I find that this relative price effect has led to a significant increase in Mexican real income inequality due to NAFTA over the period 1994-2000.

Keywords: Market integration; inequality; product quality
JEL Classification: D63; F13; F15; O24

*I am grateful to Robin Burgess, Dave Donaldson, Henry Overman, Steve Redding and Daniel Sturm for their advice and support on this project. I would like to thank Etienne Gagnon at the Federal Reserve Board in Washington D.C. for allowing me access to the more recent years of the Mexican store price data, and Gabriel Romero Velasco at INEGI in Aguascalientes, Mexico for his tireless support on the Mexican plant data. I am grateful to Pol Antras, David Atkin, Michael Best, Francisco Costa, Arnaud Costinot, Swati Dhingra, Esther Duflo, Jason Garred, Horacio Larreguy, Guy Michaels, John Morrow, Thomas Sampson and Anthony Venables for very helpful conversations. I would also like to thank seminar participants at Harvard, LSE, MIT and UCL for comments and suggestions. This project was supported by research funding from the UK Economic and Social Research Council.

†London School of Economics and Centre for Economic Performance; Email: b.s.faber@lse.ac.uk.

1 Introduction

How does globalization affect inequality? The canonical approach to analyze this question is through the lens of Stolper-Samuelson, whereby trade across homogeneous final goods affects the relative returns to domestic factors of production. This paper departs from this tradition in three ways. First, I consider household price indexes in the denominator of real income, rather than nominal incomes in the numerator, as a channel through which trade liberalization can affect inequality. Second, I emphasize access to imported inputs, rather than final consumer goods. Third, I analyze relative price changes across vertically differentiated products within disaggregate product groups, rather than across sectors.

In particular, I consider product quality choice by households in consumption and by plants in production as a channel that links differences in the consumption baskets between the rich and the poor to differences in imported input shares in production. Trade liberalization can thus affect cost of living inflation asymmetrically across the income distribution because product quality differentiation gives rise to a correlation between differences in household expenditure shares in consumption and differences in plant technologies in production. Drawing on a new collection of microdata covering Mexican households, plants, and stores, I test this channel empirically in the context of Mexican trade liberalization under NAFTA.

The paper's main contributions can be summarized as follows. Motivated by the observed pattern of vertical differentiation in Mexican microdata on household consumption and plant production at the beginning of NAFTA in 1994, I propose a model of quality choice in a setting with heterogeneous consumers and producers. Guided by this theoretical framework, I present evidence that cheaper access to US intermediate inputs reduces the relative price of higher quality products in Mexican cities. In turn, because richer households consume higher product quality, I find that this relative price effect has led to a significant increase in Mexican real income inequality due to NAFTA over the period 1994-2000. Finally, I find that the observed pattern of quality sorting also has one more general implication for real income inequality. The estimations suggest that differences in household quality choices translate into differences in weighted average plant productivities, so that the consumption baskets of the poor embody significantly higher quality adjusted prices compared to the rich. This increases real income inequality in the cross-section of households compared to a world under conventional assumptions without quality sorting by households and plants.

The paper's analysis is motivated by three more general observations about inequality and trade. The first of these concerns what we already know about real income inequality. There is pervasive evidence that relative consumer price changes can have significant implications for real income inequality due to differences in cost of living inflation between the rich and the poor (e.g. Muellbauer, 1974; Deaton, 2003; Moretti, 2011; Handbury, 2012).[1] However, when moving from the measurement of inequality to analyzing the effects

[1] Rather than focusing on prices over time, Handbury (2012) evaluates relative prices across locations.

of policy or market shocks on inequality, economists have mainly focused on nominal wages while treating differences in cost of living inflation as exogenous.[2] The second and third observations concern developing country imports. Developing country imports are strongly dominated by intermediate inputs, rather than final consumer goods. The case of Mexico is no exception to this, with roughly 90% of total imports from the US between 1994-2000 accounted for by intermediates. Finally, the majority of the variation in use of imported inputs across the product space appears to be between plants within disaggregate product groups, rather than across sectoral averages.[3] Taken together, these three insights suggest i) that relative consumer prices matter for inequality, and ii) that access to imported inputs and price changes across varieties within product groups play a significant role in capturing the consequences of developing country trade liberalization.

The analysis proceeds in several steps. I begin by documenting a set of stylized facts about vertical differentiation in Mexican consumption and production at the beginning of NAFTA in 1994. A meaningful analysis of relative prices, production technologies, and household consumption within consumer product groups requires data on unit values (prices per physical unit), plant characteristics, and household expenditures at a very fine level of product aggregation. I draw on Mexican plant surveys, including rich product line level information, in combination with household consumption surveys, including individual purchase prices and quantities, to document a set of relationships between unit values and plant characteristics in production, and between unit values and household characteristics in consumption. In particular, the Mexican microdata suggest that: 1) Plant product line unit values are increasing in imported input shares in production; 2) plant product line unit values are increasing in product sales; and 3) household purchase unit values are increasing in household income in consumption.

To capture these observed moments in the microdata, I propose a model of quality choice by households in consumption and plants in production. The model serves two main objectives. First, it formalizes a product quality interpretation of the stylized facts. Second, it yields testable predictions on NAFTA's effect on Mexican consumer prices, and guides the estimation of the cost of living implications of these relative price changes across the Mexican income distribution. While several existing theoretical frameworks have been proposed to capture separately either the consumption side or the production side of the stylized facts, the two have so far not been considered in a unified framework of quality choice.[4] I show that the introduction of heterogeneous household quality evaluations into

[2]For discussions of the literature on the causes of inequality see for example Leamer (1996), Aghion and Williamson (1998), and Goldin and Katz (2008). Notable exceptions to the focus on nominal incomes are Deaton (1989), Porto (2006), Broda and Romalis (2008), and Cortes (2008). See discussion of related literature at the end of this section.

[3]Appendix Figures A.1.1 and A.1.2 provide a graphical illustration of these insights.

[4]Existing models have focused either on quality choice across households while abstracting from plant heterogeneity (e.g. Choi et al., 2009; Fajgelbaum et al., 2011; Handbury, 2012), or on quality choice across plants while abstracting from household heterogeneity (e.g. Johnson, 2011; Kugler and Verhoogen, 2011; Feenstra and Romalis, 2012). See the discussion of related literature at the end of this section for further

a model of quality choice by heterogeneous plants poses one key challenge, which is that *physical* product quality is distinct from *perceived* product quality in the market place. The definition of this distinction is of interest more generally because it reveals the concept of product quality that has been implicitly estimated from product prices and market shares in a prominent strand of empirical work in industrial organization (e.g. Berry, 1994; Berry, Levinsohn and Pakes, 1995) and international trade (e.g. Khandelwal, 2010; Hallak and Schott, 2011; Feenstra and Romalis, 2012), when acknowledging that we live in a world with non-homothetic tastes.

The theoretical framework is then used to guide the empirical estimation in the three remaining sections. Section 5 draws on the barcode level store price microdata of the Mexican Consumer Price Index to empirically test the model's predictions on the effect of NAFTA's tariff cuts on Mexican consumer prices. Section 6 estimates the cost of living implications of these relative price effects across the Mexican income distribution by drawing on observable moments in the household consumption microdata. Finally, Section 7 imposes additional parameter assumptions in order to estimate differences in product quality and quality adjusted prices across the income distribution.

In support of the predictions, the store price regressions suggest that products with initially higher unit values experience a stronger reduction in their relative price within product groups that are subject to larger tariff cuts on their intermediate inputs over the period 1994-2000. That is, the relative price of initially more expensive items decreases in product groups that gain cheaper access to US inputs. These results are based on a novel identification strategy to relate import access to domestic outcomes. It is a common concern that tariff changes may be correlated with omitted factors that also affect mean sectoral outcomes.[5] Focusing instead on relative price changes within disaggregated product groups (e.g. antibiotic pills, electric irons) allows me to rely on the much weaker identifying assumption that tariff cuts are plausibly exogenous at the level of individual barcode product lines, especially in the case of intermediate inputs which are shared throughout the domestic economy. To address potential concerns that tariff cuts may have been targeted at particular segments of the plant distribution within product groups, I also propose an instrumental variable strategy that can be applied more generally in the context of input tariff changes. The instrument is based on the insight that tariff targeting should be of much less concern for a subset of intermediate inputs, such as basic chemicals, which are used widely across domestic output sectors. Finally, I exploit the richness of the collected store price microdata to report three different placebo falsification tests.

The model also makes predictions about the heterogeneity of the relative price effect of input tariff cuts. In particular, the observed effect should be driven by differentiated product groups in which initial price differences provide stronger signals about differences

references.

[5]See for example discussions in Goldberg and Pavcnik (2005; 2007). See Slaughter (1998) for a survey of the empirical literature on trade induced relative price changes across industries.

in quality and plant technologies. To test this prediction empirically, I follow a two stage procedure. In the first stage, I use the model's estimation equation for sectoral scopes of quality differentiation in terms of observable moments in the plant microdata. In the second stage, regression results then confirm that the observed average effect of input tariff cuts on relative store prices is driven by product groups that have been estimated to be differentiated in the Mexican plant microdata. A final prediction concerns the effect of access to imported inputs on the reallocation of market shares towards higher quality product lines. To test this prediction, I draw on detailed monthly listings of product entry and exit in the store price surveys, and present evidence in support of this effect.

To evaluate the consequences of NAFTA's observed store price effects for differences in household cost of living inflation, the model yields a convenient estimation equation in terms of observable moments in the household consumption microdata. I discuss the two key assumptions underlying this expression and outline the empirical strategy to estimate it from the data. I find that the average tariff cut under NAFTA between 1993-2000 (12 percentage points) has led to at least 1.7 percentage points higher cost of living inflation in tradable consumption of the poorest urban income quintile compared to the richest over the six year period 1994-2000. This estimate increases to 2.6 percentage points in what I refer to as the baseline specification, and to 3.9 percentage points in what I refer to as the upper bound estimate. In terms of real income inequality, these effects are equivalent to approximately 25-55% of the total observed increase in nominal income inequality among the same groups of households over the period 1994-2000.[6]

In the final section, I impose a parameter assumption on the elasticity of substitution in demand in order to estimate differences in weighted average quality as well as quality adjusted prices across households. In particular, the model captures the observed pattern in the Mexican microdata which suggest that more productive plants sort into higher quality products. The implication is that differences in household quality choices translate into differences in quality adjusted prices. I estimate that the poorest quintile of urban Mexicans consume 20-50% lower weighted average product quality among differentiated product groups. In turn, I find that these observed consumption differences translate into at least 2-7% lower weighted average plant productivities, and thus higher quality adjusted prices, in tradable consumption of the poorest quintile compared to the richest.

The paper relates and contributes to several strands of literature. It is related to empirical work on trade and inequality in developing countries. A comprehensive review of this literature is given in Goldberg and Pavcnik (2007), and more recent contributions include Verhoogen (2008) and Topalova (2010). The focus of this literature has been on trade induced differences in nominal income growth across skill or income groups. This paper, on the other hand, analyzes a distributional channel of developing country trade liberalization that links changes in the relative price of quality to differences in household cost of living

[6] The comparison is adjusting for the fact that tradable consumption accounts for 54% of total consumption in 1994.

inflation.[7]

There are a number of notable exceptions to the focus on nominal income inequality. Porto (2006) combines scheduled Argentinian tariff changes under Mercosur with household expenditure shares across seven consumption sectors to predict household inflation differences. Broda and Romalis (2008) analyze the link between consumer good imports from China and household inflation differences using homescanner data in the US. Outside the focus on trade, Cortes (2008) analyzes the price index implications of low-skilled immigration in US metropolitan areas, and earlier work by Deaton (1989) predicts the cost of living implications of agricultural price changes using household consumption surveys.[8] To the best of my knowledge, this paper is the first to i) look at the cost of living implications of relative price changes within product groups in a developing country context, ii) empirically estimate these relative price changes in the context of a major trade liberalization episode, and iii) guide this analysis within a theoretical framework of quality choice by heterogeneous households and plants.

The paper also relates to recent contributions on quality choice in a setting with ex ante heterogeneous firms (Mandel, 2010; Johnson, 2011; Kugler and Verhoogen, 2011; Feenstra and Romalis, 2012; Gervais, 2012). This paper introduces quality choice by heterogeneous households into this setting and draws attention to, and estimates empirically the distributional implications that arise when differences in consumption baskets across the income distribution are linked to differences in plant technologies through quality choice.

Finally, the paper is related to existing literature on non-homotheticity in international trade. Non-homothetic preferences were originally introduced to explain part of the variation of cross-country trade flows left unaccounted for by neoclassical trade theory (Markusen, 1986; Bowen *et al.*, 1987; Trefler, 1995; Matsuyama, 2000; Choi *et al.*, 2009; Fieler, 2011).[9] Rather than focusing on the consequences for trade flows, this paper analyzes the implications of non-homotheticity for the distributional effects of trade in a developing country.[10]

The remainder of the paper is structured as follows. Section 2 describes the background and data. Section 3 documents stylized facts about vertical differentiation in Mexican consumption and production at the beginning of NAFTA in 1994. Section 4 presents the model. Section 5 presents the empirical estimation of NAFTA's effect on Mexican consumer prices. Section 6 presents the estimation of the cost of living implications of these relative price effects. Section 7 presents the estimation of differences in quality and quality adjusted

[7] The paper also relates to recent literature on the effects of access to imported inputs in a developing country context (e.g. Amiti and Konings, 2007; Goldberg et al., 2010; Halpern et al., 2011).

[8] More recent contributions on the measurement of cost of living inflation in developing countries include Deaton and Dupriez (2011) and Li (2012).

[9] Atkin (2010) investigates the implications of regional taste differences for the gains from trade, rather than trade flows.

[10] In a recent theoretical contribution, Fajgelbaum *et al.* (2011) consider the implications of non-homotheticity in the context of both across and within country income distributions. While focusing on trade in final goods, their model yields predictions that are in line with the empirical findings presented in this paper, that trade increases inequality in the poorer country through inflation differences across the domestic income distribution.

prices across the Mexican income distribution. Section 8 concludes.

2 Background and Data

2.1 Mexican Trade Liberalization

Mexican trade liberalization began as part of government stabilization efforts in response to the severe economic crisis at the beginning of the 1980s. When Mexico joined the General Agreement on Tariffs and Trade (GATT) in 1986, it initially agreed to bind tariffs at a ceiling of 50 percent. In December of 1987 the government then implemented another major consolidation of its tariff schedule whereby all non-agricultural import tariffs were set at either zero, five, ten, fifteen, or twenty percent (Kate, 1992). Following this first wave of liberalization, the Mexican tariff schedule remained largely unchanged between the end of the 1980s until the beginning of NAFTA in January 1994.

NAFTA represented a significant second wave of Mexican import tariff reductions. While in 1993 only 10% of manufacturing imports from the US fell into a tariff category of 15% or less, this fraction increased to 60% in January of 1994 (Lopez-Cordova, 2002). In contrast, NAFTA had a smaller effect on US tariffs on Mexican exports as these were at already low levels before NAFTA took effect.[11] Figure A.2 provides an illustration of average Mexican tariff changes on US imports and their sectoral variation over the period 1993-2000.

Concerning the importance of other trade partners, US imports have consistently accounted for 75-80% of total Mexican imports during the 1990s. In particular, the period under study precedes China's admission to the World Trade Organization in 2001 and the subsequent surge of Chinese imports into Mexico (Iacovone *et al.*, 2010).

Any analysis of NAFTA's consequences in Mexico must address the empirical challenge that the beginning of NAFTA coincided with a severe economic crisis that unfolded in Mexico in 1995, the adverse consequences of which are apparent in Mexican real income data until the beginning of the following decade (Attanasio and Binelli, 2010). As discussed in more detail in the following sections, the present empirical analysis addresses such concerns by focusing on parts of the variation in relative price changes that are plausibly unrelated to the consequences of economy wide macroeconomic shocks.

2.2 Data

The following subsections provide a description of the datasets used in this paper, and further details can be found in the Appendix.

[11] The average export tariff was approximately 2% in December 1993.

2.2.1 Central Bank Store Price Surveys

The great majority of countries, including Mexico, are subscribers to the ILO/IMF Consumer Price Index Dissemination Standard. This manual imposes a clear set of rules on how to compile and process data in order to report national consumer price inflation. The backbone of national CPI reporting are central bank store price surveys that are collected throughout the country and usually at several times during each month.

In a fortunate turn of events, the Articulo 20-Bis of the Mexican Codigo Fiscal de la Federacion requires the central bank since January of 1989 to publish the store price microdata on a monthly basis in the official government gazette, the Diario Oficial de la Federacion.[12] These publications are phone book like listings of individual city-store-barcode product combinations and their price quotes in a given month.[13]

Starting from 1989, each month of data contains approximately 30,000 individual price quotes across 35 Mexican cities and 284 product groups covering non-durables (e.g. Salchicha sausages, sanitary paper towels, antibiotic pills), durables (e.g. electric water boilers, bicycles), as well as services (e.g. language courses, taxi rides).[14] For the empirical analysis, I compute average price quotes of individual items across three months in the third quarters of 1989, 1993, 1994, 2000, and the first quarter of 1995.[15]

These price data have a number of notable features. First, the survey is intended to capture a representative sample of Mexican household consumption and covers street vendors, markets, convenience and specialized stores, as well as supermarkets and department stores across cities. Second, any change in the presentation, appearance, size, modality, model number or otherwise is reported in an appendix of the monthly publications in the Diario Oficial as a product item substitution. The objective of the Mexican central bank is to compute price inflation for identical product items over time. This is to say that what I refer to as "persistent" product series are identical barcode products in the identical store over time.[16] This detailed documentation also allows me to test predictions on product

[12]The commonly cited reason is concern of political influence on the computation of consumer price inflation at that point in Mexican history.

[13]While the actual product barcode is not reported, the detailed product descriptions including brand, product name, pack size, model number, and modalities (e.g. color, packaging type) provide an equivalent level of product identification.

[14]For product groups in food and beverages the reported monthly price quotes are averages across 2-4 monthly price quotes for each item. The number of cities and product groups increased in a revision in March 1995. The stated figures refer to cities and product groups that were consistently covered both before and after the revision in 1995.

[15]The price quotes of the latter three periods were provided by courtesy of Etienne Gagnon at the Federal Reserve Board in Washington D.C. A detailed description of this dataset can be found in Gagnon (2009). For the third quarter price quotes in 1989 as well as 1993, I obtain copies of archival records contained in the Archivo General de la Nacion in Mexico City, and digitize these monthly price quotes by double blind data entry. The percentage of non-identical entries was approximately 1%. These cells were then double checked and corrected by hand.

[16]To this end, I digitize the complete series of product substitution appendices from January 1989 to the end of 1993 and I obtain the more recent substitution listings between 1994-2000 by courtesy of Gagnon (2009).

replacements and additions in addition to relative price changes of persistent barcode items.

Third, price quotes are reported in prices per common physical unit for the majority of product groups. For product groups where this is not consistently the case (e.g. measured per pack of toilet paper, or measured per bottle of body lotion), I either clean the data by hand to convert it to common physical units (e.g. per roll of toilet paper reported in the product description, or per 100ml of body lotion), or I exclude the product group from the estimations where such a correction is not feasible (e.g. tortillas or clothing where reported prices are based on store sample averages within a city so that changes in product items are unobserved).

Finally, one important limitation of the store price data is that while each city-store-barcode item has a unique identifier code, the individual store identifiers cannot be recovered from these item codes.[17] Potential estimation concerns arising from this limitation will be addressed in detail in the empirical analysis.

The main estimation sample is a panel of individual city-store-barcode items within 144 processed tradable product groups that i) report individual barcode unit values rather than store sample average prices, and ii) could be matched to product groups in the plant production microdata that is described in the following subsection.[18] Table A.1 presents descriptive statistics of the store price microdata together with a breakdown of the estimation sample's coverage of total household expenditure among urban Mexicans in 1994 and 2000.

2.2.2 Manufacturing Establishment Surveys

There are two general empirical challenges when empirically investigating vertical product differentiation in plant microdata. The first is that the majority of manufacturing establishment surveys do not report physical output quantities in combination with sales revenues to compute output unit values. The second is that product quality differentiation is empirically meaningful only at a very fine level of product aggregation. Most plant surveys report two digit (e.g. food processing), four digit (e.g. meat processing), or sometimes six digit (e.g. meat products except poultry) industrial classifications, which would be insufficient to match the detailed product groups that are present in the store price surveys described above.

Fortunately, the microdata reported in the Mexican Encuesta Industrial Mensual (EIM) make it possible to address both of these challenges. Starting in 1994, the survey reports monthly physical output in combination with sales at the level of several thousands of product groups within 203 six digit CMAP manufacturing sectors. Plants on average report output and sales across several products, so that the level of aggregation present in the

[17]The original reason for encoding the store identifier in the Diario publications was confidentiality concerns. Unfortunately, the correspondence table between the published and the actual store identifiers appears to have been lost for the period before 2002.

[18]See Appendix for product group descriptions and the correspondence between product groups across the different datasets.

data can be thought of as individual product lines within an establishment. These product descriptions are used to match the finest available level of product groups in the store price surveys to the product classifications in the plant microdata, resulting in 144 processed tradable product groups (see Appendix for details).

The second plant dataset that I draw on in the analysis is the Encuesta Industrial Anual (EIA) which covers the identical plants at annual intervals. In particular, I use the EIA to complement the EIM data with annual plant level information on the use of imported inputs as well as employment. Both the EIM and the EIA microdata are administered by the Instituto Nacional de Estadistica, Geografia, e Informatica (INEGI). I obtained access to the confidential microdata for 12 months of the EIM data in 1994 and the matching annual records contained in the EIA. These establishment surveys cover all manufacturing production sectors and represent roughly 85% of total Mexican manufacturing output. The data do not cover the universe of Mexican production establishments as the surveys typically omit the tail of small producers (INEGI, 2000). Table A.2 presents descriptive statistics of the 1994 plant microdata.

2.2.3 Household Consumption Surveys

I use the microdata of the Mexican national household consumption survey in 1994 (Encuesta Nacional de Ingresos y Gastos de los Hogares, ENIGH) for information on per capita household incomes, expenditure weights and unit purchase values across 255 processed tradables product groups that overlap with the plant production product groups of the EIA/EIM plant data (see Appendix for details). These surveys are administered by INEGI from where I obtain access to the data. To be consistent with the urban only coverage of the store price surveys, I only use data on households in urban classified municipalities.[19]

There are several notable features of the household consumer surveys. First, they cover the whole of Mexico and report nationally representative sample weights. Second, they are collected to represent total household consumption expenditure during the third quarter of 1994 which coincides with the collection of the central bank store price series. Third, they provide a rich breakdown of product groups into several hundred product codes covering all types of consumption expenditure. Fourth, they report every single transaction within a product group made by members of the household. Fifth, they report unit values (e.g. per kilogram or per liter) for 118 product groups that pertain to food and beverages and tobacco expenditures out of the 255 total product categories that could be matched to product groups that are also present in the EIA/EIM plant data. Finally, the surveys report the store type linked to every single transaction. The store types include street vendors, markets, convenience and specialized stores, and supermarkets and department stores. Table A.1 presents an overview of the shares of consumption expenditures captured by the processed tradables estimation sample, and Table A.3 presents descriptive statistics

[19]In the ENIGH 1994 survey, these are municipalities with more than 2500 residents.

for the urban household ENIGH sample in 1994.

2.2.4 Input Tariffs

Intermediate input tariff changes are computed at the four digit industrial classification (NAICS) of the Mexican 2003 input output table.[20] I use total (direct and indirect) requirement coefficients from the import matrix to compute the weighted average US tariff changes across four digit input sectors for each destination sector. In the instrumental variable regressions, I also make use of direct and indirect requirement coefficients of the total use matrix (not just imports).

Out of a total of 101 tradable four digit NAICS sectors, I focus attention on 75 manufacturing input sectors to compute weighted average intermediate import tariff changes for each of the 101 tradable destination sectors. The focus on manufacturing inputs is due to concerns that agricultural and primary imports are subject to significant non-tariff barriers (e.g. Krueger, 1999). The Appendix provides the concordance table that was used to match the 144 product groups in the store price estimation sample to four digit destination sectors in the Mexican input output table.

Tariff changes at the four digit NAICS level are based on average tariff changes across eight digit tariff lines from the Mexican Secretaria de Economia. NAFTA tariff changes in the estimations refer to the difference between average applied rates during the year 2000 and December 1993. For tariff changes on final good imports and exports over the same period, I match the 144 store product groups to the eight or six digit levels of tariff lines.[21] Figure A.2 provides an illustration of average Mexican tariff changes on US imports and their cross-sectoral variation over the period 1993-2000.

3 Stylized Facts about Vertical Differentiation in Mexican Consumption and Production

This section draws on the Mexican plant and household microdata to document a set of stylized facts about vertical differentiation in production and consumption at the beginning of NAFTA in 1994. These insights serve to motivate the theoretical framework in Section 4 and the empirical estimations in Sections 5-7.

Plant Product Line Unit Values Increase in Imported Input Shares: Figure 1 depicts the first stylized fact. The graph plots the relationship between deviations of product line log unit values (prices per common physical unit) and plant level imported input shares.

[20] This is the most recent available Mexican IO table since 1979. A data request had to be filed at INEGI in order to obtain the four digit break up of the Mexican IO table.

[21] While the available input output information dictates a four digit aggregation for intermediate tariff changes, this constraint does not apply to final product tariff changes. Six or eight digit matches depend on the levels of product aggregation in store surveys relative to tariff lines. Data on US applied tariff rates is taken from Feenstra and Romalis (2002).

Estimations are based on 2656 plants reporting across 9163 unique product lines in 1018 product groups pertaining to 79 six digit manufacturing sectors that produce consumer goods. The unit value-import share elasticity is positive and statistically significant at the 1% level. On average higher unit values within disaggregated product groups embody larger shares of imported inputs in a statistically significant way.

Figure 1: Plant Product Line Unit Values and Imported Input Shares in 1994

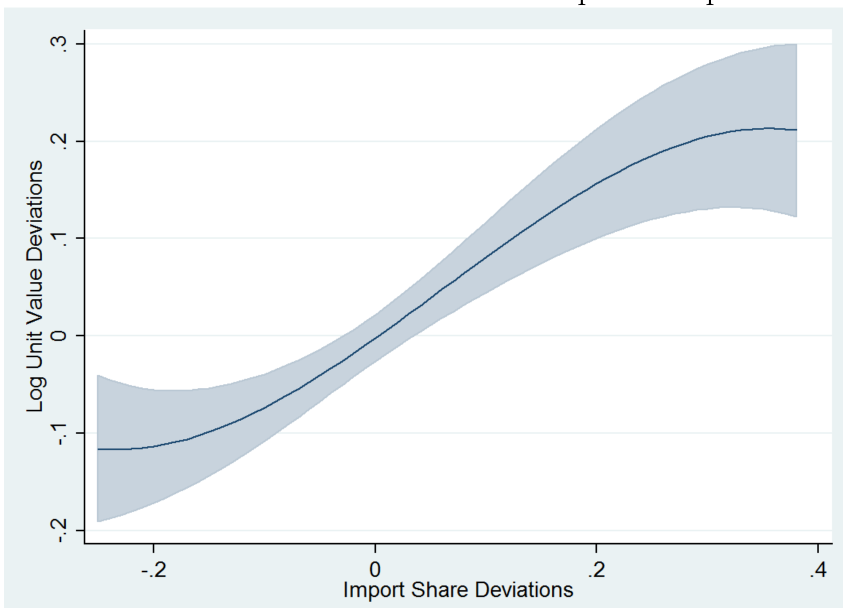

The fitted relationship corresponds to the best fitting polynomial functional form according to the Akaike Information Criterion (AIC). Standard errors are clustered at the product level and the shaded area indicates 95% confidence intervals. The y-axis depicts the residuals of a regression of log monthly product line unit values on month-by-product and month-by-state fixed effects. Estimations are based 2656 plants in 79 six digit final good manufacturing sectors and 1018 products over 12 months in 1994. The number of observations is 94741. The x-axis depicts the residuals of a regression of annual 1994 plant level imported input shares on product and state fixed effects. The number of observations in this regression is equal to the number of unique product lines (9163). The bottom and top 0.5% on the x-axis are excluded from the graph.

Plant Product Line Unit Values Increase in Market Shares: Figure 2 depicts the second stylized fact. The graph plots the relationship between deviations of product line log unit values and product line log sales from month-by-product and month-by-state means during 12 months in 1994. The estimation is based on the identical sample of plants and product lines reported for the previous figure. The unit value-sales elasticity is positive, close to log linear, and statistically significant at the 1% level. On average higher unit values within disaggregated product groups embody larger market shares in a statistically significant way.

Figure 2: Plant Product Line Unit Values and Product Market Shares in 1994

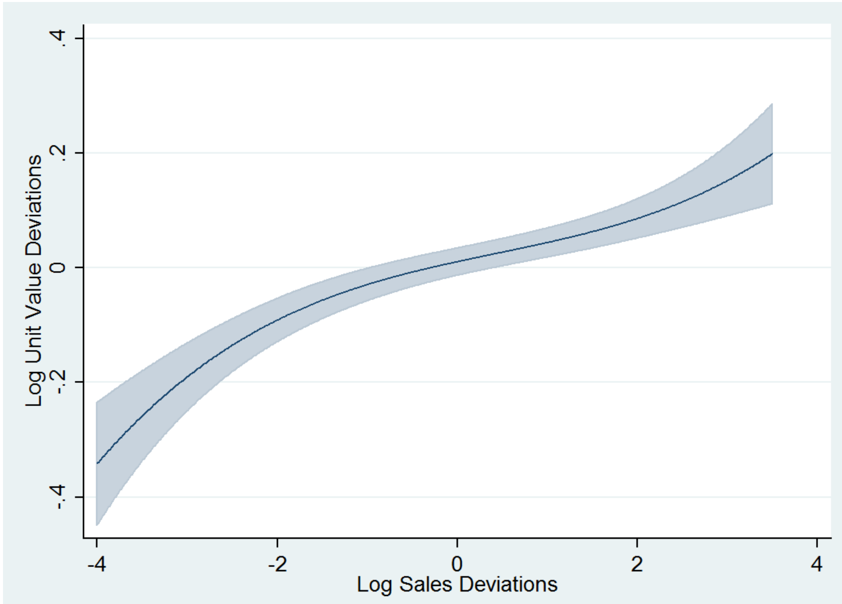

The fitted relationship corresponds to the best fitting polynomial functional form according to the Akaike Information Criterion (AIC). Standard errors are clustered at the product level and the shaded area indicates 95% confidence intervals. The y-axis and the x-axis depict the residuals of two regressions of log product-line unit values or log product line sales on month-by-product and month-by-state fixed effects. Estimations are based 2656 plants in 79 six digit final good manufacturing sectors and 1018 products over 12 months in 1994. The number of observations is 94741. The bottom and top 0.5% on the x-axis are excluded from the graph.

Household Purchase Unit Values Increase in Household Income: Figure 3 depicts the third stylized fact. The graph plots the relationship between deviations of household weighted average log unit values from city-by-product-by-store type means and deviations of log household incomes from the national mean in 1994. The weights are given by households expenditure weights attached to each reported purchase. Reported store types are markets, street vendors, convenience and specialized stores, and supermarkets and department stores.[22] For purchases in the same city-by-product-by-store type cell, unit values of the average household expenditure items are statistically significantly increasing in household per capita incomes.

[22]In Section 6 I report the estimated price gaps between rich and poor households both before and after the inclusion of store type fixed effects in order to learn about the potential role of heterogeneous store markups for identical barcode items in Figure 3. As discussed in more detail, the inclusion of store type fixed effects increases, rather than decreases, the estimated unit value differences.

Figure 3: Household Purchase Unit Values and Household Income in 1994

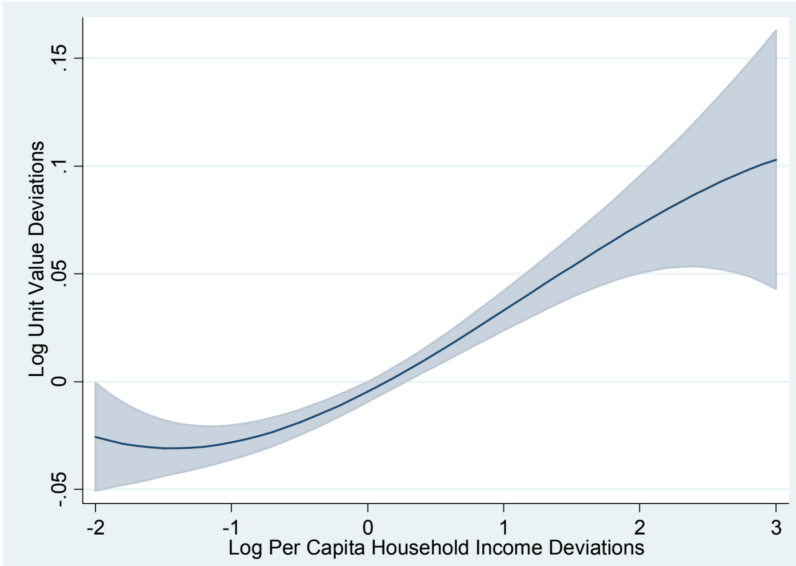

The fitted relationship corresponds to the best fitting polynomial functional form according to the Akaike Information Criterion (AIC). Standard errors are clustered at the municipality level and the shaded area indicates 95% confidence intervals. The y-axis depicts the residuals of a regression of log unit purchase values on city-by-product-by-store type fixed effects. These residuals are then averaged at the household level using reported expenditure weights. The x-axis depicts mean deviations of log household per capita incomes. Estimations are based on urban Mexican households in 1994 and subject to nationally representative sample weights. The bottom and top 0.5% on the x-axis are excluded from the graph.

To summarize, these stylized facts motivate a theoretical framework in which differences in household expenditures across the income distribution are linked to differences in plant technologies through vertical differentiation in consumption and production. While the documented moments in the Mexican microdata are neither individually novel[23], nor particularly suprising, they have so far not been been considered together in a unified framework. The remainder of the paper has the two-fold objective to formalize a product quality interpretation of the documented stylized facts, and to empirically test the implications for welfare inequality that arise from this setting in the context of NAFTA in Mexico.

4 Theoretical Framework

This section proposes a model of quality choice in a setting with heterogeneous households in consumption and heterogeneous plants in production. The model serves two main objectives. First, it formalizes a product quality interpretation of the documented stylized facts. Second, it yields testable predictions on NAFTA's effect on Mexican consumer prices, and guides the estimation of the cost of living consequences of these relative price effects across the Mexican income distribution.

[23]On the production side, the presented results confirm recent findings in Kugler and Verhoogen (2009; 2011) from similarly rich Colombian plant microdata. On the consumption side, for example Deaton (1997) discusses evidence that in the same village unit values within agricultural product groups are increasing in household incomes.

While several existing theoretical frameworks have been proposed to capture separately either the consumption side of quality choice in Figure 3, or the production side in Figures 1 and 2, the two have not been modeled in a unified framework. In this paper, I show that the introduction of heterogeneous household quality evaluations into a model of quality choice across heterogeneous plants poses one key challenge, which is that *physical* product quality is distinct from *perceived* product quality in the market place. The definition of this distinction is of interest more generally because it reveals the concept of product quality that has been implicitly estimated from product prices and market shares in a prominent strand of empirical work in industrial organization (e.g. Berry, 1994) and international trade (e.g. Khandelwal, 2010; Hallak and Schott, 2011; Feenstra and Romalis, 2012) when acknowledging that we live in a world with non-homothetic tastes (Figure 3).

The key features of the model are as follows. On the consumption side, I introduce non-homothetic preferences that allow household quality choice to differ across the income distribution. On the production side, I follow Kugler and Verhoogen (2011) who propose complementarity between plant efficiency and input quality in the production of output quality, and introduce the assumption that input quality is increasing in the use of imported inputs. The following provides a summary of the key features of the model while a more detailed exposition is provided in the Appendix.

4.1 Physical and Perceived Product Quality under Non-homothetic Preferences

A household h's preferences are given by a two-tier Dixit-Stiglitz utility function in which the upper tier is Cobb Douglas across product groups denoted by subscript k, while the subutility index U_{hk} is a CES function over varieties denoted by subscript i within the product group.[24]

$$U_h = \int_{k=0}^{K} U_{hk}^{\mu_{hk}} \mathrm{dk} \quad U_{hk} = \left(\int_{i=0}^{I_k} (q_{ki}^{\varphi_h} x_{hki})^{1-1/\sigma} \, di \right)^{\frac{1}{1-1/\sigma}} \qquad 0<\mu_{hk}<1 \quad \sigma>1 \tag{1}$$

For ease of exposition, product group subscripts k are suppressed in the remainder of this subsection. Household utility is a function of physical units consumed, x_{hi}, and a variety's quality $q_i^{\varphi_h}$, where $\varphi_h>0$ is a household specific taste parameter that determines the intensity of preferences for product quality. Product quality thus enters as a shift in utility derived from consuming a given amount of physical units, and the extent of this shift is allowed to vary across household valuations of quality. To introduce non-homotheticity across the

[24]While the Dixit-Stiglitz structure is standard and will be convenient to solve the model, all results carry through to, for example, a nested logit demand structure as long as the source of non-homotheticity is modeled in the same way as in (1), namely through household quality evaluations rather than heterogeneity in price elasticities. In a recent contribution, Handbury (2012) draws on US homescanner data to separately estimate non-homotheticity in quality evaluations and differences in price elasticities. The evidence presented therein is in support of this assumption.

income distribution in a reduced form approach, I let φ_h be a positive function of household per capita income, as for example proposed by Hallak (2006), and empirically estimated by Handbury (2012).[25] Differences in household income, in turn, enter the model through differences in household endowments of effective labor units.

The first implication of (1) is that a household's expenditure shares within product groups increase in φ_h for products with above average quality, and decrease in φ_h for below average quality products.[26] As a consequence, the weighted average quality of the household's consumption basket increases in its quality valuation φ_h.

The second implication of (1) concerns the relationship between changes in product quality and observed market demand. For clarity, let us first consider the conventional approach to the stylized facts in Figures 1 and 2 by assuming that preferences in (1) pertain to one single representative household. In this case, the elasticity of market demand with respect to quality becomes a function of this representative agent's φ_h^* parameter, which can be set to unity without loss of generality.[27] Holding prices constant, total sales increase or decrease subject to the elasticity $(\sigma\text{-}1)\varphi_h^*$ when product quality changes. This insight has been used extensively in empirical industrial organization and international trade to infer differences in product quality by observing both prices per physical unit and product market shares.

Following this approach, the positive relationship between unit values and market shares in Figure 2 would imply that prices per physical unit increase close to log linearly with product quality.[28] It would also imply that producers offering a better quality adjusted price sort into higher quality products.[29] Taken together, the microdata in Figures 1-3 would thus suggest that higher quality products are **1)** produced using higher shares of imported inputs, **2)** produced by more productive plants, and **3)** consumed by richer households.

The particularly convenient assumption of the representative agent framework is that by setting $\varphi_h^*{=}1$, one can avoid any distinction between *physical* product quality (q_i) (e.g. degree of shininess or number of screws securing a handle) and *perceived* quality $\left(q_i^{\varphi_h^*}\right)$ without any loss of generality. In the presence of non-homothetic preferences in (1), however, one cannot avoid this distinction because the observed market valuation of a product's quality characteristics reflects heterogeneous household evaluations of the identical features. To see this, I derive this reference evaluation from the elasticity of the horizontal summation of household demands with respect to physical product quality. Let y_{hi} and y_i indicate household h's expenditure on product item i within a product group and total market sales

[25] A straight forward way to microfound this structure would be to assume complementarity between the consumption of a (normal) outside good and higher quality within differentiated product groups.

[26] See Appendix for details.

[27] With a representative agent, setting $\varphi_h^*{=}1$ is a monotonic transformation of utility in (1).

[28] To see this, we can write the elasticity of sales $(\sum_H y_{hi})$ with respect to price as: $\frac{\partial \ln(\sum_H y_{hi})}{\partial \ln \mathrm{p}_i}{=}(\sigma-1)\left(\frac{\partial \ln \mathrm{q}_i}{\partial \ln \mathrm{p}_i}-1\right)$.

[29] To see this, we note that $\frac{\partial \ln(\sum_H y_{hi})}{\partial \ln \mathrm{p}_i}{=}(\sigma-1)\left(\frac{\partial \ln \mathrm{q}_i}{\partial \ln \mathrm{p}_i}-1\right)>0$ implies that $\frac{\partial \ln \mathrm{q}_i}{\partial \ln \mathrm{p}_i}>1$, and so $\frac{\partial \ln(\mathrm{q}_i/\mathrm{p}_i)}{\partial \ln \mathrm{p}_i}>0$.

of variety i respectively. We get:

$$\frac{\partial \ln(\sum_H y_{hi})}{\partial \ln q_i} = (\sigma - 1)\left(\sum_H \frac{y_{hi}}{y_i}\varphi_h\right) = (\sigma - 1)\,\varphi_i^* \tag{2}$$

Expression (2) reveals the concept of product quality that we infer from prices and market shares when acknowledging the presence of non-homotheticity. In particular, ${q_i}^* = {q_i}^{\varphi_i^*}$ is an expenditure weighted average valuation of quality. Quality is thus measured in units of market valuation, rather than in physical units of product attributes, because the reference evaluation in (2), φ_i^*, is not constant across the spectrum of physical product qualities. In fact, (1) implies that expenditure shares $\left(\frac{y_{hi}}{y_i}\right)$ of households with lower (higher) quality evaluations are decreasing (increasing) in an item's physical quality, so that the distribution of product quality estimated from unit values and market shares will have a larger variance and be skewed to the right compared to the underlying distribution of differences in physical quality characteristics.[30]

In the present setting, expression (2) has two main implications. First, it matters for how we model quality choice across plants on the production side in order to derive the observed unit value relationships in the microdata. That is, while all three conclusions about the stylized facts from the representative agent approach above still hold, their log linear functional forms do not, unless we relabel product quality from a physical product concept to one in terms of perceived market valuation. To see this difference more clearly, note that under non-homothetic tastes in (1) and (2), holding prices constant, a percentage increase in physical quality leads to a larger percentage increase in sales for an initially higher quality product compared to a lower quality item. In contrast, this elasticity is constant and equal to $(\sigma - 1)$ with respect to changes in perceived quality $\left({q_i}^* = {q_i}^{\varphi_i^*}\right)$.

Second, expression (2) can in principle be a source for great inconvenience when modeling quality choice on the producer side. In particular, while the finding that market shares reflect expenditure weighted household evaluations of product quality is intuitive and not particular to preferences in (1), the CES structure does impose the convenient assumption that rich and poor households respond to price changes in the same way. If this were not the case, then expenditure shares $\left(\frac{y_{hi}}{y_i}\right)$, and thus the market evaluation of a given quality characteristic, would vary across a firm's pricing decisions.[31]

[30] To see this, we can write: $\varphi_i^* = \sum_H \frac{y_{hi}}{y_i}\varphi_h = N_H Cov\left(\frac{y_{hi}}{y_i}; \varphi_h\right)_i + \bar{\varphi_h}$. The intuition for this result is that the initial total sales of higher (lower) quality products derive to a higher (lower) extent from consumers who attach greater value to a given percentage change in quality, so that sales respond more (less) to a given change in quality.

[31] In other words, the CES structure assures that the change in sales due to a change in product quality is independent on a firm's pricing decision.

4.2 Technology

On the production side, the final goods sector consists of a continuum of monopolistically competitive plants that produce horizontally and vertically differentiated varieties within any given product group. Production of the final good can be separated into a production function of physical units, and a production function of quality that depends on plant characteristics and input quality. The production function of final goods is given by:

$$F_{Fi}{=}\lambda_i m_i \tag{3}$$

λ_i is a plant specific productivity parameter that in this context of quality differentiation I will refer to as technical efficiency. It defines the efficiency at which a plant converts a given number of intermediate inputs, m_i, into units of final products. Following Kugler and Verhoogen (2011), the production of final product quality is subject to complementarity between input quality, z_i and technical plant efficiency:[32]

$$q_i^*{=}\left(\alpha\lambda_i^{\psi\vartheta}+(1-\alpha)z_i^{\gamma\vartheta}\right)^{1/\vartheta} \quad \vartheta{<}0 \quad \psi{>}0 \quad \gamma{>}0 \quad 0{<}\alpha{<}1 \tag{4}$$

Here, ϑ determines the degree of complementarity between technical plant efficiency and input quality, and the assumption $\vartheta{<}0$ imposes log supermodularity of quality in plant efficiency and intermediate quality.[33] The intermediate input is produced subject to perfect competition using labor hours denoted by l. The input production function is:

$$F_{Mj}{=}\frac{l_j}{z_j^{\delta(\tau)}} \tag{5}$$

Intermediate unit costs thus increase in input quality. Substituting intermediate unit costs into final product unit costs, we get $c_i{=}\lambda_i^{-1}p_{im}{=}\lambda_i^{-1}\mathrm{w}z_i^{\delta(\tau)}$, where w is the wage rate. Kugler and Verhoogen (2011) interpret the input characteristic z_i as intermediate quality which is complementary to plant efficiency in producing output quality. The simplest possible way to introduce foreign inputs into this setting is by letting input quality be increasing in shares of imported inputs. (5) captures this assumption in a simple reduced form approach by letting the elasticity of unit input costs with respect to input quality (δ) be an increasing function of foreign input costs (τ).[34]

[32]This is consistent with empirical findings in Kugler and Verhoogen (2009) and Manova and Zhang (2012).

[33]See Costinot (2009). Intuitively, $\vartheta{<}0$ assures that a marginal increase in input quality leads to a greater increase in final product quality for a higher λ_i plant. Note that in (4) product quality enters the production side not as physical concept (q_i), but in terms of market valuation (q_i^*). This convenient functional form assumption yields equilibrium relationships between unit values, market shares, and input usage that are consistent with the documented stylized facts from the Mexican plant microdata.

[34]Since I deliberately abstract from relative factor income effects of import access, this is convenient but without loss of generality. To see this, consider a unit input cost function c=c$(w,\tilde{w})$ of any linearly homogeneous input production function where w and $\widetilde{w}$ are prices of a domestic and a foreign factor of production. Using Shephard's lemma, we get the unit factor requirements: a=$\frac{\partial c}{\partial w}$ and b=$\frac{\partial c}{\partial \tilde{w}}$. Differentiating

As in Melitz (2003), to enter the final-good sector, plants pay an investment cost f_e measured in domestic labor units in order to receive a technical efficiency draw λ. The distribution of this parameter is assumed to be Pareto with a c.d.f. $G(\lambda)=1-\left(\frac{\lambda_m}{\lambda}\right)^{\xi}$, where $\lambda_m<\lambda$, and ξ is the shape parameter. There is a fixed cost of production, f in each period, and plants exit with exogenous probability χ each period. Given zero cost of horizontal differentiation, each plant choosing the same product quality produces a distinct variety so that λ can be used to index both plants and varieties.

4.3 Predictions and Estimation Equations

In equilibrium, plants simultaneously choose output quality and prices to maximize profits, while households maximize utility in (1). In this setting, the model guides the empirical estimation in three main respects that I summarize here in the order of the three subsequent empirical sections. First, the model yields a series of predictions on NAFTA's effect on Mexican consumer prices that I am able to test empirically by drawing on the barcode level store price panels (Predictions 1.1-1.3). Second, it yields an estimation equation of the cost of living implications of these price effects in terms of observable moments in the household consumption microdata (Prediction 2). Third, it guides the estimation of differences in quality and quality adjusted prices across the Mexican income distribution in terms of observable moments in the household and plant microdata (Prediction 3).

In equilibrium, the elasticity of product unit values with respect to perceived quality is given by: $\frac{\partial lnp_{ki}}{\partial lnq^*_{ki}} = \eta_k=\frac{\delta(\tau)}{\gamma}-\frac{1}{\psi_k}$. Here, the parameter ψ_k from the quality production function (4) represents the equilibrium elasticity of perceived quality with respect to plant efficiency $\left(\frac{\partial lnq^*_{ki}}{\partial ln\lambda_{ki}}\right)$. Following Kugler and Verhoogen (2011) and earlier work by Sutton (1998), this parameter can be thought of as a product group specific scope for quality differentiation. A given distribution of ex ante plant heterogeneity leads to a wider range of product quality if the scope parameter ψ_k is greater. Intuitively, the first term in η_k represents the unit cost-quality elasticity in absence of endogenous plant sorting, while the second term captures the equilibrium link between plant efficiency and quality.

In this setting, we can derive three testable predictions about the effect of cheaper access to imported inputs on Mexican consumer prices. The first prediction concerns the average effect of input tariff cuts on the relative prices of initially more or less expensive products within product groups.

Prediction 1.1: *Input tariff cuts decrease the relative price of higher quality products.*

$$\frac{\partial^2 lnp_{ki}}{\partial lnq^*_{ki}\partial\tau_k}=\frac{\partial\eta_k}{\partial\tau_k}>0 \tag{6}$$

$c=c(w,\tilde{w})$, we get: dc=a dw+b $d\tilde{w}$. Rearranging, we get: $\frac{dc}{c}=\frac{aw}{c}\frac{dw}{w}+\frac{b\tilde{w}}{c}\frac{d\tilde{w}}{\tilde{w}}$ (e.g. Bhagwati *et al.*, 2009, pp. 143-144). The elasticity of input unit costs with respect to input quality can then be written as: $\left(\frac{dc}{c}/\frac{dz_j}{z_j}\right)=\frac{d\left(\frac{b\tilde{w}}{c}\right)}{dz_j/z_j}\frac{d\tilde{w}}{\tilde{w}}$, which is increasing in foreign factor costs as long as the foreign input cost share $\left(\frac{b\tilde{w}}{c}\right)$ is increasing in input quality z_j.

Because the production of higher quality is more intensive in imported inputs, cheaper access to foreign inputs reduces the equilibrium elasticity of unit values with respect to product quality (η_k). Coupled with the observation in Figures 1 and 2 that on average $0 < \eta_k < 1$ across consumer product groups, this first prediction implies that input tariff cuts should lead to a relative reduction in the price of initially more expensive products within product groups.

The second prediction concerns the heterogeneity of this observable effect as a function of sectoral scopes for product differentiation. In particular, I can exploit a common feature of existing models of quality choice across plants in production - that unit values embody technological heterogeneity differently across product groups with different scopes for differentiation - for empirical estimation.

Prediction 1.2: *The observed average effect of input tariff cuts on relative prices is driven by differentiated product groups.*

$$\frac{\partial^2 lnp_{ki}}{\partial lnp_{ki}^{t0}\partial\tau_k} = \begin{cases} 0 & \text{for } \eta_k = 0 \text{ (”Undifferentiated”)} \\ \frac{1}{\eta_k}\frac{\partial\eta_k}{\partial\tau_k} > 0 & \text{for } \eta_k > 0 \text{ (”Differentiated”)} \end{cases} \tag{7}$$

To empirically estimate differences in sectoral scopes for quality differentiation, the model yields a convenient estimation equation in terms of observable moments in the plant micro-data:

$$\frac{\partial lnp_{ki}}{\partial ln\left(\frac{q_{ki}^*}{p_{ki}}\right)} = \frac{\eta_k}{1-\eta_k} = (\sigma - 1)\frac{\partial lnp_{ki}}{\partial lns_{ki}} \tag{8}$$

where the final term is a product group specific elasticity between unit values and market shares. In particular, notice that product group differences in this observable elasticity provide a sufficient statistic to separate product groups into sectors that show a statistically significant relationship between initial unit values and quality ($\eta_k > 0$) which I refer to as differentiated, and sectors without a significant relationship between unit values and quality (η_k close to 0) which I refer to as undifferentiated.[35]

The third prediction concerns the effect of input tariff cuts on the distribution of product market shares across the quality distribution.

Prediction 1.3: *Input tariff cuts lead to a reallocation of market shares towards higher quality products.*

$$\frac{\partial^2 lns_{ki}}{\partial lnq_{ki}^*\partial\tau_k} = \frac{\partial\left((\sigma-1)\left(1-\eta_k\right)\right)}{\partial\tau_k} < 0 \tag{9}$$

[35] Intuitively, as the link between plant efficiency and product quality (captured by scope parameter ψ_k) decreases, η_k decreases. In differentiated sectors unit values increase with product quality, but less than one for one due to higher plant productivities associated with quality. In non-differentiated sectors more productivity improvements are necessary to produce units of quality so that prices are insignificantly related to higher plant capability and product quality.

It is apparent from (4) and (5) that cheaper access to higher quality inputs benefits producers of higher final good quality relatively more. By reducing η_k this increases the elasticity of plant revenue with respect to quality in the final good sector $\left(\frac{\partial lns_{ki}}{\partial lnq^*_{ki}}=(\sigma-1)(1-\eta_k)\right)$. The testable implication is that product groups with higher tariff cuts on their imported inputs should experience a shift of market shares towards the higher end of the quality distribution. In the next section, I draw on detailed monthly listings of product exit and entry in the store price microdata to test this prediction empirically across the unit value distribution.

Because household quality evaluations (φ_h) are increasing in household income, the consumption baskets of richer households embody higher weighted average product quality than those of poorer households. Cheaper access to foreign inputs thus gives rise to differences in cost of living inflation across the income distribution. Under two assumptions that the model makes explicit, this effect can be expressed as a function of initial expenditure share differences and price growth. Following Konus (1939), a household's cost of living index is defined as the ratio of expenditures necessary to reach a reference utility level u^* subject to store price vectors at two periods p^{t_0} and p^{t_1}: $\frac{\mathrm{e}(u^*,p^{t1})}{\mathrm{e}(u^*,p^{t0})}$. Denoting a poor and a rich household by subscripts P and R respectively and taking log differences, we get:

Prediction 2: *The relative price effect of input tariff cuts increases real income inequality through differences in cost of living inflation.*

$$\ln\left(\frac{e\left(u^*,p^{t1}\right)}{e\left(u^*,p^{t0}\right)}\right)_P\text{-}\ln\left(\frac{e\left(u^*,p^{t1}\right)}{e\left(u^*,p^{t0}\right)}\right)_R=\sum\nolimits_I\left(s^{t0}_{kiP}-s^{t0}_{kiR}\right)\ln\left(\frac{p^{t1}_{ki}}{p^{t0}_{ki}}\right)>0 \qquad (10)$$

Within the structure of the model, (10) presents the difference in the exact ideal price index due a change in the price of quality. Two assumptions underlie this convenient result. First, the model abstracts from general equilibrium effects on relative incomes of rich and poor households, which would affect the ideal expenditure weights due to non-homotheticity in (1) (Diewert, 1979). Second, the CES functional form in (1) abstracts from differences in the elasticity of substitution across households. That is, while households are allowed to substitute away from higher price increases (the source of the traditional CPI substitution bias), they do so at the same rate so that by taking the difference in household cost of living inflation, second period expenditure shares drop out of the expression. Notice that if either of these assumptions were violated, (10) would remain a first order approximation of the difference in cost of living inflation.[36]

Finally, the equilibrium elasticity of product quality with respect to quality adjusted productivity is $\frac{\partial lnq_{ki}}{\partial ln(q^*_{ki}/\mathrm{c}_{ki})}=\frac{1}{1\text{-}\eta_k}>0$, so that differences in household quality choices are predicted to translate into differences in weighted average plant productivities and, thus, quality adjusted prices across the income distribution. Denoting log differences in weighted average product quality and (inverse) quality adjusted prices by (lnQ_P-lnQ_R) and $\left(ln\left(\frac{Q}{P}\right)_P-ln\left(\frac{Q}{P}\right)_R\right)$ respectively, we get:

[36] The Appendix provides a more detailed derivation of (10).

Prediction 3: *Differences in household quality choices translate into differences in quality adjusted prices.*

$$\ln\left(\frac{e\left(u^{*},p^{t1}\right)}{e\left(u^{*},p^{t0}\right)}\right)_{P}\text{-}\ln\left(\frac{e\left(u^{*},p^{t1}\right)}{e\left(u^{*},p^{t0}\right)}\right)_{R}=\ln\left(\frac{Q}{P}\right)_{P}\text{-}\ln\left(\frac{Q}{P}\right)_{R}=\sum\nolimits_{K}\left(1-\eta_{k}\right)\left(lnQ_{Pk}-lnQ_{Rk}\right)<0 \tag{11}$$

Expression (11) represents the following comparative static result. For any given observed differences in household consumption expenditure between rich and poor people, (11) presents the difference in household cost of living inflation that would have to occur if the conventional assumption was true that quality differences do not embody differences in plant productivity. In particular, under standard assumptions η_k=1 would imply that unit values increase proportionally in product quality, so that higher quality is unrelated to quality adjusted prices. (11) is thus equivalent to a quality adjusted price index in a cross-section of households that face the identical store prices in the identical location.

To summarize, the theoretical framework guides the empirical estimation in three main respects. It yields testable predictions on the effect of NAFTA's tariff cuts on Mexican consumer prices. It provides an estimation equation to evaluate the cost of living implications of these relative price changes. And it guides the estimation of differences in quality and quality adjusted prices in household consumption. In the remainder of the paper, I draw on the Mexican microdata to empirically estimate these in the stated order.

5 Testing NAFTA's Effect on Relative Consumer Prices

This section draws on the central bank store price surveys to empirically test the predictions on NAFTA's effects on Mexican consumer prices. First, I test for the average effect of input tariff cuts on the relative price of initially more expensive relative to less expensive products within product groups. Second, I test whether this observed average effect is driven by product groups that are estimated to be differentiated in terms of observable moments in the plant microdata. Finally, I draw on detailed monthly listings of product entry and exit to test for the reallocation effects of input tariff cuts on product market shares within product groups.

5.1 NAFTA's Average Effect on Relative Prices

5.1.1 Empirical Strategy

To test for the effect of NAFTA's input tariff cuts on the relative price of initially more expensive relative to less expensive barcode items within product groups in (6), I run the following baseline regression equation:

$$\mathrm{dln}p_{ick}^{94-00} = \alpha_{ck} + \beta_1 \ \mathrm{ln}p_{ick}^{94} + \beta_2 \ \mathrm{ln}p_{ick}^{94} * \mathrm{d}\tau_k^{93-00} + \varepsilon_{ick} \tag{12}$$

$\mathrm{dln}p_{ick}^{94-00}$ is the log price change of a unique barcode-store combination i in product group k and city c from the third quarter in 1994 to the third quarter in 2000, and $\mathrm{d}\tau_k^{93-00}$ is the weighted average intermediate import tariff change under NAFTA in percentage points across four digit input industries of product group k. α_{ck} indicates city-by-product group fixed effects. Price growth is thus regressed on initial log price levels and their interaction with a product group's intermediate input tariff change within city-by-product group cells. The coefficient β_2 captures how the relative price growth of initially higher or lower unit values within city-by-product group cells differs across product groups with higher or lower intermediate import tariff cuts. To address the concern of correlated error terms (ε_{ick}) across barcode items in the same product group, standard errors are clustered at the level of 144 product groups.

5.1.2 Baseline Results

Table 1 presents the baseline estimation results. Column 1 reports results before including the interaction of initial log prices with intermediate tariff changes. Store prices within city-by-product group cells appear to have significantly converged in Mexico over the period 1994-2000. This result could be driven by a number of economic stories including trade and the relative price of quality, as well as, for example, the very significant economic crisis that unfolded over this period. Alternatively, $\beta_1<0$ might just be a consequence of measurement error or temporary store price hikes (drops) in the initial period, so that initially high (low) prices within a city-product cell would have a mechanical tendency towards lower (higher) price growth.[37]

Column 2 of Table 1 then introduces the product group's tariff interaction of interest. Product groups with higher tariff cuts on their intermediates are characterized by lower relative price growth of initially higher unit values in a statistically significant way. In Column 3 the point estimate of β_2 is unaffected by the inclusion of contemporaneous import and export tariff cuts on final consumer products. The precision of the β_2 estimate slightly increases, while no statistically significant effect of tariff changes on final goods is found. This result is consistent with the minor share of Mexican consumption expenditure on US imported final consumer goods documented in Figure A.1.1, and the fact that export tariffs to the US had already been at low levels before NAFTA.[38]

[37]This would be analogous to a case of Galton's fallacy as discussed by Quah (1993) in the context of the empirical literature on growth and convergence.

[38]See also Verhoogen (2008) for a discussion of the relatively minor importance of export tariff changes compared to real exchange rate movements due to the Peso crisis in 1995.

5.1.3 An Instrumental Variable Strategy for Input Tariff Changes and Additional Robustness Results

The identification of β_2 is based on comparing the relative price growth of barcode items within product groups across product groups that have been exposed to different degrees of intermediate import tariff cuts. By focusing on the within product group dimension of relative price changes, the identifying assumption is that intermediate tariff cuts are not targeted at particular product lines within sectors. Given that imported intermediate input tariffs affect a wide range of producers because inputs are shared across sectors, this assumption appears plausible.

To address remaining concerns about potentially omitted factors that could affect relative price growth systematically across the initial price distribution, while also being correlated with the weighted average intermediate tariff cuts across product groups, I also propose an instrumental variable (IV) strategy. In particular, I notice that there is a subset of intermediate sectors, such as basic chemicals, that have significant shares of input use across a wide range of domestic destination sectors. Since endogeneity concerns revolve around the strategic targeting of tariff cuts at particular establishments within industries, such concerns are less likely with respect to input categories that are widely shared across the Mexican economy.

Guided by this logic, I adjust input-output requirements to sum to 100% for each destination product group over a subset of 20% of intermediate input sectors that have the highest median input requirement coefficients across all four digit destination sectors in the total (not just import flows) Mexican input output table. Table A.4 in the Appendix provides an overview of these sectors. I then construct an instrument for the overall weighted average intermediate input tariff cut using the weighted average across these commonly shared input categories for each product group. The IV estimation results reported in Column 4 of Table 1 confirm both the size and statistical significance of the OLS point estimate of β_2 which provides evidence against endogeneity concerns of the tariff treatments.

Table 1: Testing the Average Effect of Input Tariffs on Mexican Store Prices 1994-2000

Dependent Variable: Change ln(Store Price) 1994-00	(1) OLS	(2) OLS Baseline	(3) OLS Controls	(4) IV	(5) OLS No Identical Items	(6) Placebo 1 Only Identical Items	(7) Placebo 2 1989-1993	(8) Placebo 3 1994-1995
ln(Store Price 1994)	-0.183***	0.271	0.292	0.336	0.366	-1.054	-0.232	0.0417
	(0.0248)	(0.250)	(0.229)	(0.264)	(0.252)	(0.966)	(0.205)	(0.0690)
ln(Store Price 1994) * Change Intermed Imp Tariff 93-00		0.0379*	0.0442**	0.0434*	0.0456**	-0.0414	-0.00894	0.00834
		(0.0214)	(0.0199)	(0.0222)	(0.0215)	(0.0810)	(0.0163)	(0.00570)
ln(Store Price 1994) * Change Export Tariff 93-00			-0.00325					
			(0.00452)					
ln(Store Price 1994) * Change Final Imp Tariff 93-00			-0.00285					
			(0.00639)					
City Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Product Group Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
City-By-Product Group FX	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
City-By-Barcode FX	No	No	No	No	No	Yes	No	No
Obs	13,589	13,589	13,589	13,589	11,143	13,589	11,029	13,589
1^{st} Stage F-Stat				735.35				
N(Product Groups)	144	144	144	144	144	144	144	144
Within-R^2	0.093	0.099	0.099		0.101	0.211	0.070	0.042

All regressions include city-by-product group fixed effects. Intermediate tariff changes are weighted averages in percentage points, where weights are total (direct and indirect) requirement coefficients across four digit sectors in the import matrix of the Mexican IO table. Final product import and export tariffs are at the HS six digit level. Instrumental variable results are 2nd stage estimates after instrumenting for input tariff changes with the weighted average tariff changes of the 20% of input sectors with the highest median input requirement across destination sectors in the Mexican IO table. In Columns 5 and 6 "Identical Items" refers to multiple counts of the identical product within a city. Standard errors are clustered at the level of 144 final product groups. ***1%, **5%, and *10% significance levels.

As noted in the data section, one potentially important limitation of the Mexican store price microdata is that store identifiers cannot be recovered from the unique item identifiers. This gives rise to the concern that the estimated pattern of price dynamics within city-product group cells could be driven by relative price changes across stores rather than across vertically differentiated products. Fortunately, the detailed product descriptions can be used to estimate a robustness check on this question. Column 5 reports estimation results after excluding multiple counts of the same barcode items within a city.[39] In turn, Column 6 reports results after including city-by-barcode fixed effects so that the estimation is restricted to variation across multiple counts of identical barcode items within a city. The fact that the exclusion of multiple product counts slightly increases both precision and size of the point estimate of β_2, and that no effect is found on identical items in Column 6, provide assurance that the observed tariff effects are driven by relative price changes across vertically differentiated products, rather than omitted changes in store markups.

In a final set of robustness checks, I estimate two additional placebo falsification tests. Columns 7 and 8 report regressions of store price changes during 1989-1993 and during 1994-1995 on NAFTA tariff changes 1993-2000. The first falsification test reported in Column 7 is estimated off price changes during the preceding four year period during which Mexican import tariffs remained practically unchanged (e.g. Kate, 1992). The second falsification test reported in Column 8 is estimated off price changes between the 3rd quarter of 1994 and the first quarter of 1995 which captures the spike of inflation that occurred in Mexico in the immediate aftermath of the Peso crisis in December 1994 (see Appendix Figure A.3).

These specifications address two particular concerns. The first is that NAFTA tariff changes might be associated to particular product groups that, in general, are characterized by different price distributional changes across stores and/or barcode items. The fact that the point estimate of the tariff interaction in Column 7 is close to zero and insignificant in the preceding period of price changes provides evidence against this concern. Second, tariff changes might have been correlated with product groups whose price distributions were differently affected by the Peso crisis. The fact that the β_2 point estimate in Column 8 is close to zero and insignificant provides evidence against this concern.

5.2 Testing the Heterogeneity of Tariff Effects

To test the model's prediction on the heterogeneity of the observed average effect of input tariff changes in (7), I extend the baseline specification (12) in the following way:

$$\begin{aligned} \mathrm{dln}p_{ick}^{94-00} &= \alpha_{ck} + \beta_1 \ln p_{ick}^{94} + \beta_2 \ln p_{ick}^{94} * \mathrm{d}\tau_k^{93-00} \\ &+ \beta_3 \ln p_{ick}^{94} * \mathrm{Tech}_k + \beta_4 \ln p_{ick}^{94} * \mathrm{d}\tau_k^{93-00} * \mathrm{Tech}_k + \varepsilon_{ick} \end{aligned} \tag{13}$$

[39]For items with multiple counts, one of the identical product items was randomly selected to remain in the regression in order not to waste information.

β_1 captures the average relative price growth between product items with initially higher versus lower unit values within city-by-product group cells. β_2 captures the effect of intermediate input tariff cuts on this relative price change among the non-differentiated product groups in the reference category ($Tech_k = 0$). β_3 captures how on average price changes of initially higher and lower unit values differ between non-differentiated and differentiated ($Tech_k = 1$) sectors. The coefficient of interest, β_4, captures how the relative price effect of intermediate tariff cuts differs in differentiated product groups relative to the non-differentiated reference category. The prediction of the model outlined above is that β_4>0.

5.2.1 Estimating Technology Parameters from Plant Microdata

To empirically estimate differences in sectoral scopes for quality differentiation captured by the Tech_k indicator in (8), I follow the model's estimation equation in (13). In particular, I estimate the following specification separately across six digit industries in the 1994 monthly plant surveys in order to parameterize the final term $\left(\frac{\partial lnp_{ki}}{\partial lns_{ki}}\right)$ in (8):

$$\text{ln}p_{igmkrt}\text{=}\ \alpha_{rt} + \alpha_{mt} + \beta\ \text{ln}s_{igmkrt} + \varepsilon_{igmkrt} \tag{14}$$

Subscript i indexes a plant-product line combination, g indexes a plant, m indexes several thousands of manufacturing product groups, k indexes six digit production sectors, r indexes 32 Mexican states, and t indexes 12 months in 1994. lns_{igmkrt} are log monthly sales of different product lines within a plant. Log monthly output unit values are thus regressed on product group-by-month fixed effects, state-by-month fixed effects, and the product line's log sales.[40]

Following from (8), the β coefficient yields an estimate of $\left(\frac{1}{\sigma-1}\frac{\eta_k}{1-\eta_k}\right)$, either pooled across all product groups, or estimated individually for each six digit manufacturing sector. In order to empirically distinguish differentiated sectors where initial unit values in 1994 are positively related to product quality and quality adjusted productivity, and non-differentiated sectors, where no such relationship is observed, I define a binary identifier variable Tech_k which takes the value 1 if β is statistically significantly positive at the 10% level.[41][42]

To address the concern of correlated error terms within the same product category,

[40]By construction of the plant surveys, sectoral fixed effects would be redundant in (14) because a plant cannot be assigned to product codes in different six digit sectors. In such cases, a novel product group is added in the principal six digit sector. See also Iacovone (2008) for a discussion of the EIM plant data.

[41]Kugler and Verhoogen (2011) estimate the scope for differentiation pooling across product groups within four digit industries, rather than at six digit level. I also replicate all subsequent results using the alternative four digit level of aggregation to assign product differentiation. Results are unaffected in size or statistical significance.

[42]The model relates differences in the magnitude of the unit value-sales elasticity to the scope for differentiation captured by the parameter ψ_k. As a robustness check I follow Kugler and Verhoogen (2011) and verify that the technology estimates are related to existing "off-the-shelf" measures of vertical differentiation in statistically significant way. Appendix Figure A.4 provides graphs and Table A.6 reports regression results in which either the log sales regressor or intermediate input shares in Figure 1 are interacted with Sutton's (1998) measure of differentiation in terms of advertising and R&D intensities across sectors.

standard errors are clustered at the level of m product groups.[43] Finally, following Deaton (1988), both unit values on the left hand side and sales on the right embody measurement error in prices. To address the concern of non-traditional measurement error, I follow Kugler and Verhoogen (2011) and instrument for a product line's log monthly sales by the log of its establishment's employment in 1994.

Table A.5 presents the pooled unit value-sales elasticities for all reporting plants, as well as for final goods sectors, and for differentiated final goods industries only. Unit value-sales elasticities are estimated in OLS for log sales and log employment, as well as by IV when instrumenting for log sales with employment in the third column. Out of the 203 reporting manufacturing six digit industries, 79 sectors can be matched to final consumption product groups present in the consumer surveys and/or store samples. Differentiated sectors refer to product groups within six digit sectors for which the unit value-sales elasticity estimate is statistically significantly greater than zero at the 10% level. This cutoff identifies about one third of processed tradable household consumption as differentiated.[44]

5.2.2 Results and Robustness

The estimated technology parameters allow me to estimate regression specification (13) in order to test the prediction on the product group heterogeneity of the tariff effect. In particular, the plant data estimates in Table A.5 suggest that product groups significantly differ in the degree to which observed unit values in 1994 are related to differences in quality and plant technologies. The results reported in Table 2 confirm the prediction. In particular, the first interaction term (β_2) in (13) becomes statistically insignificant, indicating that the previously estimated average effect of intermediate import access on within product group store prices in Table 1 is indeed driven by differentiated product groups.

The point estimate of β_4 is confirmed in size and statistical significance in Column 2 when instrumenting for intermediate input tariff cuts with the weighted average input tariff cut across the 20% of input sectors with the highest median requirement across destination sectors in the Mexican input output table. Finally, Columns 3-6 report the identical battery of robustness tests as discussed for the average tariff effect in Table 1. In particular, the point estimate of $\beta_4$slightly increases when excluding identical barcode items, and it becomes close to zero and statistically insignificant when estimated off price changes in the preceding period, or price changes in the immediate aftermath of the Peso crisis.

[43]Reported results are unaffected by clustering at the plant level instead.

[44]This proportion is close to identical when choosing to pool all product groups in the same four digit industry to estimate the technology parameters, rather than the less aggregated six digit level chosen here.

Table 2: Testing the Heterogeneity of Tariff Effects

Dependent Variable: Change ln(Store Price) 1994-00	(1) OLS 1994-2000	(2) IV 1994-2000	(3) OLS No Identical Items	(4) Placebo 1 Only Identical Items	(5) Placebo 2 1989-1993	(6) Placebo 3 1994-1995
ln(Store Price 1994)	0.135	0.0678	0.191	-1.115	-0.327	-0.0629
	(0.310)	(0.268)	(0.311)	(0.807)	(0.388)	(0.104)
ln(Store Price 1994) * Change Intermed Imp Tariff 93-00	0.0246	0.0191	0.0292	-0.0767	-0.0164	3.18e-05
	(0.0252)	(0.0211)	(0.0253)	(0.0687)	(0.0296)	(0.00817)
ln(Store Price 1994) * Differentiated Prod Dummy	0.676*	0.670*	0.696*	1.254	0.200	0.211
	(0.402)	(0.389)	(0.394)	(1.051)	(0.466)	(0.131)
ln(Store Price 1994) * Change Intermed Imp Tariff 93-00 * Differentiated Prod Dummy	0.0671**	0.0662**	0.0678**	0.112	0.0159	0.0168
	(0.0337)	(0.0325)	(0.0333)	(0.0896)	(0.0373)	(0.0106)
City Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Product Group Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
City-By-Product Group Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
City-By-Barcode FX	No	No	No	Yes	No	No
Obs	13,589	13,589	11,143	13,589	11,029	13,589
1st Stage FStat		165.94				
N(Product Groups)	144	144	144	144	144	144
Within-R^2	0.108		0.110	0.219	0.070	0.044

All regressions include city-by-product group fixed effects. Intermediate tariff changes are weighted averages in percentage points, where weights are total (direct and indirect) requirement coefficients across four digit sectors in the import matrix of the Mexican IO table. Instrumental variable results are 2nd stage estimates after instrumenting for input tariff changes with the weighted average tariff changes of the 20% of input sectors with the highest median input requirement across destination sectors in the Mexican IO table. Quality differentiation is estimated from observed unit value-sales elasticities in Mexican plant microdata across 12 months in 1994. Standard errors are clustered at the level of 144 final product groups. In Columns 3 and 4 "Identical Items" refers to multiple counts of the identical product within a city. ***1%, **5%, and *10% significance levels.

5.3 Testing Reallocation Effects on Product Entry and Exit

Finally, to test for the effect of tariff cuts on the reallocation of market shares stated in (9), I draw on detailed monthly records of product additions and replacements in the central bank store price microdata. In particular, I estimate logit regressions of the form:

$$\mathrm{Entr}y_{ikc}^{94-00}\text{or } \mathrm{Exi}t_{ikc}^{94-00} = \alpha_{ck} + \beta_1 \ln p_{ick}^{t} + \beta_2 \ln p_{ick}^{t} {*} d\tau_k^{93-00}$$

$$+ \beta_3 \ln p_{ick}^{t} {*} \mathrm{Tec}h_k + \beta_4 \ln p_{ick}^{t} {*} d\tau_k^{93-00} {*} \mathrm{Tec}h_k + \varepsilon_{ick} \tag{15}$$

$\mathrm{Entr}y_{ikc}^{94-00}$ and $\mathrm{Exi}t_{ikc}^{94-00}$ are a binary indicators of reported product additions or disappearances over the period 1994-2000 respectively. Superscript t indicates the third quarter in 1994 when the dependent variable is $\mathrm{Exi}t_{ikc}^{94-00}$, and the third quarter in 2000 when the dependent variable is $\mathrm{Entr}y_{ikc}^{94-00}$. Exit propensities are thus estimated as a function of initial prices, whereas entry propensities are estimated as a function of prices in 2000.

Table 3 reports logit estimation results. The estimation results provide empirical support of the model's predictions on market share reallocations towards the higher end of the quality spectrum. While for both entry and exit regressions in Columns 1 and 3, the average effect of tariff cuts across all product groups is not statistically significant, the tariff effect is significant and of expected opposite sign for entry and exit among differentiated product groups. Higher intermediate tariff cuts appear to have increased the propensity of exit at the lower end of the initial price distribution, whereas they increased the entry propensity at the higher end of the price distribution in 2000.

Table 3: Testing Predictions on Market Share Reallocation

Dependent Variable: Additions/Replacements 1994-00	(1) Entry	(2) Entry	(3) Exit	(4) Exit
ln(Store Price 2000)	0.0720	1.255		
	(1.039)	(1.397)		
ln(Store Price 2000) * Change Intermed Imp Tariff 93-00	0.00535	0.106		
	(0.0833)	(0.110)		
ln(Store Price 2000) * Differentiated Prod Dummy		-3.415		
		(2.129)		
ln(Store Price 2000) * Change Intermed Imp Tariff 93-00 * Differentiated Prod Dummy		-0.303*		
		(0.171)		
ln(Store Price 1994)			1.065	-6.056**
			(1.369)	(2.817)
ln(Store Price 1994) * Change Intermed Imp Tariff 93-00			0.0747	-0.491**
			(0.109)	(0.223)
ln(Store Price 1994) * Differentiated Prod Dummy				13.29***
				(3.087)
ln(Store Price 1994) * Change Intermed Imp Tariff 93-00 * Differentiated Prod Dummy				1.031***
				(0.243)
City Fixed Effects	Yes	Yes	Yes	Yes
Product Group Fixed Effects	Yes	Yes	Yes	Yes
City-By-Product Group Fixed Effects	Yes	Yes	Yes	Yes
Obs	19,277	19,277	15,591	15,591

The table presents logit regression results. All specifications include city-by-product group fixed effects. Intermediate tariff changes are weighted averages in percentage points, where weights are total (direct and indirect) requirement coefficients across four digit sectors in the import matrix of the Mexican IO table. Quality differentiation is estimated from observed unit value-sales elasticities in Mexican plant microdata across 12 months in 1994. Standard errors are clustered at the level of 144 final product groups. ***1%, **5%, and *10% significance levels.

6 NAFTA's Effect on Cost of Living Inflation across the Income Distribution

In this section, I draw on the household consumption microdata to evaluate the cost of living implications of NAFTA's observed relative price effects across the Mexican income distribution.

6.1 Empirical Strategy

The cost of living expression in (10) requires information on household expenditure share differences across product items in combination with trade induced relative price changes. The empirical strategy combines observed expenditure shares linked to purchase unit values in household consumption surveys with the causal estimate of NAFTA's relative price effect within product groups from the store price panels. Observed expenditure shares reported for each individual purchase allow the estimation of the first term on the right hand side of (10), while the linked unit purchase prices allow to estimate predicted price growth in the second term as a function of an item's position in the unit store price distribution. Intuitively, the estimation strategy combines the observed expenditure weighted household unit value differences depicted in Figure 3 with the estimated effect of tariff changes on relative store prices presented in Tables 1 and 2.

It is apparent from (10) that this estimation is subject to bias if errors in predicted price growth are correlated with expenditure share differences between the rich and the poor across product items.[45] The empirical challenge is that while both survey data on store prices as well as survey data on household purchase prices stem from the identical population of points of purchase in the third quarter of 1994, the collected microdata do not report a direct link between store price items and household expenditure weights across these datasets. In other words, the Mexican microdata on store price panels and household consumption are not available in homescanner format.[46]

The particular concern that arises is that part of the observed unit value differences between rich and poor households depicted in Figure 3 could be driven by price differences across identical items due to rich people consuming at more expensive stores. In that case, the predicted price changes derived from store price regressions are based on initial store unit value differences in 1994 that reflect quality differentiation, whereas observed price differences between rich and poor households could simply reflect differences in store markups of identical items. The resulting bias would lead to an over-estimate of NAFTA's true implication on differences in household cost of living inflation for the poor relative to the rich.[47]

Fortunately, Mexican consumer surveys contain information that can be used to estimate a robustness test on this concern. In particular, the surveys report point of purchase types (street vendors, markets, convenience and specialized stores, supermarkets and department

[45]Formally, from (10) we get:
$$\left(\sum_I (s^{t0}_{kiP} - \widehat{s^{t0}_{kiR}})\, ln\left(\frac{p^{t1}_{ki}}{p^{t0}_{ki}}\right)\right) = \sum_I \left(s^{t0}_{kiP} - s^{t0}_{kiR}\right)\left(ln\left(\frac{p^{t1}_{ki}}{p^{t0}_{ki}}\right) + \epsilon_{ki}\right) = \sum_I \left(s^{t0}_{kiP} - s^{t0}_{kiR}\right) \ln\left(\frac{p^{t1}_{ki}}{p^{t0}_{ki}}\right)$$
$$+ N_I \mathrm{Cov}\left(\left(s^{t0}_{kiP} - s^{t0}_{kiR}\right), \epsilon_{ki}\right), \text{ where } \epsilon_{ki} = \left(ln\,\widehat{\left(\frac{p^{t1}_{ki}}{p^{t0}_{ki}}\right)} - ln\left(\frac{p^{t1}_{ki}}{p^{t0}_{ki}}\right)\right).$$

[46]Recent contributions making use of US homescanner data from AC Nielsen include Broda et al. (2009), Broda and Romalis (2009), Handbury and Weinstein (2011), and Handbury (2012).

[47]Formally, this would lead to positive correlation between expenditure share differences and prediction errors of relative price growth expressed in (10): $\mathrm{Cov}\left(\left(s^{t0}_{kiP} - s^{t0}_{kiR}\right), \epsilon_{ki}\right) > 0$.

stores) alongside household expenditures and product unit values. Table A.7 reports regressions of log purchase unit values on household income per capita quintile dummies both before and after including city-by-product-by store type fixed effects. If store markups were driving unit value differences, then one would expect the inclusion of store type fixed effects to significantly reduce the estimated price gap across income quintiles. The fact that the estimated unit value differences slightly increase due to the inclusion of store type fixed effects in Column 2 provides evidence against this concern.[48] Several explanations for this finding have been advocated, including the cost of mobility to reach cheaper stores (Caplovitz, 1963), or bulk discounting (Attanasio and Frayne, 2006).

6.2 Baseline Estimation

Following these insights, Figure 4 proceeds to present the baseline estimation results of the household price index effect on total tradable consumption of a 12 percentage point US import tariff cut (average of NAFTA tariff cuts 1993-2000) across the urban Mexican income distribution in 1994.[49] The baseline results are based on observed deviations of household purchase unit values from city-by-product-by-store type means in combination with the preferred estimate of NAFTA's average relative price effect reported in Column 4 of Table 1.[50]

[48]This finding is consistent with results reported in Broda et al. (2009) using US barcode homescanner data. While that paper's main conclusion is that on average poorer US households consume at slightly lower prices compared to richer households (incomes above US$ 60,000 in 2005), Figure 2 of their paper shows that this finding is reversed over the real income range reported in Mexican consumer surveys.

[49]As reported in Table 1, total tradable (i.e. non-services) household consumption accounts for on average 54% of Mexican household consumption in 1994. The store price estimation sample covers processed tradables which account for 70% of total tradable consumption. The reported estimation results are scaled to total tradable consumption, under the (conservative) assumption that no relative price effects occur among tradable products outside the estimation sample.

[50]Since common physical units can be hard to define, consumer surveys report unit values only for food products, beverages, and tobacco products. Out of the 255 processed tradable product groups, 118 report unit values. I assign the weighted average household mean unit value deviation to household expenditures with missing unit information, where the weights reflect the share of household expenditures across product groups with reported unit values. This strategy is likely to be conservative as food and beverages product groups are estimated to be on average less differentiated in the plant microdata. In confirmation of this argument, the reported regressions in Table A.8 of the Appendix show that product groups that are estimated to have higher scopes for quality differentiation in the plant microdata have statistically significantly higher estimated unit purchase value gaps between rich and poor households in the consumption surveys.

Figure 4: Cost of Living Effect of Average NAFTA Import Tariff Cut 1994-2000

The fitted relationship corresponds to the best fitting polynomial functional form according to the Akaike Information Criterion (AIC). Standard errors are clustered at the municipality level and the shaded area indicates 95% confidence intervals. The y-axis depicts mean deviations of estimated household cost of living inflation of tradable consumption due to a 12% tariff cut on US imports. These estimates are based on the average effect of input tariff cuts in Column 4 of Table 1. The x-axis depicts mean deviations of log household per capita incomes. Estimations are based on urban Mexican households in 1994 and subject to nationally representative sample weights. The bottom and top 0.5% on the x-axis are excluded from the graph.

Table 4 presents the same estimates after collapsing the data to mean outcomes across five nominal income quintiles subject to nationally representative household survey weights. The reported result is that the average tariff cut under NAFTA has led to a 2.6 percentage point increase in tradable consumption inflation of the poorest quintile of urban Mexican households compared to the richest quintile.

6.3 Accounting for Product Group Heterogeneity

Because the estimation results in Figure 4 are based on the average store price effect of tariff cuts across product groups in Table 1, the implicit assumption is that all product groups are characterized by the same average scope for quality differentiation. Following the plant data technology estimates in Table A.5 and the heterogeneity of NAFTA's store price effects in Table 2, this assumption is clearly rejected in the data.

How does the observed product group heterogeneity affect the estimation results in Figure 4? In this subsection, I report two alternative estimations which are based on different assumptions about what the observed price differences between rich and poor households within municipality-by-product-by-store type cells measure in terms of differences in product quality choices.[51]

[51]The underlying empirical challenge is that product market shares are observable in combination with unit values (to estimate quality) only in the plant microdata, while consumption surveys do not report

The first approach is based on the assumption that the observed unit value differences between rich and poor households only reflect differences in product quality choices in product groups in which it is also true that prices are correlated strongly enough with product quality to estimate a statistically significant ($\eta_k > 0$, $Tech_k = 1$) in the plant microdata. That is, we assume that any effect on the relative price of quality is only present in sectors in which we are able to proxy for product quality with unit value differences in the plant data, and for which we find significant observable store price effects in Table 2. This approach thus applies the estimated store price effect in Column 2 of Table 2 only to those household consumption product groups that are estimated to be differentiated in the plant microdata. Since the plant data estimates identify only around 30% of consumer product groups as quality differentiated in this respect, I refer to this estimation approach as a lower bound estimate of NAFTA's true effect on differences in cost of living inflation across the Mexican income distribution.

The second approach is based on the opposite assumption that price differences between the rich and the poor within municipality-by-product-by-store type reflect differences in quality choices across all 255 sample product groups, despite the fact that the relationship between unit values and quality might not be strong enough to be captured in the plant production or store price microdata. The argument is that the plant and store price data include price variation across the full product space, whereas product purchase variation between the rich and poor households within the same store types is more informative to capture quality differences.[52] This approach thus applies the estimated store price effect among differentiated sectors in Table 2 to observed household consumption price differences across all processed tradable product groups. In the estimation results reported below, I refer to this as an upper bound estimate of NAFTA's effect on differences in household cost of living inflation.

Appendix Figure A.5 presents the estimation results under these alternative assumptions in addition to the baseline estimates that are based on the average tariff effect depicted in Figure 4, and Table 4 presents these estimations after collapsing the data to mean outcomes across five nominal income quintiles subject to nationally representative household survey weights. The lower bound estimate suggests that NAFTA caused a 1.7 percentage point higher cost of living inflation for the poorest income quintile compared to the richest over the period 1994-2000. The upper bound estimate of this effect is 3.9 percentage points. As expected, these alternative estimation approaches fall on different sides of the baseline estimate of 2.6 percentage points.

---

product barcodes alongside expenditures and unit values.

[52]Acknowledging noise in the price data, this conception would be fully consistent with the predictions of the model. The larger a given set of households is apart in terms of incomes (and thus quality valuations), the larger should be the signal to noise ratio of quality differences embodied in observed price differences.

Table 4: Cost of Living Effects across the Income Distribution

	Differences in Price Index Effects		
	Baseline	Lower Bound	Upper Bound
2nd Income Quintile	-0.00388*	-0.00291	-0.00615*
	(0.00228)	(0.00233)	(0.00354)
3rd Income Quintile	-0.00969***	-0.00789***	-0.0149***
	(0.00272)	(0.00243)	(0.00422)
4th Income Quintile	-0.0136***	-0.00873***	-0.0210***
	(0.00274)	(0.00286)	(0.00424)
5th Income Quintile	-0.0259***	-0.0170***	-0.0391***
	(0.00311)	(0.00263)	(0.00475)
Household Obs	7632	7632	7632

The table presents regression results of estimated outcomes on national urban income quintile dummies across 7632 urban Mexican households. Point estimates are based on nationally representative sample weights. Standard errors are clustered at the municipality level. ***1%, **5%, and *10% significance levels. Price index effects are based on reported household unit purchase values in combination with the estimated relative store price effect of US import tariff cuts, and based on the average NAFTA tariff reduction (12 percentage points). "Baseline" is estimated using the observed average relative price effect of input tariff cuts in Column 4 of Table 1. "Lower Bound" is estimated under the assumption that no relative price effects are present in undifferentiated sectors. "Upper Bound" is estimated under the assumption that the observed relative price effect in differentiated sectors (Column 2 of Table 2) operates in all processed tradable product groups.

In terms of the direction and magnitude of NAFTA's consequences for inequality, the estimated effect reinforces the observed increase in nominal inequality in urban Mexico over the same period, and is equivalent to approximately 25-55% of the total observed difference in nominal income growth between the richest and the poorest income quintiles over the period 1994-2000.[53]

7 Quality and Quality Adjusted Prices across the Income Distribution

In the final section, I shift attention from the consumer price effects of NAFTA in Mexico to a more general implication of quality sorting by households and plants for real income inequality. In particular, I impose an additional parameter assumption on the elasticity of substitution in estimation equation (8) in order to parameterize η_k across product groups. In the following, I discuss the empirical strategy to use these parameter estimates in order

[53]This comparison adjusts for the fact tradable consumption accounts for 54% of total consumption in 1994 (i.e. I do not assume similar effects outside the estimation sample). I estimate nominal income growth differences from mean incomes by quintile in 1994 and 2000 using the identical household sample and population weights as in the cost of living estimations reported above. See also Attanasio and Binelli (2010) for a descriptive summary of changes in Mexican income inequality during the 1990s using a number of different microdata and measures of inequality.

to evaluate differences in the consumption of product quality and quality adjusted prices (plant productivity) across the income distribution in 1994.

From (11), the estimation of household differences in weighted average quality and quality adjusted prices requires information on expenditure shares and product characteristics. To estimate (11), I draw on observed unit value distributions in the household consumption surveys in combination with product group estimates of η_k from plant microdata.

The main empirical concern in this estimation is that the technology estimates described in Section 5 are based on factory gate prices in plant microdata, whereas household consumer surveys report point of purchase store prices. In particular, if distribution costs embody not just ad valorem but also per unit cost components then the factory gate price distribution would be compressed in the store price data. Because estimated differences in product quality and plant productivity are based on unit values at the factory gate in the plant microdata, the compression of the price distribution reported in consumer surveys would thus lead to an under-estimation of household consumption differences. In this light, the results reported in this section can be regarded as conservative estimates.[54]

The results reported in Table 5 and Figure 5 are based on a parameter value of σ=2.5, while I report estimation results across a range of commonly estimated demand parameterizations of σ=2-5.[55] As in the cost of living estimations of the previous section, I report estimation results subject to alternative assumptions. As before, I will refer to lower bound estimates when assuming that observed household purchase price differences only reflect quality differences in "differentiated" sectors which are estimated to have a statistically significant relationship between price differences and product quality in the plant microdata in 1994. And, as before, I will refer to upper bound estimations when assuming that purchase price differences between rich and poor households capture product quality differences in all processed tradable product groups (70% of total tradable consumption).

Table 5 presents the results of these estimations after averaging across households in five income quintiles subject to nationally representative household weights. These results are based on observed deviations of household purchase unit values from city-by-product-by-store type means in combination with the estimated unit value-quality elasticities η_k from the plant microdata in Section 5. Among differentiated product groups, the poorest quintile is estimated to consume on average approximately 30% lower weighted average quality among vertically differentiated goods sectors. From (8) this estimate is decreasing in σ and varies between 20-50% across the parameter range σ=2-5.

Finally, Figure 5 depicts the resulting differences in weighted average quality adjusted

[54]Formally, this leads to a positive correlation between expenditure share differences (poor minus rich) and prediction errors of quality and technology differences. Added noise in product retail prices compared to factory gate prices would have the same effect to under-estimate quality and technology differences embodied in consumption. Furthermore, to the extent not accounted for by store type fixed effects, the same holds true with respect to the estimated propensity of poorer households to consume identical barcodes at slightly higher prices compared to the rich in Table A.7.

[55]This range corresponds to, for example, results reported in Broda et al. (2006).

productivity that are embodied in tradable household consumption according to the lower bound estimation. When averaging the quality adjusted productivity differences across five income quintiles reported in Table 5, the poorest quintile is estimated to source tradable consumption at roughly 4.5% lower weighted average productivity compared to the richest quintile. From expression (8) this estimate is decreasing in σ and varies between 2-7% across the parameter range σ=2-5.

Table 5: Quality and Quality Adjusted Prices across the Income Distribution

	Log Quality Differences		Log Productivity Differences	
	Lower Bound	Upper Bound	Lower Bound	Upper Bound
2^{nd} Income Quintile	0.0142	0.0389	0.00719	0.0198
	(0.0123)	(0.0240)	(0.00623)	(0.0122)
3^{rd} Income Quintile	0.0423***	0.103***	0.0214***	0.0524***
	(0.0128)	(0.0285)	(0.00649)	(0.0145)
4^{th} Income Quintile	0.0467***	0.144***	0.0237***	0.0733***
	(0.0152)	(0.0290)	(0.00769)	(0.0147)
5^{th} Income Quintile	0.0912***	0.285***	0.0463***	0.145***
	(0.0142)	(0.0339)	(0.00718)	(0.0172)
Household Obs	7632	7632	7632	7632

The table presents regression results of estimated outcomes on national urban income quintile dummies across 7632 urban Mexican households. Point estimates are based on nationally representative sample weights. Standard errors are clustered at the municipality level. ***1%, **5%, and *10% significance levels. Quality differences are estimated using reported household unit purchase values in combination with the product group's estimated unit value-quality elasticity (η_k). Productivity deviations are estimated using reported household unit purchase values in combination with the product group's estimated unit value-productivity elasticity $\left(\frac{\eta_k}{1-\eta_k}\right)$. Both parameter estimates are based on σ=2.5 in demand. "Lower Bound" is estimated under the assumption that no quality differences in consumption exist in product groups with $\eta_k = 0$ in the plant microdata. "Upper Bound" is estimated under the assumption that the same observed quality differences among sectors with $\eta_k > 0$ exist in all sampled processed consumer good sectors (but nowhere outside this sample).

Figure 5: Plant Productivity and Household Income in 1994

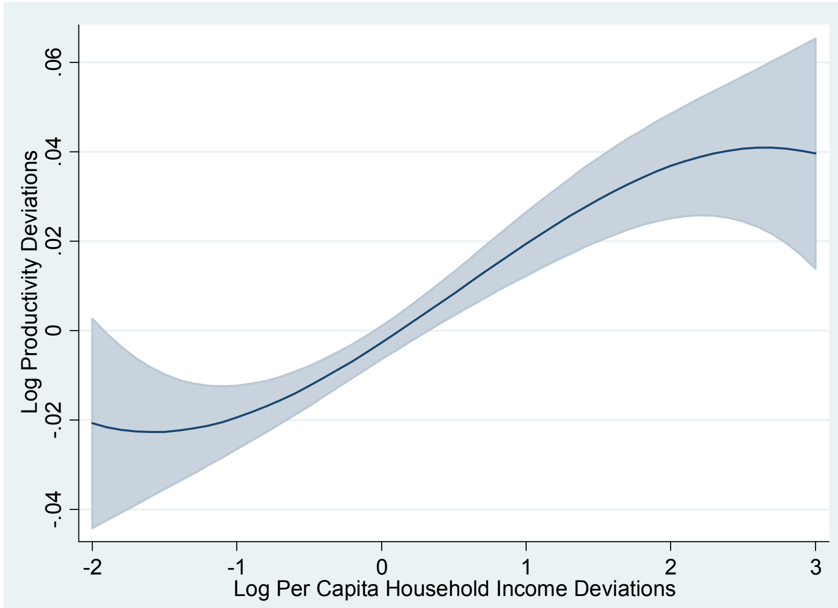

The fitted relationship corresponds to the best fitting polynomial functional form according to the Akaike Information Criterion (AIC). Standard errors are clustered at the municipality level and the shaded area indicates 95% confidence intervals. The y-axis depicts mean deviations of estimated weighted average plant productivities embodied in tradable consumption across households. Weights are reported household expenditure shares. The depicted estimates are based on a parameter value of σ=2.5 in demand, and according to the "Lower Bound" estimate in the text. The x-axis depicts mean deviations of log household per capita incomes. Estimations are based on urban Mexican households in 1994 and subject to nationally representative sample weights. The bottom and top 0.5% on the x-axis are excluded from the graph.

8 Conclusion

The question of how globalization affects real income inequality in developing countries has been a prominent policy subject in the study of international trade. This paper considers household cost of living in the denominator of real income as a channel through which trade liberalization can affect inequality in a developing country. Drawing on a new collection of microdata covering Mexican households, plants, and stores, I empirically estimate this channel in the context of Mexico's trade liberalization under NAFTA.

The paper presents evidence in favor of the hypothesis that access to intermediates from developed countries reduces the relative price of higher quality products in a developing country. In turn, because quality choices differ across the income distribution, this relative price effect is found to have significantly increased real income inequality in urban Mexico due to NAFTA over the period 1994-2000. In terms of the direction and magnitude of this effect, NAFTA's estimated cost of living implications reinforce the observed increase in nominal inequality in urban Mexico over the same period, and are equivalent to at least 25% of the total observed difference in nominal income growth between the richest and the poorest income quintiles.

The paper also highlights that quality choice in a setting with heterogeneous households and plants has a more general implication for real income inequality. In particular, the observed pattern of quality sorting is such that higher quality products embody cheaper quality adjusted prices, because more productive plants sort into higher quality product lines. The empirical analysis confirms that differences in household quality choices translate into economically significant disparities in weighted average plant productivities and, thus, quality adjusted prices of low income households relative to the rich. This finding suggests that the same observed distribution of nominal incomes leads to more pronounced differences in real incomes in a world with quality sorting in consumption and production compared to a world under conventional assumptions.

For policy analysis, the presented findings serve to highlight the importance of potential price index effects in addition to the conventional focus on nominal incomes when analyzing the general equilibrium consequences of policy or market shocks for the distribution of real incomes. In this respect, the paper points to a number of interesting unanswered research questions concerning, for example, the price index effects of globalization in developed as opposed to developing countries, and the cost of living implications of other policies, such as transport infrastructure or retail sector entry (de-)regulation, in both developing and developed economies.

References

Aghion, P., & Williamson, J. (1998). *Growth, inequality, and globalization: Theory, history, and policy.* Cambridge Univ Pr.

Amiti, M., & Konings, J. (2007). Trade liberalization, intermediate inputs, and productivity: Evidence from Indonesia. *The American Economic Review*, 1611–1638.

Atkin, D. (2010). Trade, tastes and nutrition in India. *Yale University Economic Growth Center Discussion Paper No. 986, Yale Economics Department Working Paper No. 80*.

Attanasio, O., & Binelli, C. (2010). Mexico in the 1990s: The main cross-sectional facts. *Review of Economic Dynamics*, *13*(1), 238–264.

Attanasio, O., & Frayne, C. (2006). Do the poor pay more? *University College London mimeo*.

Berry, S. (1994). Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics*, *25*(2), 242–262.

Berry, S., Levinsohn, J., & Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica*, 841–890.

Bhagwati, J., Panagariya, A., & Srinivasan, T. (1998). *Lectures on international trade.* MIT press.

Bowen, H., Leamer, E., & Sveikauskas, L. (1987). Multicountry, multifactor tests of the factor abundance theory. *The American Economic Review*, 791–809.

Broda, C., Greenfield, J., & Weinstein, D. (2006). *From groundnuts to globalization: A structural estimate of trade and growth* (Tech. Rep.). National Bureau of Economic Research.

Broda, C., Leibtag, E., & Weinstein, D. (2009). The role of prices in measuring the poor's living standards. *The Journal of Economic Perspectives*, *23*(2), 77–97.

Broda, C., & Romalis, J. (2008). Inequality and prices: Does China benefit the poor in America? *University of Chicago mimeo*.

Broda, C., & Romalis, J. (2009). The welfare implications of rising price dispersion. *University of Chicago mimeo*.

Caplovitz, D., & of Applied Social Research, C. U. B. (1963). *The poor pay more: Consumer practices of low-income families.* Free Press of Glencoe.

Choi, Y., Hummels, D., & Xiang, C. (2009). Explaining import quality: The role of the income distribution. *Journal of International Economics*, *78*(2), 293–303.

Cortes, P. (2008). The effect of low-skilled immigration on us prices: Evidence from CPI data. *Journal of Political Economy*, *116*(3), 381–422.

Costinot, A. (2009). An elementary theory of comparative advantage. *Econometrica*, *77*(4), 1165–1192.

Deaton, A. (1988). Quality, quantity, and spatial variation of price. *The American Economic Review*, 418–430.

Deaton, A. (1989). Rice prices and income distribution in Thailand: A non-parametric analysis. *The Economic Journal*, *99*(395), 1–37.

Deaton, A. (1997). *The analysis of household surveys: A microeconometric approach to development policy.* Johns Hopkins University Press.

Deaton, A. (2003). Prices and poverty in India, 1987-2000. *Economic and Political Weekly*, 362–368.

Deaton, A., & Dupriez, O. (2011). Purchasing power parity exchange rates for the global poor. *American Economic Journal: Applied Economics*, *3*(2), 137–166.

Diewert, W. (1979). *The economic theory of index numbers: A survey.* Department of Economics, University of British Columbia.

Fajgelbaum, P., Grossman, G., & Helpman, E. (2011). Income distribution, product quality, and international trade. *Journal of Political Economy*, *119*(4), 721–765.

Feenstra, R., & Romalis, J. (2012). International prices and endogenous quality. *unpublished paper*.

Feenstra, R., Romalis, J., & Schott, P. (2002). *Us imports, exports, and tariff data, 1989-2001* (Tech. Rep.). National Bureau of Economic Research.

Fieler, A. (2011). Nonhomotheticity and bilateral trade: Evidence and a quantitative explanation. *Econometrica*, *79*(4), 1069–1101.

Gagnon, E. (2009). Price setting during low and high inflation: Evidence from Mexico. *The Quarterly Journal of Economics*, *124*(3), 1221.

Gervais, A. (2012). Product quality and firm heterogeneity in international trade. *University of Notre Dame mimeo*.

Goldberg, P., Khandelwal, A., Pavcnik, N., & Topalova, P. (2010). Imported intermediate inputs and domestic product growth: Evidence from India. *The Quarterly Journal of Economics*, *125*(4), 1727–1767.

Goldberg, P., & Pavcnik, N. (2005). Trade, wages, and the political economy of trade protection: Evidence from the Colombian trade reforms. *Journal of International Economics*, *66*(1), 75–105.

Goldberg, P., & Pavcnik, N. (2007). Distributional effects of globalization in developing countries. *Journal of Economic Literature*, *45*, 39–82.

Goldin, C., & Katz, L. (2008). *The race between education and technology.* Harvard University Press.

Hallak, J. (2006). Product quality and the direction of trade. *Journal of International*

Economics, *68*(1), 238–265.
Hallak, J., & Schott, P. (2011). Estimating cross-country differences in product quality. *The Quarterly Journal of Economics*, *126*(1), 417–474.
Halpern, L., Koren, M., & Szeidl, A. (2011). Imports and productivity. *Central European University mimeo*.
Handbury, J. (2012). Are poor cities cheap for everyone? Non-homotheticity and the cost of living across US cities. *Columbia University mimeo*.
Handbury, J., & Weinstein, D. (2011). *Is new economic geography right? Evidence from price data* (Tech. Rep.). National Bureau of Economic Research.
Hufbauer, G., & Schott, J. (1992). *North American free trade: Issues and recommendations.* Peterson Institute.
Iacovone, L. (2008). Exploring Mexican plant data. *University of Sussex mimeo*.
Iacovone, L., Rauch, F., & Winters, L. (2010). *Trade as an engine of creative destruction: Mexican experience with Chinese competition.* London school of economics and political science (LSE). Centre for economic performance (CEP).
INEGI. (2000). Sintesis metodologica de la encuesta industrial anual. *Instituto Nacional de Estadistica, Geographia, e Informatica*.
Johnson, R. (2011). Trade and prices with heterogeneous firms. *Journal of International Economics*.
Kate, A. (1992). Trade liberalization and economic stabilization in Mexico: Lessons of experience. *World Development*, *20*(5), 659–672.
Khandelwal, A. (2010). The long and short (of) quality ladders. *Review of Economic Studies*, *77*(4), 1450–1476.
Konüs, A. (1939). *The problem of the true index of the cost of living.*
Krueger, A. (1999). Trade creation and trade diversion under NAFTA. *National Bureau of Economic Research Working Paper*.
Kugler, M., & Verhoogen, E. (2009). Plants and imported inputs: New facts and an interpretation. In *American economic review papers and proceedings* (Vol. 99, pp. 501–507).
Kugler, M., & Verhoogen, E. (2011). Prices, plant size, and product quality. *Review of Economic Studies*, *Forthcoming*.
Leamer, E. (1996). Wage inequality from international competition and technological change: Theory and country experience. *The American Economic Review*, *86*(2), 309–314.
Li, N. (2012). An Engel curve for variety: Household food consumption and welfare in India. *University of Toronto mimeo*.
López-Córdova, J. (2002). NAFTA and Mexico's manufacturing productivity: An empirical investigation using micro-level data. *Inter-American Development Bank Working Paper). IDB, Washington, DC*.
Mandel, B. (2010). Heterogeneous firms and import quality: Evidence from transaction-level prices. *International Finance Discussion Papers*.
Manova, K., & Zhang, Z. (2012). Export prices across firms and destinations. *The Quarterly Journal of Economics*, *127*(1), 379–436.
Markusen, J. (1986). Explaining the volume of trade: An eclectic approach. *The American Economic Review*, 1002–1011.
Matsuyama, K. (2000). A ricardian model with a continuum of goods under nonhomothetic preferences: Demand complementarities, income distribution, and North-South trade.

Journal of Political Economy, *108*(6), 1093–1120.

Melitz, M. (2003). The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica*, *71*(6).

Moretti, E. (2011). Real wage inequality. *UC Berkeley mimeo*.

Muellbauer, J. (1974). Prices and inequality: The United Kingdom experience. *The Economic Journal*, *84*(333), 32–55.

Porto, G. (2006). Using survey data to assess the distributional effects of trade policy. *Journal of International Economics*, *70*(1), 140–160.

Quah, D. (1993). Galton's fallacy and tests of the convergence hypothesis. *The Scandinavian Journal of Economics*, 427–443.

Rodriguez, C. (1979). The quality of imports and the differential welfare effects of tariffs, quotas, and quality controls as protective devices. *The Canadian Journal of Economics*, *12*(3), 439–449.

Sato, K. (1976). The ideal log-change index number. *The Review of Economics and Statistics*, *58*(2), 223–228.

Slaughter, M. (1998). What are the results of product-price studies and what can we learn from their differences? *National Bureau of Economic Research Working Paper*, *No 6591*.

Sutton, J. (1998). *Technology and market structure: Theory and history.* The MIT Press (Cambridge, Mass).

Topalova, P. (2010). Factor immobility and regional impacts of trade liberalization: Evidence on poverty from India. *American Economic Journal: Applied Economics*, *2*(4), 1–41.

Trefler, D. (1995). The case of the missing trade and other mysteries. *The American Economic Review*, 1029–1046.

Vartia, Y. (1976). Ideal log-change index numbers. *Scandinavian Journal of Statistics*, 121–126.

Verhoogen, E. (2008). Trade, quality upgrading, and wage inequality in the Mexican manufacturing sector. *The Quarterly Journal of Economics*, *123*(2), 489–530.

Appendix

Appendix Figures

Figure A.1: Two Notable Features about Mexican and Developing Country Imports

Figure A.1.A: What Do Developing Countries Import?

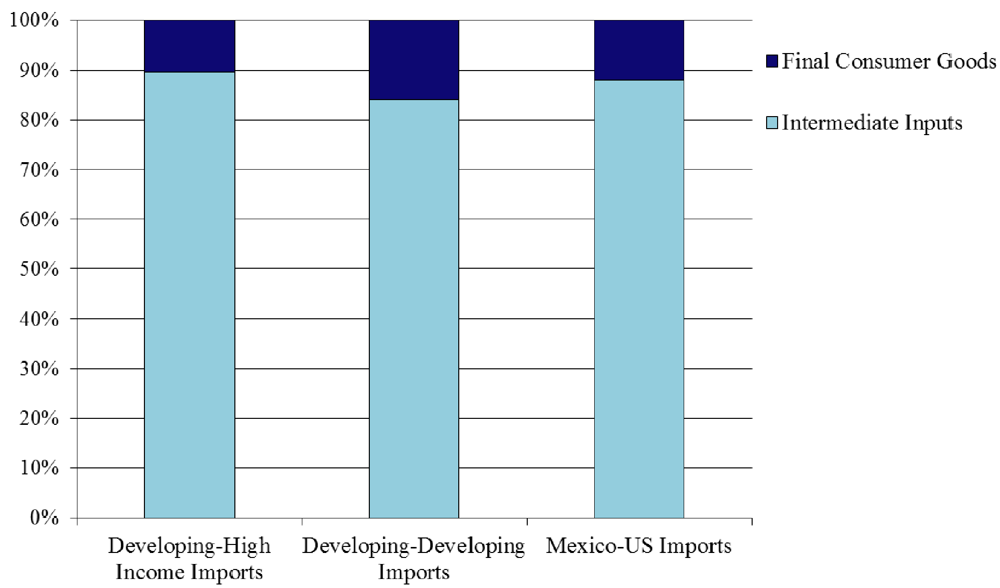

Figure A.1.B: Variation in Use of Imported Inputs

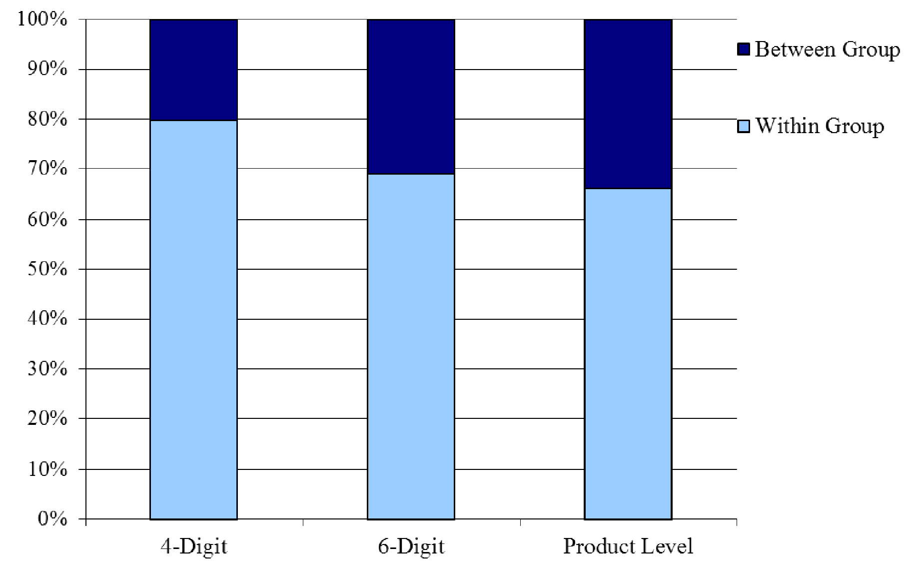

Figure A.1.A depicts end use shares in the sum of imports across different pairs of countries over the period 1994-2000. From left to right the bars depict developing country imports from high income countries, developing country imports from developing countries, and Mexican imports from the US. “Developing” and “High Income” refer to low-and-middle income and high income countries according to the World Bank’s classification in 2010 respectively. Data on import flows are from the World Bank’s WITS database. End-use classifications into final consumption and intermediate goods are based on BEC classifications. Figure A.1.B depicts the variance decomposition of plant level imported input shares into between and within product group components. The estimates are from Mexican plant data in 1994 and conditional on state fixed effects. “4-digit”, “6-digit”, and “Product Level” refer to 80, 203, and roughly 3234 manufacturing product groups respectively.

Figure A.2: Mexican Tariffs on US Imports 1993-2000

Figure A.2.1: Average Applied Tariff Rates

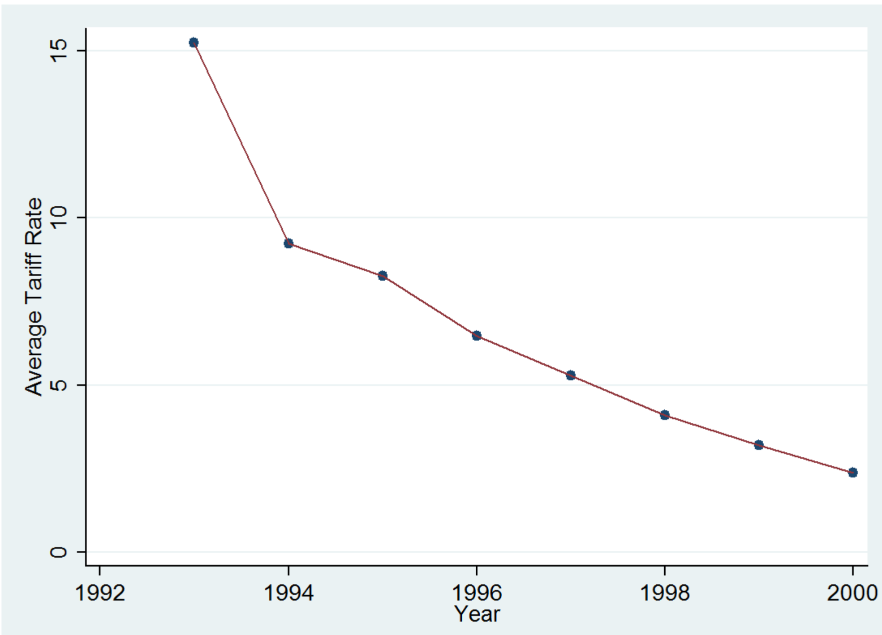

Figure A.2.2: Sectoral Variation

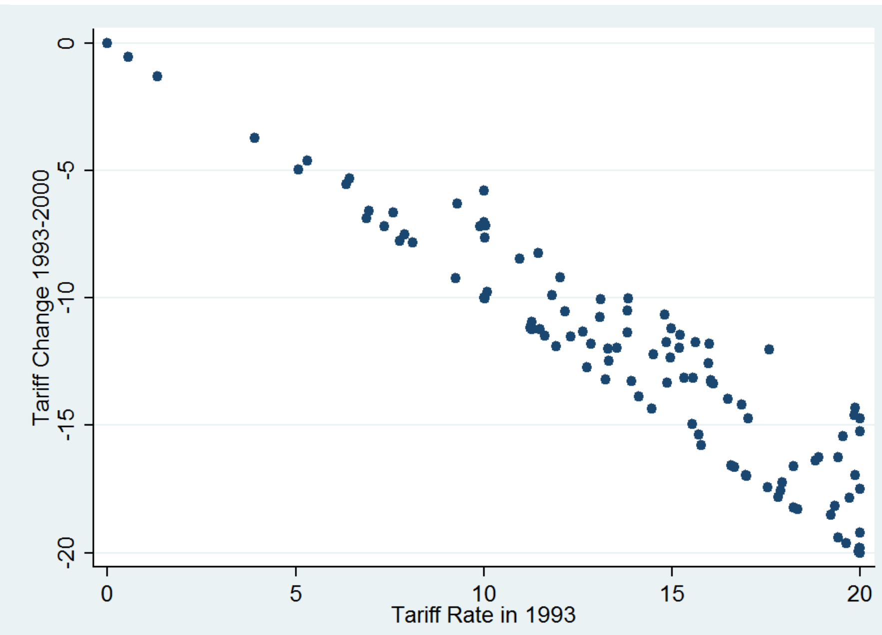

Figure A.2.A depicts average applied tariff rates on US manufacturing imports over the period 1993-2000. Figure A.2.B depicts the relationship between changes and initial levels of average applied tariff rates across four digit US manufacturing sectors between 1993-2000. Source: Secretaria de Economia.

Figure A.3: The Peso Crisis and Spike of Price Changes 1994-1995

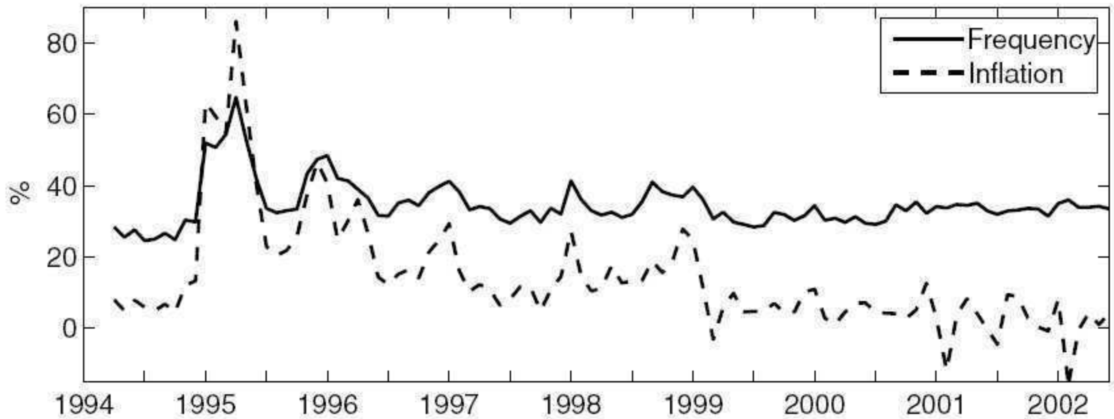

The figure is taken from Gagnon (2009, pp. 1233). It depicts monthly frequencies of price changes and CPI inflation for non-regulated goods and services.

Figure A.4: Unit Value Elasticities and Alternative Measures of Vertical Differentiation

Figure A.4.1: Unit Values and Sales

Figure A.4.2: Unit Values and Imported Input Shares

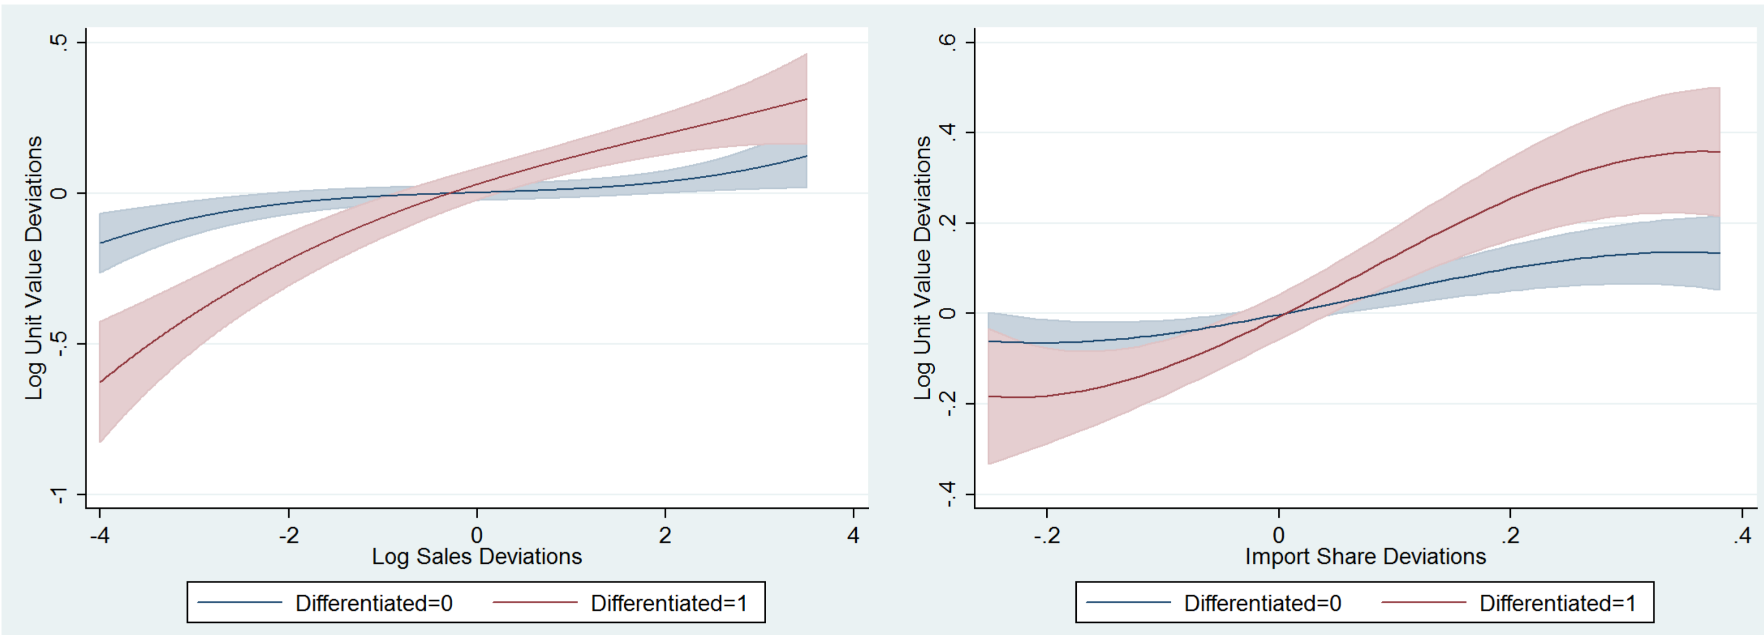

The fitted relationships correspond to the best fitting polynomial functional form according to the Akaike Information Criterion (AIC). The y-axis in both graphs depict the residuals of a regression of log product-line unit values on month-by-product and month-by-state fixed effects. The x-axis in A.4.1 depicts residuals of a regression of log product line sales on the same fixed effects. The x-axis in A.4.2 depicts residuals of a regression of annual plant level imported input shares on state and product fixed effects. The sample is broken down into four-digit sectors with above or below mean shares of R&D and advertisement over sales following Sutton (1998) (Differentiated=1,0 respectively). The bottom and top 0.5% are excluded from the graph. Standard errors are clustered at the product level and the shaded area indicates 95% confidence intervals.

Figure A.5: Alternative Estimates of NAFTA's Effect on Cost of Living Inflation

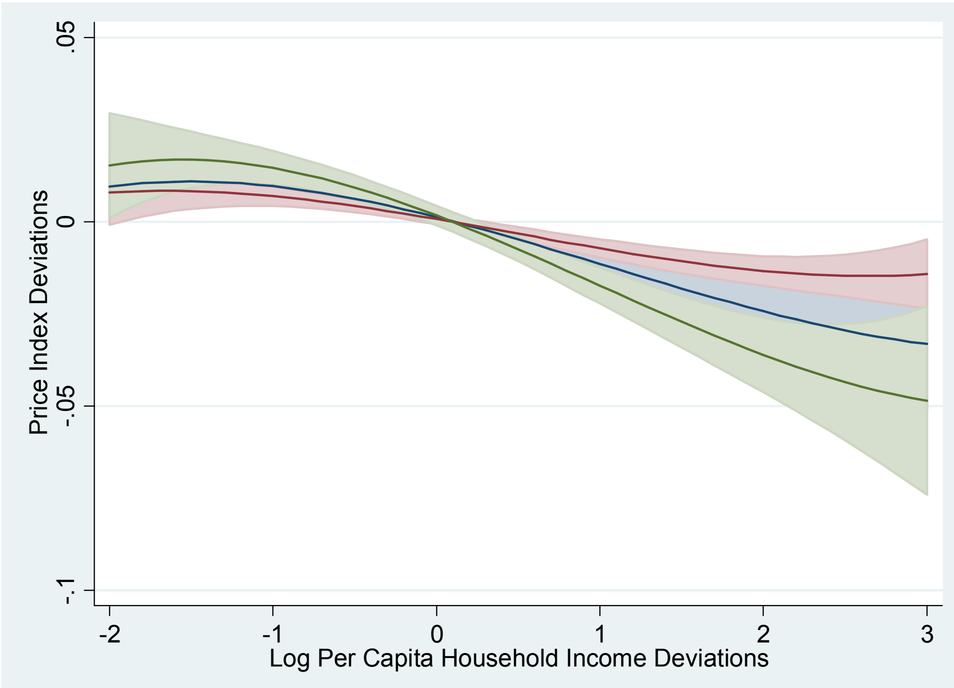

The fitted relationship corresponds to the best fitting polynomial functional form according to the Akaike Information Criterion (AIC). Standard errors are clustered at the municipality level and the shaded area indicates 95% confidence intervals. The y-axis depicts mean deviations of estimated household cost of living inflation of tradable consumption due to a 12% tariff cut on US imports. The steepest, middle, and flattest functions are based on "upper bound", "baseline", and "lower bound" estimations respectively. Estimations are based on urban Mexican households in 1994 and subject to nationally representative sample weights. The bottom and top 0.5% on the x-axis are excluded from the graph.

Appendix Tables

Table A.1: Coverage of Store Price Sample in Total Household Consumption

	July, August, September 1994	July, August, September 2000	1994-2000
Number of Monthly Store Price Quotes in 35 Cities	28515	40280	
Tradable Sample	24089	33699	
Processed Tradeable Estimation Sample	16792	22235	
Persistent Barcode Series in Estimation Sample	13589	13589	13589
Product Replacements in Estimation Sample			3768
Basket Net Expansion in Estimation Sample			5866
Tradable Coverage in Total Urban Household Consumption (from ENIGH Consumer Surveys)	0.54	0.50	
Processed Tradable Sample Coverage in Tradable Consumption	0.70	0.66	

Table A.2: Plant Data Descriptive Statistics

	Full Sample	Final Good Sectors Only
Number of 6-Digit Sectors	203	79
Number of Establishments	6341	2762
Number of Products Reported Over 12 Months	3234	1331
Number of Month * Establishment * Product Observations	257736	136440
Average Number of Products Per Establishment	3.4	4.1
Median Employment Size	103	122

Table A.3: Household Consumption Survey Descriptive Statistics

	Urban Household Sample
Number of Households	7764
Number Of Municipalities	236
Total Number of Reported Transactions Across all Expenditure Categories	524782
Number of Reported Transactions In Processed Tradables Sample (255 Product Groups)	279584
Number of Transactions In Processed Tradables Sample With Unit Values	122572
Share of Processed Tradables Transactions at Markets	0.081
Share of Processed Tradables Transactions at Street Vendors	0.112
Share of Processed Tradables Transactions at Convenience and Specialized Stores	0.485
Share of Processed Tradables Transactions at Supermarkets and Department Stores	0.264

Table A.4: Input Sectors with Highest Median Total Requirement Coefficients across Mexican Destination Sectors

Rank	NAICS 4-Digit Sector	NAICS Description
1	3251	Fabricación de productos químicos básicos
2	3241	Fabricación de productos derivados del petróleo y del carbón
3	3261	Fabricación de productos de plástico
4	3222	Fabricación de productos de papel y cartón
5	3221	Fabricación de celulosa, papel y cartón
6	3363	Fabricación de partes para vehículos automotores
7	3252	Fabricación de hules, resinas y fibras químicas
8	3231	Impresión e industrias conexas
9	3259	Fabricación de otros productos químicos
10	3255	Fabricación de pinturas, recubrimientos, adhesivos y selladores
11	3311	Industria básica del hierro y del acero
12	3399	Otras industrias manufactureras
13	3132	Fabricación de telas
14	3211	Aserrado y conservación de la madera
15	3312	Fabricación de productos de hierro y acero de material comprado

Table A.5: Technology Parameter Estimates from Plant Microdata

Estimation Samples	Dependent Variable: ln(Unit Value)	(1) OLS	(2) OLS	(3) IV
All 6-Digit Manufacturing Sectors	ln(Sales)	0.0486***		0.0444***
		(0.00594)		(0.0123)
	ln(Employment)		0.0363***	
			(0.0102)	
	Obs	170240	167449	160835
	N(Plants)	5779	5665	5665
	Within-R^2	0.017	0.009	
Final Goods 6-Digit Manufacturing Sectors	ln(Sales)	0.0557***		0.0733***
		(0.00781)		(0.0159)
	ln(Employment)		0.0598***	
			(0.0133)	
	Obs	94741	93154	91064
	N(Plants)	2656	2602	2602
	Within-R^2	0.024	0.013	
Differentiated Final Goods 6-Digit Manufacturing Sectors	ln(Sales)	0.104***		0.240***
		(0.0159)		(0.0315)
	ln(Employment)		0.209***	
			(0.0300)	
	Obs	24312	24123	23392
	N(Plants)	672	666	666
	Within-R^2	0.065	0.066	

All regressions include state-by-month and month-by-product fixed effects. Product grous refer to several thousand disaggregate product descriptions within 203 6-digit manufacturing sectors. Unit values and sales vary across plants, product lines within plants, and months. Annual employment varies across plants. The first stage regressions of ln(Sales) on ln(Employment) not reported here are highly statistically significant. Standard errors are clustered at the product level. ***1%, **5%, and *10% significance levels.

Table A.6: Unit Value-Sales Elasticities and Alternative Measures of Vertical Differentiation

Dependent Variable: ln(Unit Value)	(1)	(2)	(3)	(4)	(5)	(6)
ln(Sales)	0.0557***	-0.00720				
	(0.00781)	(0.00831)				
ln(Sales) * R&D/Advert Intensity		1.091***				
		(0.126)				
ln(Employment)			0.0598***	-0.0410***		
			(0.0133)	(0.0150)		
ln(Employment) * R&D/Advert Intensity				1.880***		
				(0.279)		
Import Share					0.576***	0.00849
					(0.103)	(0.108)
Import Share * R&D/Advert Intensity						10.23***
						(1.708)
State Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
State*Month Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Product*Month Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Obs	94741	94741	93154	93154	93996	93996
N(Plants)	2656	2656	2602	2602	2625	2625
Within-R^2	0.024	0.041	0.013	0.026	0.026	0.026

All regressions include state-by-month and month-by-product fixed effects. "R&D/Advert Intensity" refers to US shares of R&D and advertising expenditures in firm sales, averaged to the four digit SITC industry sectors. These measures are reported in Kugler and Verhoogen (2011) and were matched to Mexican six digit final goods manufacturing industries. Standard errors are clustered at the product level. ***1%, **5%, and *10% significance levels.

Table A.7: Do the Rich and the Poor Consume Identical Items at Different Prices?

Dependent Variable: ln(Unit Value)	(1)	(2)
2nd Per Capita Household Income Quintile	0.00297	0.00962
	(0.0106)	(0.00873)
3rd Per Capita Household Income Quintile	0.0288**	0.0426***
	(0.0125)	(0.0110)
4th Per Capita Household Income Quintile	0.0454***	0.0573***
	(0.0148)	(0.0122)
5th Per Capita Household Income Quintile	0.0803***	0.0983***
	(0.0197)	(0.0146)
City Fixed Effects	Yes	Yes
Product Group Fixed Effects	Yes	Yes
City-Product Group Fixed Effects	Yes	Yes
Store Type Fixed Effects	-	Yes
City-Store Type Fixed Effects	-	Yes
Product-Store Type Fixed Effects	-	Yes
City-Product-Store Type Fixed Effects	-	Yes
Obs	122,572	122,572
N(Households)	7632	7632
R^2	0.894	0.929

The dependent variable is reported purchase unit values in 118 out of a total of 255 processed tradable product groups in the 1994 household consumption survey. Income quintiles are based on per capita household incomes. Household consumption surveys include the following store types: Markets, street vendors, convenience and specialized stores, and supermarkets and department stores. Standard errors are clustered at the municipality level. ***1%, **5%, and *10% significance levels.

Table A.8: Rich-Poor Price Gaps and Product Group Differentiation

Dependent Variable: ln(Unit Value)	(1)
ln(Income per Capita)	0.0317***
	(0.00559)
ln(Income per Capita)*Scope	0.0750***
	(0.0278)
City Fixed Effects	Yes
Product Group Fixed Effects	Yes
City-Product Group Fixed Effects	Yes
Store Type Fixed Effects	Yes
City-Store Type Fixed Effects	Yes
Product-Store Type Fixed Effects	Yes
City-Product-Store Type Fixed Effects	Yes
Obs	122,572
N(Households)	7632
R^2	0.930

The dependent variable is reported purchase unit values in 118 out of a total of 255 processed tradable product groups in the 1994 household consumption survey. Income quintiles are based on per capita household incomes. Household consumption surveys include the following store types: Markets, street vendors, convenience and specialized stores, and supermarkets and department stores. "Scope" is the product group specific estimate of the unit value-sales elasticity observed in the plant microdata in 1994. Standard errors are clustered at the municipality level. ***1%, **5%, and *10% significance levels.

Additional Results of the Model

For ease of exposition product group subscripts k are suppressed unless indicated otherwise.

Preferences

From preferences in (1) of the paper, consumer optimization yields the following expression for household expenditure:

$$p_i x_{hi} = y_{hi} = \frac{p_i^{1-\sigma}\left(q_i^{\varphi_h}\right)^{\sigma\text{-}1}}{\int_{j=0}^{N} p_j^{1-\sigma}\left(q_j^{\varphi_h}\right)^{\sigma\text{-}1} dj} y_h = s_{hi} y_h$$

s_{hi} is the household's expenditure share on variety i and y_h is its total expenditure on product group k. Within product group expenditure shares increase in quality valuation for products with above average quality, and decrease in quality valuation for below average quality products:

$$\frac{\partial s_{hi}}{\partial \varphi_h} = (\sigma\text{-}1) s_{hi}\left(\ln q_i \text{-} \sum\nolimits_I s_{hi} ln q_i\right)$$

Weighted average quality of the household's consumption basket increases in quality valuation:

$$\frac{\partial\left(\sum_I s_{hi}(ln q_i - \left(\frac{1}{I}\sum_I ln q_i\right)\right)}{\partial \varphi_h} = \text{I*Cov}\left(\left(\frac{\partial s_{hi}}{\partial \varphi_h}\right), \left(ln q_i - \overline{ln q_i}\right)\right) > 0$$

Finally, I solve for the elasticity of market sales with respect to product quality to get:

$$\frac{d\left(\sum_H s_{hi}y_h\right)}{\left(\sum_H s_{hi}y_h\right)}/\frac{dq_i}{q_i}=(\sigma-1)\left(\sum_H\frac{y_{hi}}{y_i}\varphi_h\right)=(\sigma-1)\,\varphi_i^*$$

And following the discussion in the text, I define: $q_i^*={q_i}^{\varphi_i^*}$.

Technology

The unit cost functions for production technologies in (3) and (5) of the paper are $c_i=\lambda_i^{-1}p_{im}$ and $c_j=\mathrm{w}z_j^{\delta}$ respectively. Substituting intermediate unit costs into final product unit costs, we get: $c_i=\lambda_i^{-1}\mathrm{w}z_i^{\delta}$. Plant profits are given by:

$$\pi_i=(p_i\text{-}c_i)x_i\text{-f}=(p_i\text{-}\lambda_i^{-1}\mathrm{w}z_i^{\delta})x_i\text{-f}$$

In quality adjusted terms, profits can be written as:

$$\pi_i=\left(\frac{p_i}{q_i^*}-\frac{c_i}{q_i^*}\right)q_i^*x_i\text{-f}=\left(\frac{p_i}{q_i^*}-\lambda_i^{-1}wz_i^{\delta}q_i^{*\text{-}1}\right)q_i^*x_i\text{-f}$$

Equilibrium

Firms simultaneously choose product quality and quality adjusted prices to maximize profits. From the profit equations it is clear that maximizing profits with respect to product quality, implies minimizing $\frac{c_i}{q_i^*}$ with respect to q_i^*, that is minimizing the average variable cost per unit of product quality.[56] The additional parameter restriction $\gamma>\delta$ assures a well behaved optimum. This provides an expression for equilibrium product quality as a function of technical efficiency:

$$q_i^*=\left(\frac{\alpha\gamma}{\gamma-\delta}\right)^{1/\theta}\lambda_i^{\psi}$$

We next solve for intermediate input quality as a function of plant efficiency:

$$z_i=\left(\frac{1}{1-\alpha}\left(\frac{\alpha\gamma}{\gamma-\delta}\right)^{\psi/\vartheta}-\frac{\alpha}{1-\alpha}\right)^{1/\vartheta\gamma}\lambda_i^{\psi/\gamma}$$

Equilibrium intermediate input quality is thus increasing in plant efficiency. Alternatively, we can solve for λ_i as a function of product quality and substitute back into the unit cost function to derive the equilibrium relationship between final product quality and unit costs:

$$c_i=\mathrm{w}\left(\frac{1}{\alpha}-\frac{\delta}{\alpha\gamma}\right)^{-1/\vartheta\psi}\left(\frac{1}{1-\alpha}\frac{\delta}{\gamma}\right)^{\delta/\vartheta\gamma}q_i^{*\eta}$$

$\eta=\frac{\delta}{\gamma}\text{-}\frac{1}{\psi}$ is the elasticity of unit costs with respect to final product quality. The equilibrium relationship between product quality and the inverse of quality adjusted marginal costs becomes:

[56] The same condition is present in Mandel (2010), Johnson (2011), and Feenstra and Romalis (2012). The latter paper attributes this finding to Rodriguez (1979).

$$\frac{\partial ln q_i^*}{\partial ln\left(q_i^*/c_i\right)}=\frac{1}{1-\eta}$$

Given CES preferences, the equilibrium relationship between observed unit values and product quality can then be expressed as:

$$p_i=\frac{\sigma}{\sigma-1}c_i=\frac{\sigma}{\sigma-1}\mathrm{w}\left(\frac{1}{\alpha}-\frac{\delta}{\alpha\gamma}\right)^{-1/\vartheta\psi}\left(\frac{1}{1-\alpha}\frac{\delta}{\gamma}\right)^{\delta/\vartheta\gamma}q_i^{*\eta}$$

Finally, the equilibrium relationship between unit values and quality adjusted productivity becomes:

$$\frac{\partial ln p_i}{\partial ln\left(q_i^*/c_i\right)}=\text{-}1+\frac{\partial ln q_i^*}{\partial ln\left(q_i^*/c_i\right)}=\frac{\eta}{1-\eta}$$

Starting from an initial equilibrium outcome of input quality and output quality choices across plants, the observed final product unit value elasticity with respect to product quality is as derived in the above: $\frac{\partial ln p_i}{\partial ln q_i^*}=\frac{\partial ln c_i}{\partial ln q_i^*}=\eta=\frac{\delta}{\gamma}\text{-}\frac{1}{\psi}$. Import access then lowers the relative cost of inputs with higher foreign value shares which is captured by a change in δ. Denote import tariffs by τ_k, then the cross-derivative expression of interest becomes:

$$\frac{\partial^2 ln p_i}{\partial ln q_i^*\partial\tau}=\frac{\partial^2 ln c_i}{\partial ln q_i^*\partial\tau}=\frac{1}{\gamma}\frac{\partial\delta}{\partial\tau}>0$$

The general equilibrium solution of the model closely follows Melitz (2003). To assure finite means in efficiency draws and final product plant revenues, the shape parameter of the pareto distribution needs to have a lower bound at $\xi>\max(\psi(\sigma-1)(1-\eta),1)$. The cut-off values are determined by two conditions.

First, profits of the marginal plant must be zero: $\pi(\lambda^*)=\frac{r^*(\lambda)}{\sigma}\text{-f}=0$. Second, free entry implies that ex ante expected profits are zero: $(1\text{-G}(\lambda^*))\sum_{t=0}^{\infty}(1\text{-}\chi)^t\left(\frac{E(r^*(\lambda))}{\sigma}-f\right)\text{-}f_e=0$. Using these two conditions, and that $\frac{r^*(\lambda)}{r^*(\lambda^*)}=\left(\frac{\lambda}{\lambda^*}\right)^{\psi(\sigma-1)(1-\eta)}$, we get: $\mathrm{E}(r^*(\lambda))=\frac{\xi}{\xi-\psi(\sigma-1)(1-\eta)}\sigma\mathrm{f}$. It follows that:

$$\lambda^*=\lambda_m\left(\frac{f}{f_e\chi}\left(\frac{\xi}{\xi-\psi(\sigma-1)(1-\eta)}-1\right)\right)^{1/\xi}$$

Finally, the free entry condition in combination with the condition that in steady state the mass of new entrants is equal to the mass of exiting firms $M_e(1\text{-G}(\lambda^*))=\chi\mathrm{M}$, and labor market clearing ($L=MEr^*(\lambda)-\Delta+M_ef_e$), where L is labor supply and Δ is the difference between final sector revenues and profits, pin down the mass of final good producers in steady state:

$$\mathrm{M}=\frac{L(\xi-\psi(\sigma-1)(1-\eta))}{\xi\sigma f}$$

Cost of Living Implications

Based on the work of Sato (1976) and Vartia (1976), the ideal price index for a homothetic CES utility function is:

$$\left(\frac{e\left(u^{*},p^{t1}\right)}{e\left(u^{*},p^{t0}\right)}\right)_{h}=\prod\nolimits_{I}\left(\frac{p_{ki}^{t1}}{p_{ki}^{t0}}\right)^{\omega_{hki}}\text{, where }\omega_{hki}=\frac{s_{hki}^{t1}-s_{hki}^{t0}}{ln\left(s_{hki}^{t1}-s_{hki}^{t0}\right)}/\left(\sum\nolimits_{I}\frac{s_{hki}^{t1}-s_{hki}^{t0}}{ln\left(s_{hki}^{t1}-s_{hki}^{t0}\right)}\right)$$

I is the number of all varieties pooled across all product groups k in the economy. Household cost of living inflation is a weighted geometric mean of price changes where the weights are ideal log changes of household budget shares. In the following, I will refer to two representative consumers that can be thought of as a poor and a rich household denoted by subscripts P and R. Taking log differences in household cost of living inflation between a poor and a rich household, we then get:

$$\ln\left(\frac{e\left(u^{*},p^{t1}\right)}{e\left(u^{*},p^{t0}\right)}\right)_{P}\text{-}\ln\left(\frac{e\left(u^{*},p^{t1}\right)}{e\left(u^{*},p^{t0}\right)}\right)_{R}=\sum\nolimits_{I}\left(s_{kiP}^{t0}-s_{kiR}^{t0}\right)\ln\left(\frac{p_{ki}^{t1}}{p_{ki}^{t0}}\right)$$

which is (10) in the text. In the presence of non-homotheticity in (1), the Sato-Vartia ideal price index, in principle, does not hold because income changes affect expenditure shares so that the ideal weights (ω_{hki}) cease to hold (e.g. Diewert, 1979). Within the structure of the model, however, (10) represents the difference in the exact ideal price index between two representative agents due to a *ceteris paribus* change in the relative price of quality. The two underlying assumptions are that i) CES preferences in (1) hold so that elasticities of substitution are the same across households, and ii) we abstract from general equilibrium consequences of import access on relative incomes. If either of the model's assumptions is violated, then (10) remains an approximation of the true difference in cost of living to the first order, because as in a Laspeyres price index its weights are based on differences in initial expenditure weights.

The second welfare distributional implication concerns price levels rather than price changes. Using the same notation as above, we start with an expression for log differences in weighted average product quality between a poor and a rich household denoted by $\ln Q_P$-$\ln Q_R$:

$$\ln Q_{P}\text{-}\ln Q_{R}=\sum\nolimits_{I}\left(s_{kiP}^{t0}-s_{kiR}^{t0}\right)\left(lnq_{ki}^{*}-\overline{lnq_{ki}^{*}}^{k}\right)$$

Substituting for product quality by the equilibrium relationship to quality adjusted costs, we get:

$$\ln Q_{P}\text{-}\ln Q_{R}=\sum\nolimits_{K}\frac{1}{1-\eta_{k}}\sum\nolimits_{Ik}\left(s_{kiP}^{t0}-s_{kiR}^{t0}\right)\left(ln\left(\frac{q_{ki}^{*}}{c_{ki}}\right)-\overline{ln\left(\frac{q_{ki}^{*}}{c_{ki}}\right)}^{k}\right)$$

Denoting differences in weighted average inverse quality adjusted prices by $\ln\left(\frac{Q}{P}\right)_P$-$\ln\left(\frac{Q}{P}\right)_R$, we get:

$$\ln\left(\frac{e\left(u^{*},p^{t1}\right)}{e\left(u^{*},p^{t0}\right)}\right)_{P}\text{-}\ln\left(\frac{e\left(u^{*},p^{t1}\right)}{e\left(u^{*},p^{t0}\right)}\right)_{R}=\ln\left(\frac{Q}{P}\right)_{P}\text{-}\ln\left(\frac{Q}{P}\right)_{R}=\sum\nolimits_{K}\left(1-\eta_{k}\right)\left(lnQ_{Pk}-lnQ_{Rk}\right)$$

Chapter 3

A Peculiar Trade: Tourism and Industrialization in a Developing Country*

Benjamin Faber†

Abstract

Tourism is one of the most visible and fastest growing facets of globalization. Despite its impressive scale as a channel of market integration, existing empirical evidence of tourism's economic implications is rather limited. This paper proposes a novel empirical strategy to contribute to our understanding of this question. To address the many endogeneity concerns that arise when causally relating tourism to economic outcomes, I propose an instrumental variable strategy that exploits geological and oceanographic variation in beach quality along the Mexican coastline. I then apply this methodology to answer one particular question of policy interest. From a trade theory perspective, tourism constitutes exports of non-tradable amenities and services via traveling consumers rather than shipping goods. Tourist arrivals thus lead to an increase in local consumption expenditure that resembles the spending effect in neoclassical models of the Dutch disease. The arising question is whether tourism negatively affects the competitiveness of tradable production sectors. I find that while conventional regressions display a strongly significant positive correlation between tourism GDP and municipality industrialization, this relationship is reversed in a statistically and economically significant way in instrumental variable estimations. In the light of additional evidence that I document on factor immigration as a consequence of tourism, these results should be interpreted as a lower bound compared to a cross-country setting where factors are less mobile to follow tourism induced relative price movements.

Keywords: Market integration; tourism; Dutch disease
JEL Classification: F15; O24

---

*I am grateful to Gene Feldman and Norman Kuring at NASA for their technical support on the satellite remote sensing data. I thank Dave Donaldson, Esther Duflo, Horacio Larreguy Abresu, Daniel Sturm, and Anthony Venables for very helpful conversations. All remaining errors are mine alone. This project was supported by research grants from the UK Economic and Social Research Council.

†London School of Economics and Centre for Economic Performance; Email: b.s.faber@lse.ac.uk

1 Introduction

Over the past decades tourism exports have grown to be a quantitatively important channel of global integration. Figure 1 depicts world tourism exports relative to global services exports, merchandise exports, and foreign direct investment flows over the decade 1995-2005. Tourism is the single largest sector of global services trade, accounting for 30-40 percent of total trade. In terms of magnitude it has been close to equivalent to world trade in food and agriculture, and world foreign direct investment flows over the same period. Perhaps as a consequence of these figures, tourism has also attracted vast policy attention in both developed and developing countries. Virtually every country in the world has one or several publically funded tourism promotion agencies.[1]

Despite tourism's impressive scale as a form of market integration and remarkable policy interest, our existing empirical evidence of its economic implications is rather limited. This paper proposes a novel empirical strategy to contribute to our understanding of this question. To address the many endogeneity concerns that arise when causally relating tourism activity to economic outcomes, I propose an instrumental variable strategy based on geological and oceanographic variation in beach quality along the Mexican coastline.

I then apply this strategy to answer one particular question of policy interest. From a trade theory perspective, tourism is equivalent to exporting non-tradable amenities and services via traveling consumers rather than shipping goods. As first noted by Copeland (1991), tourist arrivals thus lead to an increase in local consumption expenditure that resembles the spending effect in neoclassical models of the Dutch disease (e.g. Corden and Neary, 1982). The arising question is whether tourism reduces the competitiveness of tradable production sectors.

The main empirical motivation of this paper is captured in Figure 2, which depicts cross country scatter plots of log manufacturing production or manufacturing export shares in total domestic GDP on log tourism exports. There is a strong and statistically significant positive correlation between the national scale of tourism exports and measures of industrialization and manufacturing export specialization. The central question that the empirical analysis seeks to answer is whether this positive relationship is causal and remains robust after addressing the many endogeneity concerns underlying these observed relationships.

[1] Switch on CNN International or BBC Worldwide in a hotel room and you will notice from the advertisements.

The empirical analysis is based on variation in tourism activity and sectoral production in a within country setting across Mexican municipalities. The explanatory variable of central interest will be municipality level GDP of the hotel and lodging sector reported in the establishment surveys of the Mexican Censos Economicos in 1998 and 2003. Figure 3 presents an overview of the relative importance and historical development of tourism related GDP in the Mexican economy. In Mexico tourism related activities have accounted for close to 10% of overall GDP over the past decade. As shown in Figure 3.B the great majority of this activity is actually driven by domestic tourism expenditure rather than international tourists. While international tourism is more visible and more frequently discussed, it is domestic tourists that account for the great majority of tourism expenditure.[2] Figure 3.C depicts the evolution of tourist arrivals from the historical archive of the Mexican central bank since 1962. Mexican tourism appears to have started reaching a relevant economic scale sometime during the 1980s.

To address the empirical concern that municipality level hotel and lodging sector activity is not randomly assigned, I construct two instrumental variables that the tourism management literature has mentioned among the determinants of touristic potential (e.g. Leatherman, 1997). The first concerns the scenic beauty of a beach location and is proxied by the presence of an off-shore island within 10 kilometer off the coastline, which I construct from geological maps of the Mexican maritime territory. The second is a measure of water turbidity (the opposite of clarity) in close proximity to the shoreline which I construct from remote sensing satellite data. Figure 4 provides a cartographic illustration of these variables in a geographical information system (GIS) across Mexican municipalities.

The identifying assumption underlying this strategy is that these natural characteristics affect municipality level economic outcomes only through their effect on tourism potential. To test the robustness of the exclusion restriction I estimate a placebo falsification test to verify that the instrumental variables have no effect on observable county outcomes at a time before Mexican tourism had reached a relevant economic scale in 1960.

The main findings can be summarized as follows. I find that while conventional regressions display a strong positive correlation between local tourism GDP and municipality industrialization, this relationship is reversed in a statistically and economically significant way in instrumental variable estimations. In particular, the tourism induced expansion of the services

[2] This finding is not particular to Mexico and holds for the majority of countries. This is apparent in summary statistics provided in the annual reports of the UN World Tourism Organization at www.unwto.org/en.

sector appears to be at the expense of industrial production rather than the agricultural sector. These findings become stronger as one starts conditioning on levels of local infrastructure endowments that are themselves affected by tourism development. In the light of parallel empirical evidence of strong factor immigration responses to local tourism activity, these results should be interpreted as a lower bound to a cross-country setting where factors are less able to offset tourism induced relative product and factor market price changes.

This paper is most closely related to the existing empirical literature on the determinants and economic consequences of tourism. Eilat and Einav (2004) use variation of bilateral country pair tourism flows over time to estimate the partial effects of factors such as political risk or exchange rates on bilateral tourism demand. Nunes and Sequeira (2008) use time series variation across countries to test the effect of tourism specialization on country growth. Arezki et al. (2009) apply an instrumental variable strategy based on UNESCO World Heritage sites to estimate this relationship in a cross-section of countries. Both studies find a positive relationship between tourism specialization and a country's economic growth.

The present paper, on the other hand, proposes a novel empirical strategy to contribute to our understanding of the economic consequences of tourism within a developing country. The paper's conceptual focus is on the effects of tourism on tradable goods production within a developing country. In ongoing work in progress, I use the presented empirical strategy to shed light on the consequences of local tourism activity on average real incomes of native workers and migrants, as well as heterogenous real income effects across income and skill distributions within municipalities.

The paper proceeds in four sections. Section 2 provides a brief discussion of the microeconomic channels through which local tourism activity can affect the composition of domestic production. Section 3 describes the data. Section 4 discusses the empirical strategy and presents estimation results. Section 5 concludes.

2 Tourism and Tradable Goods Production

Corden and Neary (1982) distinguish between a "resource movement effect" and a "spending effect" of a domestic resource boom. The former effect captures that the demand for the factors of production employed in the booming sector will increase, thereby drawing away factors from other sectors of production. The latter effect captures that the resulting increase of factor incomes will lead to an increase of the relative price of non-tradables over tradables under the assumption that the income elasticity of services is positive while tradable goods prices are pinned down by world market prices. Both of these mechanisms are modeled with a focus on the cross-country level without factor mobility.

The Dutch disease result is that the resource boom hurts the competitiveness of the tradable goods sector (manufacturing). In principle, both of the above effects can contribute to this outcome. However, it is only the spending effect that is a priori unambiguous in this respect. The only necessary assumption for this result is that there is a non-tradable sector with an income elasticity of demand that is greater than zero. In contrast, the consequences of the resource movement effect on sectoral outputs and output shares will in general depend on the structure of factor intensities, the number of sectors and the degree of factor mobility across sectors.

The simplest comparative static to conceptualize tourism in this setting is as an exogenous increase to local consumer incomes and spending. The assumption behind this comparative static is that tourists are drawn from the identical pool of representative agents as the destination's local population. Tourists then arrive at the destination in order to benefit from a range of non-tradable local amenities that are both non-excludable and non-rival (e.g. beaches), and during the time of their visit they bring with them the same consumption expenditure that locals spend.

In this simple setting, tourism would not give rise to any direct resource movement effect as the "booming sector" refers to the consumption of local amenities with public good character. Tourism would, however give rise to a significant spending effect as the number of local consumers is effectively multiplied by the arrival of tourists (but not more laborers as these consumers just visit). This direct spending effect, would in turn give rise to a second round effect that consists of both resource movement and spending channels. Higher local consumption expenditure, holding everything else constant, leads to higher local demand for non-tradable goods thereby increasing the relative price of that sector relative to tradables. The resource movement is thus an indirect effect of the first order spending increase and is directly in favor (not ambiguously) of the non-tradable sector in contrast to Corden and Neary's resource movement effect. The second round spending effect is analogous to Corden and Neary's spending effect as the factor returns arising in the sector whose demand has risen (resources in Corden and Neary, services in the tourism case) will be partly spent on non-tradables.

Which adjustment channels could work in the opposite direction of this Dutch disease effect of tourism? The first and simplest does not require stepping outside the neoclassical framework. If there was costless factor mobility across locations for all factors of production, then one would not expect any effect of tourism demand on sectoral output shares, as perfect factor mobility would fully offset any tourism induced relative price movements. In a more realistic version of this argument, when at least some factors such as labour and capital are able to move across locations subject to non-prohibitive reallocation costs, we would expect the Dutch disease effect to be weakened, but there is no force at work that would push in the opposite direction.

When stepping outside the neoclassical framework, there can be a number of channels through which an increase in local consumption spending could lead to positive effects on tradable goods production. Tourism development fundamentally increases the local market size relative to other regions. In a framework with transport costs and increasing returns to scale in production for at least some of the tradable goods sectors there would be a straight forward reference to a Linder type home market effect in new trade theory and new economic geography models (Krugman, 1980; Fujita et al., 1999).[3] Alternatively, one could think of potential productivity effects on both the non-tradable and the tradable sectors as a result of the increased scale of local activity. Both of these alternative adjustment channels could work against the demise of tradable goods production.

A more practical consideration is that tourism development requires substantial investments in transportation, sanitary, and other physical infrastructures. In this respect, a location's potential for tourism demand would lead to all kinds of investments into the local economy that could themselves have direct effects on the composition and levels of local production. Empirically, we would estimate these effects as part of the reduced form treatment effect of higher tourism activity, whereas in reference to theory we would lump together the direct effect of higher local tourism consumer spending and the indirect effect of better and more local infrastructure on economic outcomes. I will address this point in the empirical section by observing how the estimated effects of tourism activity are affected by the inclusion of additional infrastructure endowment variables in the regression specification.

[3] For example Zeng and Zhu (2008) propose a model of tourism in an increasing returns to scale setting which yields agglomeration effects of local tourism spending.

3 Data

3.1 Municipality Level Economic Variables

The municipality level economic outcomes of central interest are local population and local gross domestic product broken up into industrial production, agricultural production, services production, and hotel and lodging sector activity. I obtain data on municipality population from three Mexican population census available on the IPUMS International online database for the years 1960, 2000, and 2005. Non-agricultural municipality level production data stem from Mexican establishment surveys (Censos Economicos) in 1998 and 2003. These surveys are administered by the Instituto Nacional de Estadistica, Geografia e Informatica (INEGI) and cover every non-agricultural establishment in Mexico for urban areas, and a representative subsample of establishments in rural areas.[4] I obtain access to aggregated production data at the municipality for the two national rounds 1998 and 2003.

The municipality level agricultural production data stem from two rounds of two different agricultural production surveys, the Censo Agropecuario and the Censo Ejidal of the years 1998 and the year 2003. In combination, these farm surveys cover both commercial farming as well as traditional farming activities on communal "ejido" estates.[5] Table 1 provides descriptive statistics across different groups of Mexican municipalities.

3.2 Geological and Oceanographic Variables

The first data input to the geographical information system (GIS) is a georeferenced polygon layer displaying the 2005 municipality boundaries of Mexico that I obtain from INEGI. I then obtain a list of the geographical coordinate pairs of the central points of all islands within the Mexican maritime territory from the national Mexican geological survey. I import these coordinates as a point layer in addition to the municipality boundary layer which allows me to compute which of the coastal municipalities is within 10 km of an off-shore island location.

I obtain ocean water turbidity scores for 250m square pixels around the Mexican coastline from satellite remote sensing data of the US National Oceanic and Atmospheric Administration's (NOAA) Ocean Color project. I import these scores attached to the centroids of each square parcel in a new point layer in addition to the previous GIS layers. This allows me

[4] Documentation and reference manuals of these data can be found at www.inegi.gob.mx.
[5] Documentation and reference manuals of these data can be found at www.inegi.gob.mx.

to compute the average ocean water turbidity score within a 10 kilometer buffer around the boundary of each coastal municipality. Figure 4 provides a cartographic display of the different layers of the GIS.

Finally, I obtain airport and seaport coordinate locations, as well as georeferenced road and railway network files from the geographical database maintained by INEGI. All of these transport infrastructure layers refer to the year 2000. I exclude non-paved roads from the road network layer, and retired railway lines from the railway layer in the GIS.

4 Empirical Estimation and Results

4.1 Estimation Strategy

The baseline regression specification is:

$$\ln y_m = \alpha + \beta \ln HotelGDP_m + \gamma CoastDummy_m + \epsilon_m \tag{1}$$

$\ln y_m$ is the natural log of an outcome in municipality m, $\ln HotelGDP_m$ is the log of the sum of reported municipality level hotel and lodging sector GDP in the economic censuses of 1998 and 2003, $CoastDummy_m$ is an indicator of whether municipality m is within 10 kilometers of the Mexican shoreline, and ϵ_m is the residual. I run this specification for every outcome before and after including state fixed effects across 32 federal territories in Mexico. To address concerns that the error term is correlated across municipalities that are affected by common policy or market shocks, I cluster standard errors at the state level in all regressions.

The outcome variables of key interest will be log levels of population, log levels of sectoral municipal GDPs, as well as log output shares across sectors. There are many reasons to be concerned about the endogeneity of log hotel activity when explaining these outcomes in the cross-section both before and after including state fixed effects. In particular, hotel activity is likely to be correlated with many omitted variables at the municipality level that also affect local economic conditions, such as ease of transport access, quality of local services, or the historical political and economic importance of a municipality. Equally plausible are concerns about reverse causality such as a link running from the economic development of a city to higher visitor numbers. *A priori*, we would expect both omitted variable as well as reverse causality concerns to create an upward bias on the β point estimate for the outcome variables of central interest: industrial output levels and the industrialization of a municipality measured by the share of industrial output to overall GDP.

To address these concerns, I propose an instrumental variables strategy that exploits geological and oceanographic variation of beach quality across coastal municipalities in Mexico. Figure 4 presents a cartographic illustration of this strategy. More than half of Mexico's hotel activity is located among the 143 coastal municipalities, a fact that reflects the beach oriented nature of both domestic and international tourism activity in Mexico (see Table 1). The instrumental variable strategy is based on the idea that beach quality affects the level of local tourism demand while not being directly or indirectly related to municipality level economic outcomes other than through tourism.

To this end, I construct two variables that reflect two particular amenity attributes that are frequently mentioned in the tourism management literature.[6] The first is related to ocean water turbidity (the opposite of clarity). Water clarity affects a location's suitability for tourism development both through more scenic beaches with lighter blue colors of water, as well as for scuba diving and snorkeling activities off the coast. The second factor is the presence of near off-shore islands that can be reached by swimming or boats. Islands are mainly mentioned in the tourism literature as providing scenic views from the beach and closeby hotel installations, but they also generate a number of touristic activities such as trips to the islands as well as extra beach space.

The first of these variables enters the first stage regressions as the log of the mean turbidity score within 10 kilometers of a municipalities boundary, and the second enters as an identifier variable of whether or not there is an off-shore island present within 10 kilometers of the boundaries of coastal municipalities. The identifying assumption is that island locations and water turbidity are driven by geological and oceanographic processes in a way that is exogenous to other factors affecting economic outcomes among coastal municipalities else than through tourism activity.

A first potential remaining concern could be that island locations and water turbidity are geographically clustered natural phenomena along the coastline so that the inclusion of the coastal municipality identifier is insufficient to capture economic heterogeneity across Mexico in a broader sense. For instance, one might be concerned about spurious correlations to the

[6] See for instance Leatherman (1997) and references listed therein.

distance to the Northern border with the US or to Mexico City. To this end, I run each regression before and after including state fixed effects as the concern about geographical clustering of these natural phenomena appears much less plausible within each of the 32 federal entities in Mexico.

A second potential concern is that for some unobserved reason the presence of off-shore islands and/or the turbidity of the coastal ocean waters are correlated with historical and/or contemporary determinants of local economic structure. For example, water turbidity could be correlated with the presence of mainland rivers flowing into the ocean, and island locations could provide better or worse conditions for fishing and aquacultural sectors. To address such concerns, I also estimate a falsification test of the exclusion restriction. If these natural characteristics of coastal municipalities affect economic outcomes only through tourism activity, then they should not affect municipality differences in a period before Mexican tourism activity had reached a relevant scale. I therefore match contemporary counties depicted in Figure 4 to administrative units included in the Mexican population census of the year 1960 that I obtain from the IPUMS international database.

Subject to the exclusion restriction, the 2nd stage IV estimates identify the local average treatment effect of higher tourism activity on municipality level outcomes. These estimates represent the reduced form effect (in its microeconomic rather than econometric meaning) of increasing tourism activity in the cross-section of municipalities. They should be thought of as total derivatives whereby tourism activity can affect local outcomes both through direct effects that we discussed above (increasing numbers of local consumers, income effects on local population), as well as indirect effects such as improving transport and sanitary infrastructures and the built environment of municipalities.

Tourism development usually requires significant infrastructure upgrades along the transport and other physical dimensions. As tourism is an important source of foreign reserves in many developing countries including Mexico, many of these investments are made in form of federal transfers to locations with high tourism potential. Such side effects of tourism activity do not cause traditional omitted variable concerns as they are themselves a direct function of tourism activity.

They do, however, raise an interesting empirical consideration to try and disentangle the economic effects of increased tourist visitors and spending from the indirect effects of better transport access and other infrastructure that tourism development brings with it. To this end, I construct four municipality level transport infrastructure variables: Two indicators for whether a municipality is endowed with an airport or seaport, as well as the kilometer amount of paved roads and railway lines as of the year 2000. I then run regression on the economic outcome variables before and after including these infrastructure controls. Given that these infrastructure variables are themselves partly determined by tourism activity and the instrumental variables affecting tourism, a caveat of this strategy is to expect diminished explanatory power of the instrumental variables in first stage regressions after including the infrastructure controls.

4.2 Results

Table 2 presents results for first stage regressions of log municipality hotel GDP on the island dummy and the log turbidity score in addition to coastal location identifiers. The inclusion of state fixed effects is indicated at the bottom of the table.[7] Both instruments enter with expected signs (islands are good for tourism, turbidity is bad) and the first stage F-statitistics indicate strong partial R-squares both before and after including state fixed effects for both instruments. As expected, the first stage explanatory power is diminished after including infrastructure variables that are themselves a function of tourism potential as additional controls. In particular, the first stages are too weak after including the full set of four infrastructure variables. I will therefore include the airport and seaport dummies or the road and rail kilometer variables in separate regressions in the 2SLS estimations.[8]

Table 3 presents OLS and IV results on the question of factor migration. Both OLS and IV estimates present a strong positive effect of tourism activity on municipality populations. OLS and IV estimates are very similar in their magnitude. The final column with state fixed effects and both instruments suggests that a 1% increase in hotel and lodging GDP leads to a 0.3% increase in local populations. This represents a sizeable effect on cross-municipality factor immigration as a consequence of local tourism demand, a finding that is supported

[7] Notice that while many Mexican municipalities report zero hotel activity, close to the entirety of coastal municipalities report non-zero values. Given that the identifying variation in IV estimations is entirely across coastal municipalities, selection due to taking logs on the left hand side should not be a concern.

[8] The key objective in this exercise is to see how the estimated tourism effects are affected after including separate sets of infrastructure controls subject to the condition that the first stage is still acceptably strong. As discussed above, the issue is not that infrastructure controls are omitted variables that are correlated to the instruments, but to try and disentangle the effect of tourism on local outcomes net of the infrastructure effects that tourism induces.

by the observed differences in average population numbers among touristic and non-touristic municipalities in Table 1.

This finding has important implications for the interpretation of the subsequent results on the levels as opposed to output shares of municipalities. If tourism demand leads to factor immigration, then its effect on the levels of output are conceptually only interesting in relative terms. More factors of production should lead to higher output in any sector of the economy. In a setting with partial factor mobility (but not perfect/costless mobility of all factors), the interesting question is to what extend the local tourism demand shock leads to differences in sectoral specialization, not overall output levels.

Table 4 and Table 5 present OLS and IV regression results for the log levels of sectoral municipality economic output, and Table 6 and Table 7 present OLS and IV results for log output shares. The OLS results on log output levels are positive and statistically significant across agriculture, industry, and services before and after including the additional infrastructure endowment variables. The OLS results on log output shares are strongly positive and significant across specifications for the degree of municipality industrialization and the share of services output. In particular, the effect on industrial output appears to be more pronounced than on services output as reflected in the final row with significant positive effects on the ratio of industrial output over services. The opposite result is found for the share of agricultural output in the local economy where the effect of tourism is negative and significant throughout. These OLS regression results would thus confirm the positive relationship between tourism activity and industrialization that we find in the cross-country relationship depicted in Figure 3 and discussed above.

The IV results on both log output levels and output shares lead to a very different set of results. The estimated effects on log levels of industrial output are all statistically insignificant and point estimates range from positive to negative. For agricultural output levels, most point estimates are positive and some of them marginally statistically significant. Point estimates on levels on services output are all positive and most of them are statistically significant. The absence of significant growth effects on industrial production is in itself a statement about tourism's effect on industrialization when viewed in the light of the previous strong evidence of factor immigration as a consequence of tourism activity.

This argument is confirmed by the IV results on log output shares. All point estimates on industrial shares in total production or the ratio of industrial to services GDP are negative. Almost all specifications yield statistically significant coefficients for the industry/services output ratio, while results are negative and become statistically significant for the industrial output share after including infrastructure controls.

These results provide empirical evidence of the Dutch disease effect of tourism activity on industrial production, and the fact that the point estimates increase in absolute size after including infrastructure controls points to the fact that tourism induced infrastructure investments tend to foster industrial activity, while the effect of local tourism demand on industrial production net of those side effects is negative. The opposite of these results is found for the share of services in total economic activity, while no statistically significant effect is found on the share of agriculture in total economic activity. These results point to the fact that industrial rather than agricultural activities are negatively affected among the tradable production sectors.

Finally, Table 8 presents the results of the placebo falsification test on log municipality populations. The reported results are for regressing log levels of municipality population directly on the instrumental variables. While the earlier result on factor immigration as a consequence of tourism potential is confirmed on the subsample of municipalities that could be matched to the 1960 Mexican population census, no such positive effects are present on populations reported in the year 1960. This result provides a reassuring robustness check on the exclusion restriction and supports the causal interpretation of the instrumental variable results presented on factor immigration as well as sectoral output levels and municipality production shares.

5 Conclusion

Despite the impressive growth and scale of tourism as a form of economic integration, our existing empirical evidence of its economic implications is rather limited. This paper proposes a novel empirical strategy to contribute to our understanding of this question. In particular, I exploit geological and oceanographic variation in beach quality along the Mexican coastline to instrument for local tourism activity across Mexican municipalities.

While both cross-country as well as Mexican municipality data suggest a strongly significant positive correlation between industrialization and tourism activity, I find that this relationship is reversed in a statistically and economically significant way in instrumental variable estimations.

In particular, the tourism induced expansion of the services sector appears to be at the expense of industrial production rather than the agricultural sector. These findings become stronger as one starts conditioning on levels of local infrastructure endowments that are themselves affected by tourism development. In combination with the presented empirical evidence of strong factor immigration responses to local tourism activity, these results should be interpreted as a lower bound to a cross-country setting where factors are less mobile to offset tourism induced relative price changes.

It is important to conclude this study emphasizing that no statements or conclusions have been made about the national or local welfare consequences or welfare distributional consequences of tourism development. As is generally the case in the neoclassical literature on the Dutch disease, the observed effect is theoretically consistent with a range of welfare and welfare distributional outcomes. Building on the empirical methodology outlined in this paper, I plan to address these questions in ongoing work in progress on this chapter.

References

Arezki, R., Cherif, R., & Piotrowski, J. 2009. Tourism Specialization and Economic Development: Evidence from the UNESCO World Heritage List. *IMF Staff Working Paper*, **09/176**.

Copeland, B.R. 1991. Tourism, welfare and de-industrialization in a small open economy. *Economica*, 515–529.

Corden, W.M., & Neary, J.P. 1982. Booming sector and de-industrialisation in a small open economy. *The economic journal*, **92**(368), 825–848.

Eilat, Y., & Einav, L. 2004. Determinants of international tourism: a three-dimensional panel data analysis. *Applied Economics*, **36**(12), 1315–1327.

Fujita, M., Krugman, P.R., & Venables. 1999. *The spatial economy: cities, regions and international trade*. Wiley.

Krugman, P. 1980. Scale economies, product differentiation, and the pattern of trade. *The American Economic Review*, **70**(5), 950–959.

Leatherman, S.P. 1997. Beach rating: a methodological approach. *Journal of Coastal Research*, 253–258.

Sequeira, T.N., & Nunes, P.M. 2008. Does tourism influence economic growth? A dynamic panel data approach. *Applied Economics*, **40**(18), 2431–2441.

Zeng, D.A.O.Z.H.I., & Zhu, X. 2008. Tourism and Industrial Agglomeration. *Japanese Economic Review*.

Figures and Tables

Figures

Figure 1: Tourism Exports Relative to Trade and Investment Flows 1995-2005

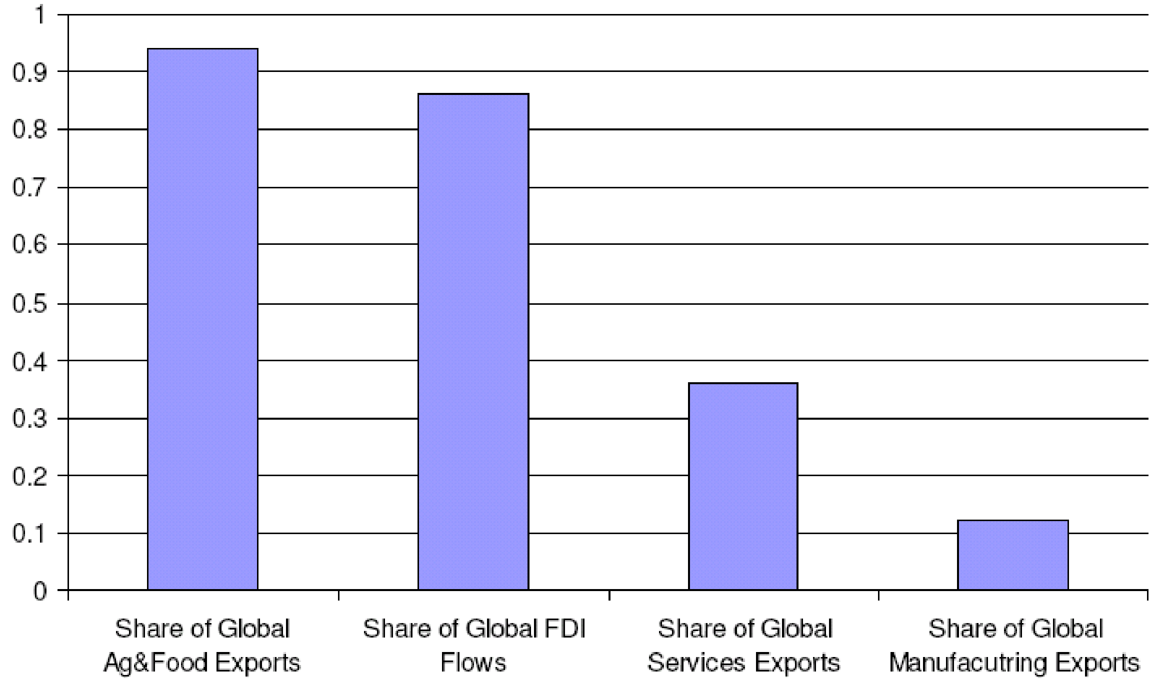

Shares (1=100%) refer to annual global tourism exports over annual global respective trade or investment flows. Foreign direct investment flows are measured as inflows net of disinvestments across countries. The data are from UNCTAD.

Figure 2: Tourism Activity and Manufacturing Productions - Cross Country Scatter Plots

A) Manufacturing Production Shares

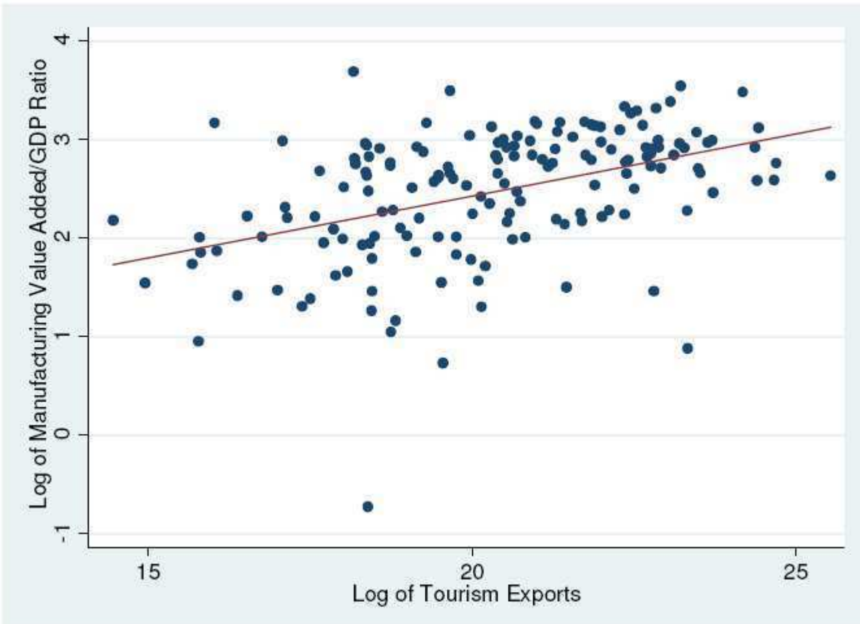

Each point represents one country observation. All data are from the World Bank Development Indicators 2011 database.

B) Manufacturing Export Shares

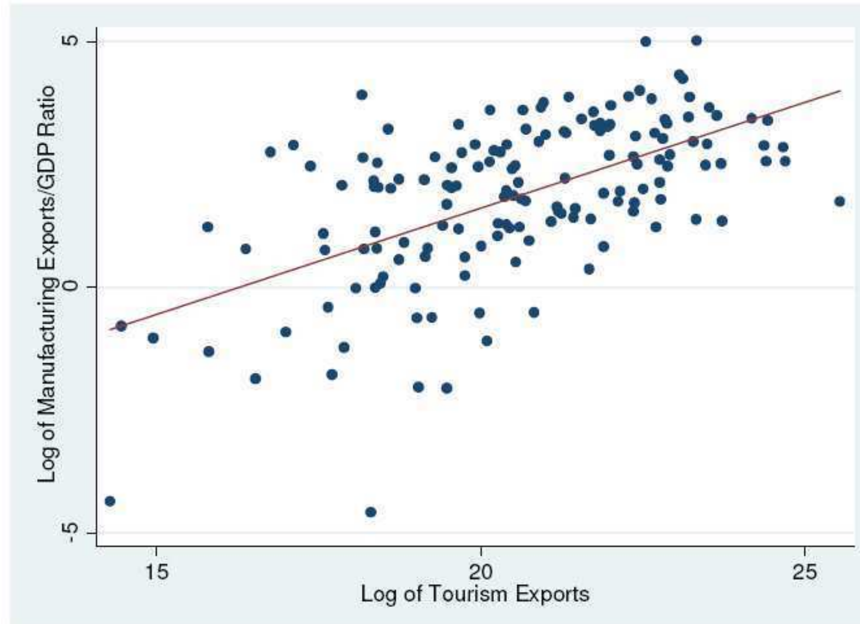

Each point represents one country observation. All data are from the World Bank Development Indicators 2011 database.

Figure 3: Tourism in Mexico

A) Tourism's Share in Mexican GDP

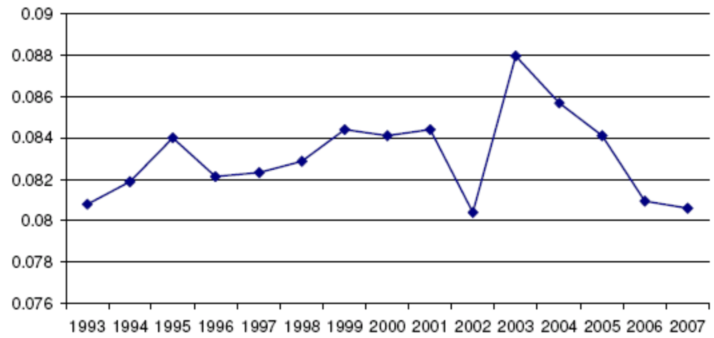

The y-axis depicts the share of tourism related Mexican GDP over total GDP (where 1=100%). Data are from Mexican National Accounts Statistics provided by INEGI.

B) Domestic and International Tourist Receipts

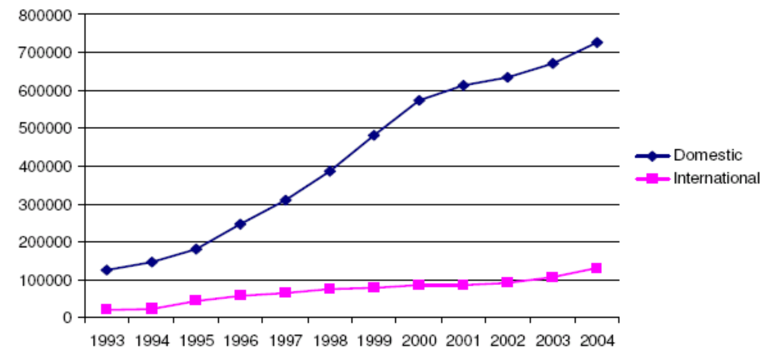

The y-axis depicts Mexican tourism receipts in Millions of current Pesos. Domestic indicates tourism activity of Mexican nationals within Mexico. International refers to cross-border tourism activity in Mexico. Data are from the Mexican Tourism Authority SECTUR.

C) Millions of International Visitors To Mexico

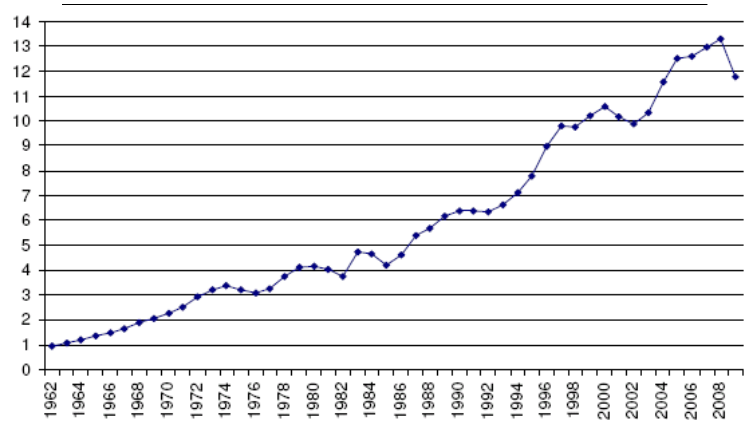

The y-axis depicts Millions of international tourist visitors in Mexico. Data are from the historical statistical archive of the Mexican Central Bank.

Figure 4: Beach Quality and Tourism in Mexico

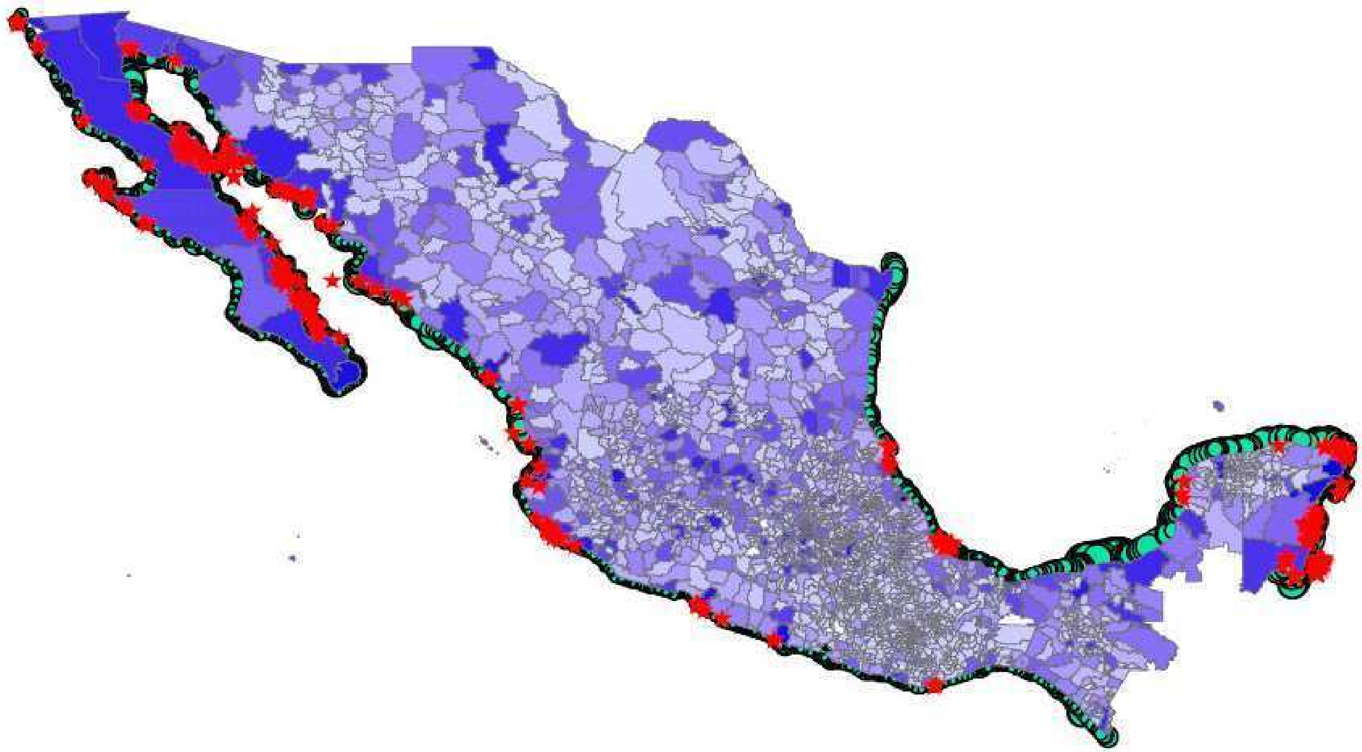

Graduated polygon colours indicate the sum of municipality hotel and lodging sector output for the years 1998 and 2003. Red stars indicate off-shore island locations within 10 km of the Mexican shoreline. The graduated size of the green point layer around the Mexican coastline indicates average ocean water turbidity of 250m square parcels within 10 km of the shore as measured by satellite remote sensors of the US National Oceanic and Atmospheric Administration's (NOAA) Ocean Color project.

Tables

Table 1: Descriptive Statistics of Mexican Municipalities

	Number of Municipalities	Population 2000 Averages	Shares of National Hotel & Lodging GDP	GDP per Capita Averages	Manufacturing GDP Share Averages	Agriculture GDP Share Averages	Services GDP Share Averages
All Municipalities	2410	40297.9	1	40893.51	0.214	0.486	0.301
Coastal Municipalities	143	99948.06	0.543	80515.21	0.227	0.322	0.451
Top 10% of Local Tourism Activity	241	233173.1	0.979	127668.2	0.374	0.51	0.115

Table 2: First Stage Results - Beach Quality and Local Tourism Demand

Dependent Variable: ln(HotelGDP)	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
Island Dummy	2.310***	1.677***	1.368***	0.951**	1.259***	1.061**	0.723	0.575								
	(0.432)	(0.480)	(0.387)	(0.465)	(0.432)	(0.517)	(0.442)	(0.534)								
ln(Ocean Turbidity)									-0.801***	-0.892***	-0.645***	-0.801***	-0.608***	-0.610**	-0.507**	-0.561**
									(0.229)	(0.212)	(0.212)	(0.242)	(0.194)	(0.231)	(0.220)	(0.262)
Coastal Dummy	1.668***	1.876***	1.237***	1.697***	1.099***	1.382***	1.132***	1.469***	-2.955*	-3.689**	-2.751*	-3.554**	-2.619**	-2.445	-2.135	-2.232
	(0.332)	(0.347)	(0.409)	(0.397)	(0.308)	(0.309)	(0.399)	(0.390)	(1.500)	(1.460)	(1.351)	(1.620)	(1.255)	(1.453)	(1.450)	(1.675)
Airport Dummy			3.603***	3.343***			2.590***	2.475***			3.676***	3.370***			2.601***	2.502***
			(0.295)	(0.334)			(0.254)	(0.257)			(0.287)	(0.321)			(0.254)	(0.251)
Seaport Dummy			0.258	-0.144			-0.428	-0.480			0.366	0.00878			-0.329	-0.358
			(0.436)	(0.499)			(0.481)	(0.513)			(0.443)	(0.437)			(0.441)	(0.456)
Paved Road km					0.00930***	0.00920***	0.00816***	0.00812***					0.00931***	0.00903***	0.00801***	0.00787***
					(0.00148)	(0.00183)	(0.00140)	(0.00166)					(0.00159)	(0.00192)	(0.00149)	(0.00174)
Railway km					0.00813***	0.0110***	0.00642***	0.00903***					0.00856***	0.0112***	0.00683***	0.00911***
					(0.00160)	(0.00216)	(0.00160)	(0.00217)					(0.00163)	(0.00215)	(0.00168)	(0.00218)
Constant	6.866***	7.975***	6.740***	7.497***	5.979***	6.388***	6.013***	6.243***	6.866***	7.975***	6.737***	7.493***	5.969***	6.406***	6.014***	6.270***
	(0.172)	(0)	(0.163)	(0.0477)	(0.183)	(0.232)	(0.183)	(0.210)	(0.172)	(0)	(0.164)	(0.0459)	(0.189)	(0.239)	(0.188)	(0.221)
State FX	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
Obs	1292	1292	1292	1292	1292	1292	1292	1292	1292	1292	1292	1292	1292	1292	1292	1292
R^2	0.103	0.215	0.188	0.285	0.240	0.327	0.280	0.363	0.099	0.220	0.192	0.293	0.243	0.329	0.284	0.367
1st Stage F-Stat	**28.57**	**12.20**	**12.52**	**4.17**	**8.50**	**4.21**	**2.68**	**1.16**	**12.26**	**17.62**	**9.30**	**10.90**	**9.81**	**6.94**	**5.32**	**4.57**

The dependent variable ln(HotelGDP) is the log of the sum of 1998 and 2003 municipality level hotel and lodging sector GDP. The island dummy is an indicator of whether the boundary of a municipality is within 10 km of an off-shore island. ln(Ocean Turbidity) is the log of the average water turbidity score within 10 km of the municipality boundary. The coastal dummy indicates municipality location within 10 km of the Mexican coastline. Airport and seaport dummies indicate respective infrastructure presences within a municipality. Paved road and railway km variables measure kilometres of respective infrastructures within a municipality as of the year 2000. Standard errors are clustered at the state level and stated in parentheses below point estimates. ***1%, **5%, and *10% significance levels.

Table 3: OLS and IV Results on Factor Migration

Dependent Variable:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ln(Population00&05)	OLS	OLS	Island IV	Island IV	Turbidity IV	Turbidity IV	Both IVs	Both IVs
ln(HotelGDP)	0.345***	0.322***	0.280***	0.346***	0.244	0.313**	0.265**	0.325***
	(0.0209)	(0.0180)	(0.0756)	(0.126)	(0.185)	(0.130)	(0.109)	(0.111)
Constant	8.445***	9.171***	8.891***	8.986***	9.141***	9.245***	8.996***	9.154***
	(0.198)	(0.143)	(0.527)	(1.002)	(1.263)	(1.034)	(0.748)	(0.887)
State Fixed Effects	No	Yes	No	Yes	No	Yes	No	Yes
Obs	1292	1292	1292	1292	1292	1292	1292	1292
R^2	0.504	0.631	0.487	0.630	0.463	0.631	0.479	0.631
1^{st} Stage F-Statistic			28.57	12.20	12.26	17.62	17.65	16.17

All regressions include an indicator of whether a municipality is within 10 km distance to the coast. The dependent variable is the log of the sum of 2000 and 2005 municipality populations. ln(HotelGDP) is the log of the sum of 1998 and 2003 municipality level hotel and lodging sector GDP. Island IV indicates 2^{nd} stage IV results after instrumenting for log hotel activity with whether or not a municipality is within 10 km of an off-shore island. Turbidity IV indicates 2^{nd} stage IV results after instrumenting for log hotel activity with the log of the average water turbidity score within 10 km of the municipality boundary. Both IVs indicates 2^{nd} stage IV results using both instruments. Standard errors are clustered at the state level and stated in parentheses below point estimates. ***1%, **5%, and *10% significance levels.

Table 4: OLS Results on Output Levels

Dependent Variable		(1)	(2)	(3)	(4)	(5)	(6)
	ln(HotelGDP)	0.760***	0.691***	0.733***	0.661***	0.672***	0.594***
		(0.0401)	(0.0415)	(0.0413)	(0.0409)	(0.0450)	(0.0414)
ln(Industry GDP)	State FX	No	Yes	No	Yes	No	Yes
	Obs	1290	1290	1290	1290	1290	1290
	R^2	0.466	0.540	0.471	0.547	0.502	0.582
	ln(HotelGDP)	0.151***	0.124***	0.144***	0.120***	0.0363	0.0139
		(0.0454)	(0.0258)	(0.0480)	(0.0252)	(0.0377)	(0.0224)
ln(Agri GDP)	State FX	No	Yes	No	Yes	No	Yes
	Obs	1279	1279	1279	1279	1279	1279
	R^2	0.077	0.415	0.077	0.415	0.214	0.550
	ln(HotelGDP)	0.697***	0.646***	0.676***	0.622***	0.646***	0.582***
		(0.0337)	(0.0269)	(0.0371)	(0.0273)	(0.0432)	(0.0298)
ln(Services GDP)	State FX	No	Yes	No	Yes	No	Yes
	Obs	1292	1292	1292	1292	1292	1292
	R^2	0.675	0.727	0.680	0.734	0.692	0.754

All regressions include an indicator of whether a municipality is within 10 km distance to the coast. ln(Industry GDP) is the log of the sum of 1998 and 2003 municiaplity industial GDP. ln(Agri GDP) is the log of the sum of 1998 and 2003 municiaplity agricultural GDP. ln(Services GDP) is the log of the sum of 1998 and 2003 municiaplity services GDP. ln(HotelGDP) is the log of the sum of 1998 and 2003 municipality level hotel and lodging sector GDP. The first two columns include no further controls. Columns 3&4 include airport and seaport location dummies. Columns 5&6 include municipality paved road km and railway track km as of the year 2000. Standard errors are clustered at the state level and stated in parentheses below point estimates. ***1%, **5%, and *10% significance levels.

Table 5: Instrumental Variable Results on Output Levels

Dependent Variable		Island IV (1)	(2)	(3)	(4)	(5)	(6)	Ocean Turbidity IV (1)	(2)	(3)	(4)	(5)	(6)	Both IVs (1)	(2)	(3)	(4)	(5)	(6)
ln(Industry GDP)	ln(Hotel GDP)	0.285	0.194	-0.233	-0.377	-0.455	-0.415	0.0231	0.0948	-0.205	0.0291	-0.280	-0.409	0.175	0.130	-0.215	-0.0250	-0.339	-0.411
		(0.214)	(0.330)	(0.358)	(0.600)	(0.367)	(0.527)	(0.304)	(0.296)	(0.396)	(0.367)	(0.350)	(0.457)	(0.194)	(0.261)	(0.297)	(0.359)	(0.294)	(0.374)
	State FX	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
	Obs	1290	1290	1290	1290	1290	1290	1290	1290	1290	1290	1290	1290	1290	1290	1290	1290	1290	1290
	1st Stage F-Stat	28.57	12.17	12.53	4.15	8.51	4.18	12.26	17.59	9.30	10.87	9.80	6.93	17.65	16.14	9.99	8.56	7.40	4.90
ln(Agri GDP)	ln(Hotel GDP)	0.226	0.163	0.242	0.179	-0.184	-0.152	0.131	0.342*	0.111	0.374*	-0.0772	0.135	0.184	0.285*	0.150	0.356*	-0.112	0.0495
		(0.197)	(0.206)	(0.313)	(0.380)	(0.363)	(0.272)	(0.245)	(0.175)	(0.294)	(0.214)	(0.316)	(0.202)	(0.177)	(0.156)	(0.248)	(0.205)	(0.303)	(0.180)
	State FX	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
	Obs	1279	1279	1279	1279	1279	1279	1279	1279	1279	1279	1279	1279	1279	1279	1279	1279	1279	1279
	1st Stage F-Stat	29.22	11.70	11.90	3.29	8.63	4.13	14.22	20.84	11.66	13.39	11.19	7.70	19.61	17.68	11.58	9.54	8.05	5.07
ln(Services GDP)	ln(Hotel GDP)	0.600***	0.616***	0.399**	0.447	0.286	0.380	0.341**	0.370***	0.213	0.332*	0.139	0.0701	0.492***	0.457***	0.278**	0.348**	0.189	0.166
		(0.112)	(0.200)	(0.161)	(0.285)	(0.186)	(0.248)	(0.157)	(0.131)	(0.210)	(0.178)	(0.183)	(0.227)	(0.101)	(0.132)	(0.140)	(0.162)	(0.140)	(0.165)
	State FX	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
	Obs	1292	1292	1292	1292	1292	1292	1292	1292	1292	1292	1292	1292	1292	1292	1292	1292	1292	1292
	1st Stage F-Stat	28.57	12.20	12.52	4.17	8.50	4.21	12.26	17.62	9.30	10.90	9.81	6.94	17.65	16.17	9.98	8.61	7.41	4.93

All regressions include an indicator of whether a municipality is within 10 km distance to the coast. ln(Industry GDP) is the log of the sum of 1998 and 2003 municiaplity industial GDP. ln(Agri GDP) is the log of the sum of 1998 and 2003 municiaplity agricultural GDP. ln(Services GDP) is the log of the sum of 1998 and 2003 municiaplity services GDP. ln(HotelGDP) is the log of the sum of 1998 and 2003 municipality level hotel and lodging sector GDP. Island IV indicates 2nd stage IV results after instrumenting for log hotel activity with whether or not a municipality is within 10 km of an off-shore island. Turbidity IV indicates 2nd stage IV results after instrumenting for log hotel activity with the log of the average water turbidity score within 10 km of the municipality boundary. Both IVs indicates 2nd stage IV results using both instruments. The first two columns include no further controls. Columns 3&4 include airport and seaport location dummies. Columns 5&6 include municipality paved road km and railway track km as of the year 2000. Standard errors are clustered at the state level and stated in parentheses below point estimates. ***1%, **5%, and *10% significance levels.

Table 6: OLS Results on Output Shares

Dependent Variable		(1)	(2)	(3)	(4)	(5)	(6)
ln(Industry/Total GDP)	ln(HotelGDP)	0.184***	0.176***	0.183***	0.175***	0.161***	0.153***
		(0.0185)	(0.0202)	(0.0190)	(0.0209)	(0.0182)	(0.0196)
	State FX	No	Yes	No	Yes	No	Yes
	Obs	1290	1290	1290	1290	1290	1290
	R^2	0.139	0.229	0.140	0.231	0.163	0.248
ln(Agri/Total GDP)	ln(HotelGDP)	-0.413***	-0.386***	-0.391***	-0.359***	-0.453***	-0.417***
		(0.0482)	(0.0310)	(0.0529)	(0.0315)	(0.0565)	(0.0376)
	State FX	No	Yes	No	Yes	No	Yes
	Obs	1279	1279	1279	1279	1279	1279
	R^2	0.289	0.489	0.295	0.499	0.306	0.505
ln(Services/Total GDP)	ln(HotelGDP)	0.120***	0.131***	0.125***	0.136***	0.133***	0.140***
		(0.0188)	(0.0130)	(0.0204)	(0.0145)	(0.0164)	(0.0133)
	State FX	No	Yes	No	Yes	No	Yes
	Obs	1292	1292	1292	1292	1292	1292
	R^2	0.146	0.277	0.148	0.279	0.158	0.283
ln(Industry/Services GDP)	ln(HotelGDP)	0.0643***	0.0456**	0.0584**	0.0390*	0.0274	0.0131
		(0.0234)	(0.0201)	(0.0242)	(0.0204)	(0.0208)	(0.0191)
	State FX	No	Yes	No	Yes	No	Yes
	Obs	1290	1290	1290	1290	1290	1290
	R^2	0.021	0.136	0.022	0.138	0.056	0.160

All regressions include an indicator of whether a municipality is within 10 km distance to the coast. The listed dependent variables are municipality level log output shares of summed sectoral GDPs in 1998 and 2003. ln(HotelGDP) is the log of the sum of 1998 and 2003 municipality level hotel and lodging sector GDP. The first two columns include no further controls. Columns 3&4 include airport and seaport location dummies. Columns 5&6 include municipality paved road km and railway track km as of the year 2000. Standard errors are clustered at the state level and stated in parentheses below point estimates. ***1%, **5%, and *10% significance levels.

Table 7: Instrumental Variable Results on Output Shares

Dependent Variable		Island IV						Ocean Turbidity IV						Both IVs					
		(1)	(2)	(3)	(4)	(5)	(6)	(1)	(2)	(3)	(4)	(5)	(6)	(1)	(2)	(3)	(4)	(5)	(6)
ln(Industry/ Total GDP)	ln(Hotel GDP)	-0.0827	-0.119	-0.286*	-0.337	-0.372***	-0.364*	-0.167	-0.160	-0.246*	-0.181	-0.261*	-0.357	-0.118	-0.145	-0.260**	-0.202	-0.298**	-0.359**
		(0.106)	(0.123)	(0.168)	(0.242)	(0.138)	(0.219)	(0.115)	(0.138)	(0.133)	(0.164)	(0.148)	(0.221)	(0.0878)	(0.106)	(0.104)	(0.154)	(0.118)	(0.173)
	State FX	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
	Obs	1290	1290	1290	1290	1290	1290	1290	1290	1290	1290	1290	1290	1290	1290	1290	1290	1290	1290
	1st Stage F-Stat	28.57	12.17	12.53	4.15	8.51	4.18	12.26	17.59	9.30	10.87	9.80	6.93	17.65	16.14	9.99	8.56	7.40	4.90
ln(Agri/ Total GDP)	ln(Hotel GDP)	-0.160	-0.159	0.163	0.246	-0.164	-0.160	-0.0744	0.0685	0.0348	0.135	-0.0790	0.171	-0.122	-0.00306	0.0731	0.145	-0.107	0.171
		(0.184)	(0.223)	(0.288)	(0.504)	(0.341)	(0.347)	(0.201)	(0.180)	(0.274)	(0.285)	(0.283)	(0.317)	(0.161)	(0.139)	(0.235)	(0.264)	(0.260)	(0.317)
	State FX	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
	Obs	1279	1279	1279	1279	1279	1279	1279	1279	1279	1279	1279	1279	1279	1279	1279	1279	1279	1279
	1st Stage F-Stat	29.22	11.70	11.90	3.29	8.63	4.13	14.22	20.84	11.66	13.39	11.19	7.70	19.61	17.68	11.58	9.54	8.05	5.07
ln(Services/ Total GDP)	ln(Hotel GDP)	0.233***	0.303***	0.346**	0.485**	0.372**	0.433**	0.151*	0.112*	0.173	0.118	0.159	0.118	0.199***	0.179***	0.233**	0.167*	0.231*	0.215**
		(0.0679)	(0.0952)	(0.138)	(0.238)	(0.160)	(0.199)	(0.0898)	(0.0670)	(0.115)	(0.0779)	(0.112)	(0.0964)	(0.0703)	(0.0661)	(0.116)	(0.0866)	(0.119)	(0.0986)
	State FX	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
	Obs	1292	1292	1292	1292	1292	1292	1292	1292	1292	1292	1292	1292	1292	1292	1292	1292	1292	1292
	1st Stage F-Stat	28.57	12.20	12.52	4.17	8.50	4.21	12.26	17.62	9.30	10.90	9.81	6.94	17.65	16.17	9.98	8.61	7.41	4.93
ln(Industry/ Services GDP)	ln(Hotel GDP)	-0.315**	-0.423**	-0.633**	-0.826*	-0.744***	-0.799**	-0.318*	-0.271	-0.419*	-0.298	-0.420**	-0.474*	-0.316***	-0.325**	-0.494**	-0.369	-0.529***	-0.574**
		(0.134)	(0.174)	(0.256)	(0.432)	(0.246)	(0.376)	(0.176)	(0.190)	(0.224)	(0.222)	(0.208)	(0.279)	(0.117)	(0.160)	(0.193)	(0.230)	(0.190)	(0.253)
	State FX	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
	Obs	1290	1290	1290	1290	1290	1290	1290	1290	1290	1290	1290	1290	1290	1290	1290	1290	1290	1290
	1st Stage F-Stat	28.57	12.17	12.53	4.15	8.51	4.18	12.26	17.59	9.30	10.87	9.80	6.93	17.65	16.14	9.99	8.56	7.40	4.90

All regressions include an indicator of whether a municipality is within 10 km distance to the coast. The listed dependent variables are municipality level log output shares of summed sectoral GDPs in 1998 and 2003. ln(HotelGDP) is the log of the sum of 1998 and 2003 municipality level hotel and lodging sector GDP. Island IV indicates 2nd stage IV results after instrumenting for log hotel activity with whether or not a municipality is within 10 km of an off-shore island. Turbidity IV indicates 2nd stage IV results after instrumenting for log hotel activity with the log of the average water turbidity score within 10 km of the municipality boundary. Both IVs indicates 2nd stage IV results using both instruments. The first two columns include no further controls. Columns 3&4 include airport and seaport location dummies. Columns 5&6 include municipality paved road km and railway track km as of the year 2000. Standard errors are clustered at the state level and stated in parentheses below point estimates. ***1%, **5%, and *10% significance levels.

Table 8: Placebo Falsification Test

Dependent Variable		(1)	(2)	(3)	(4)	(5)	(6)
	Island Dummy	0.671***	0.611**			0.621**	0.536*
		(0.252)	(0.264)			(0.258)	(0.291)
ln(Population 2000&2005)	ln(Ocean Turbidity)			-0.146	-0.218	-0.0672	-0.168
				(0.181)	(0.175)	(0.187)	(0.192)
	State FX	No	Yes	No	Yes	No	Yes
	Obs	1220	1220	1220	1220	1220	1220
	R^2	0.034	0.267	0.030	0.265	0.035	0.269
	Island Dummy	0.308	-0.0878			0.263	-0.141
		(0.331)	(0.212)			(0.365)	(0.207)
ln(Population 1960)	ln(Ocean Turbidity)			-0.0931	-0.105	-0.0596	-0.119
				(0.181)	(0.158)	(0.192)	(0.156)
	State FX	No	Yes	No	Yes	No	Yes
	Obs	1220	1220	1220	1220	1220	1220
	R^2	0.002	0.257	0.001	0.258	0.002	0.258

All regressions include an indicator of whether a municipality is within 10 km distance to the coast. The dependent variables are the log of the sum of 2000 and 2005 municipality populations or the log of 1960 municipality population. The island dummy is an indicator of whether the boundary of a municipality is within 10 km of an off-shore island. ln(Ocean Turbidity) is the log of the average water turbidity score within 10 km of the municipality boundary. Standard errors are clustered at the state level and stated in parentheses below point estimates. ***1%, **5%, and *10% significance levels.