

Chapter 5

Estimation of parameters in an economic geography model

5.1 Introduction

In the previous chapters, we have seen how models of economic geography explain the agglomeration of economic activities. The explanation involves increasing returns to scale at the firm- and industrial level. The models show that if trade is impeded by costs of transport, it is possible, depending on the parameters, that a concentration of activity arises as a result of economic mechanisms.

But while such concentrations of economic activity surely exist in practice, as many real-world examples will testify, it may be difficult to determine if the purported mechanisms of economic geography theory have anything to do with them. If we are to show that this is the case, we need to distinguish this cause for agglomeration from other possible explanations, such as natural geographical circumstances. Ellison and Glaeser (1999) claim that natural advantages, such as the presence of a natural harbor or a particular climate, can be used to explain “at least half of observed geographic concentration” (p. 316). Setting these causes, which are often called ‘first nature,’ apart from economic incentives to agglomerate (or, ‘second nature’) is a methodological challenge.

In this chapter, we will survey the empirical literature on economic geography. Furthermore, we will attempt to estimate the parameters of the model of the previous chapter, using data on shipments between American states. The payoff of this exercise will be twofold: not only will we find out about the relevance of our theoretical exercises, but the attempt to fit the theory of economic geography to real-world data will also give us a calibrated version of our model.

Firstly, the estimation of model parameters serves the purpose of validating our model as an explanation of events in the real world. The model’s

performance, its goodness of fit, will tell us if we should not stick with a simpler alternative. As an added advantage, the estimated parameters can be interpreted as variables of economic interest in their own right. As we have seen in earlier chapters, parameters such as the elasticity of substitution between differentiated goods and the costs of transport, figure in models of economic geography. An estimated value of these parameters would be interesting, even outside the context of a model.

Secondly, an economic geography model calibrated to real-world data would be a useful tool for many purposes. Models from this class, because of their explicit microeconomic foundations, lend themselves readily for the evaluation of counterfactuals. As such, they can be a tool to evaluate proposed economic policies in the light of their geographical repercussions. For instance, questions as to how the construction of a new road would affect the regions around it can be addressed using a calibrated economic geography model. Not only can we compute the new equilibrium after a change has taken place, but we can also assess the resulting change in welfare. There exists a strand of literature that employs these models for precisely this purpose, which includes Venables and Gasiorek (1996) and Bröcker (1995, 1999).

Another interesting phenomenon that could be analyzed with a calibrated model is the spatial distribution of the effects of a local shock in demand. It is assumed that such a shock will have repercussions outside the originating region, but exactly how far is unknown. We look at the analysis of this question by Hanson (1999), who analyzes the spread of a shock in demand from the center of the United States. We will evaluate a similar shock with our calibrated model below.

After the survey of the empirical literature in section 5.2, we use the rest of the chapter for an attempt to estimate the parameters of an economic geography model. The model is drawn up along the lines of chapter 4, using vertical relations between firms as the agglomerating force. In section 5.3, we discuss the model that will be used in this estimation. As before, the model explains flows of trade between different regions by relating them to the number of firms and the number of consumers in each region. The number of firms in the exporting region has a positive effect on the level of trade, because of the Armington assumption: each extra firm offers a unique product which is desired in all regions. Firms and consumers in the importing region both exercise demand, which increases when there are more of either group. Because the model also involves transport costs that increase with distance, it predicts large trade flows between states that are close. Taken together, this specification is reminiscent of the gravity trade model (by, for instance, Tinbergen 1962). Named for its mathematical analogy to that well-known force of attraction, the model predicts that trade is proportional to the size of the trading economies, and is inversely proportional to their distance. While its origins are certainly agnostic, it has been

derived in models of monopolistic competition and Deardorff (1995) shows that it can also be derived in a neoclassical framework.

But while we find that our model resembles the gravity model of trade in its specification of the trade equation, it goes beyond gravity when we take into account its general equilibrium nature. Where the gravity model only predicts the flow of trade, given the economic size of regions and their distance, the economic geography model explains trade, prices *and* economic size in one consistent framework. We will use a gravity-like equation in the first part of our estimations, but will move beyond the results to verify the coherence between wages, prices and market access for the different regions.

Our data is introduced in section 5.4. It consists of the 1997 US Commodity Flow Survey, which contains measurements of the bilateral shipments between the 51 American states. The Commodity Flow Survey is a relatively underexploited dataset. It is based on a survey conducted by the Bureau of Transportation which asks for origin and destination addresses, value and weight. The 1993 version of the survey was used by Wolf (1997) to estimate gravity equations and examine the home market effect in US states.¹

Estimation is done in Section 5.5. We proceed in two different ways. Section 5.5.1 discusses an attempt to estimate the parameters of the model using the methods of Redding and Venables (2001); section 5.5.2 applies a different method to estimate the same model from the same dataset. The latter method leaves us with a calibrated model, that can be used for the analysis of counterfactuals. We use it to compute the spatial effect of a shock in wages in the United States, and experiment with the lowering of trade costs between regions. Finally, Section 5.7 concludes.

5.2 Estimation in the literature

Overman et al. (2001) survey the empirical literature concerning economic geography. Their survey is built around a canonical theoretical model, similar to the models developed in previous chapters. Different aspects of the model are used in various exercises, all which serve the purpose of testing the model's predictions.

The analyses that are surveyed can be divided into three areas: firstly, measurements of the importance of transport costs are introduced, in order to show that distance matters for economic activity. Transport costs are defined broadly, including not only the price of actual transport but also the time in transit and costs of information. Besides measuring such costs directly, by comparing *fob* and *cif* prices, an indication of the role of distance

¹The home market effect is a result in trade theory that follows from monopolistic competition. It was discussed in section 2.B.

is given by gravity estimations. The fact that a region's trade with similar partners declines with their distance shows that transport costs indeed form a burden.

The second area of research focuses on the fact that the distribution of economic activity is usually uneven across regions. This can be shown directly, for instance with the use of Gini-coefficients. Though it is in accordance with economic geography theory, the fact that activity is distributed unevenly can be explained by other factors as well. As such, it cannot be the only evidence for the theory. However, the uneven structure of production can be used to test for the home market effect, which does point to an economic explanation. We look more closely at this line of research below, in section 5.2.1.

Finally, a strand of literature exists that relates a region's prices of immobile factors to the degree of access to other regions. The idea is that remote locations, facing high prices for intermediate goods and a bad export position, are forced to pay less for immobile factors of production. Specifically, attempts to relate the spatial wage structure to market access have been successful. We discuss this line of literature below, in section 5.2.2.

After our discussion of econometric estimation, we take a look at the more practical art of model calibration in section 5.2.3. In the papers discussed in this section, the researchers have built an economic geography model whose virtues are not under discussion. These models need to be calibrated so that initially, they are able to replicate the data. In the process of fitting the model, interesting choices have to be made on certain parameter values.

5.2.1 Measuring the home market effect

An important series of articles on the matter of testing for economic geography are Davis and Weinstein (1998a, 1998b, 1999). They use the home market effect discussed in section 2.3.3 above to put a trade model to the test. The model uses the Heckscher-Ohlin theory at the level of industries, and allows for a number of alternatives at the level of individual goods.

The home market effect may be observed if a region or a country has a large idiosyncratic demand for a particular good. In the Heckscher-Ohlin theory of trade, where the production structure is driven by factor endowments, such a region will in general be an importer of that good. The presence of an extraordinary level of demand does not affect the location of production. Even though local producers will satisfy the demand to some extent, they will not cover all of it, hence the importer status of the region.

In a world where producers only use one factor and compete within an MC framework, things are different, and the home market effect surfaces: the region with the large demand for a specific good will be a net *exporter* of that good. The producers, rather than being driven by factor scarcity,

realize that they are best off producing in only one location because of returns to scale, and prefer the region with the largest demand because of transport costs. Other regions are then serviced from this one, leading to the net export result. In terms of the model, a large demand component is matched more than one for one by the region's production. This is called the home market effect; it is discussed in Section 2.B.

In the empirical specification, the authors lump together several goods into an 'industry,' explain industry location by endowments and look how the production of particular goods within the allotted industry production is distributed. In Davis and Weinstein (1998a), the data come from the OECD and concern national manufacturing production of the member countries. Results are meagre, in the sense that most trade can be explained by the traditional Heckscher-Ohlin model. When the model is estimated with data for regions in Japan (Davis and Weinstein 1998b), the results are quite different. Estimated for each of the industries separately, two out of six feature a marked 'home market effect.'² The authors indicate that this could be the result of lower transport costs and greater factor mobility at the regional level.

The tests by Davis and Weinstein must be seen as a first coarse investigation into the relevance of economic geography models. Many assumptions are made: at what level of aggregation do 'industries' stop, and do we observe 'goods' where the market is monopolistically competitive? What goods form an industry together? How do arbitrary definitions of regions and groups of goods affect the results?

Two further objections to this strand of research are raised by Brakman et al. (2001, section 5.4.3). First of all, the home market effect that is tested by Davis and Weinstein does not exclusively identify models of economic geography. Brakman et al. argue that the effect is also present in other models that do not have the crucial endogenous market size. Secondly, the home market effect is not very robust in models of economic geography—there exist variants that do not feature the effect. This renders testing on the basis of it inconclusive at best.

5.2.2 Estimating a spatial wage structure

A result that can be obtained using models of Economic Geography, but does not appear in rival models, is the negative relationship between a region's wage and its distance from other economic activity (Brakman et al. 2001, section 5.5). Measurement of this relationship would be a clear indication of the relevance of Economic Geography models, as well as give insight into the values of its key parameters.

²These industries include: transportation equipment, general machinery, electrical machinery and precision instruments. Indeed these skill-intensive goods seem to be among the ones where differentiation is possible, as opposed to manufactured bulk goods.

The relationship can easily be derived from the expression for a firm's foreign demand in formula (2.8). If we generalize from a model with two regions to a model with R regions, we can write demand for firm i in region j as

$$x_{i,j} = p_{i,j}^{-\sigma} \sum_{s=1}^R E_s q_s^{\sigma-1} \tau_{s,j}^{\sigma-1}$$

where $\tau_{s,j}$ is the iceberg-transport cost parameter between regions s and j . It is often parametrized as $\exp(-\xi d_{s,j})$ where $d_{s,j}$ is the distance between the two locations.

Knowing that $x_{i,j}$ must be equal to its optimal level because of the free-entry assumption, and knowing that $p_{i,j}$ is usually just a multiple of the local wage level w_j , we can rewrite the above as

$$w_j = \left[\sum_{s=1}^R E_s q_s^{\sigma-1} \exp(-\xi d_{s,j})^{\sigma-1} \right]^{1/\sigma} \quad (5.1)$$

Hanson (1998, 1999) notices that this spatial wage structure, predicted by Krugman's (1991a) model, resembles the market potential function of Harris (1954). He estimates the spatial wage structure based on a market potential formulation, and goes on to estimate the reduced form of Helpman's (1998) adaptation of the Economic Geography model, which is similar to (5.1).

The estimation in Hanson (1999) is done using a first-difference form of (5.1) to account for the (unchanging) external qualities of the land. Hanson uses data for 3,075 U.S. counties and finds the expected sign for all coefficients, all of which are statistically significant. Using these coefficients, he is able to simulate the effects of a shock in demand in Chicago for the rest of the continental United States.

Brakman et al. (2001) criticize Hanson for failing to recognize the importance of international trade in his estimations. They offer an alternative study into the spatial wage structure in East and West Germany. Their estimates are similar to Hanson's (1998), but do not change very much when trade with the rest of Europe is factored into the model.

We see that two methods have been used to estimate the parameters of economic geography models using statistical methods. The works of Hanson (1999) and Brakman et al. (2000) use the wage structure of US counties and German districts, respectively, to parametrize a Krugman-type model. Davis and Weinstein estimate the coefficients of a model of international trade, interpreting the values in the light of monopolistic competition theory.

In section 5.5, we will discuss the estimation of a Venables-type model, where links between industries work through the use of intermediate goods.

We delay this discussion because it makes use of the model that is to be introduced in section 5.3 below. We will apply the methodology, introduced by Redding and Venables (2001), to a new set of data.

In the next paragraph, we look at the methods that have been used to calibrate models of Economic Geography to data. While this is not estimation *per se*, since the amount of data used is too small to disagree with the model, some insights carry over to the methodology of section 5.5.

5.2.3 Model calibration

Venables and Gasiorek (1996) use a model in which industries are linked through intermediate inputs. They apply this model to several European regions, including the Iberian peninsula and Greece. Factor shares and industry sizes are taken directly from national accounts data. Elasticities are presumably specified using estimates from the literature, but their values are not reported.

Transport costs are a function of the distance between two regions. This function is approximated iteratively: starting with an initial guess, the authors solve the model and compare (gross) trade between regions with actual data. From this comparison, the function that maps distance to transport costs is re-specified, until the outcome matches the data.

Once the gross streams of trade are fitted, the authors go on to specify sector-specific transport costs-functions, preserving size of the total streams.

Bröcker (1998) specifies a general equilibrium model that is reminiscent of the simple Venables model of chapter 4. There are some differences, however: the workers in the model do not choose their sector endogenously, and it is assumed that their wages are equal. The author finds that flows of trade between regions vary inversely with the distance between the trading partners and a coefficient $\rho = \sigma \cdot \zeta$. Here, σ is the elasticity of substitution and ζ a coefficient related to the costs of transport. Bröcker is able to estimate ρ from trade data, but then has to separate out σ and ζ in order to specify his model. He accomplishes this by looking at the share of transport costs in the total value of trade, knowing from his model that this number is equal to $1/2\sigma$. Data on transport costs thus imply that σ is smaller than 10. Furthermore, we know that the markup factor is equal to $\sigma/(\sigma - 1)$. Data on markups suggest that the value of σ must be larger than 5. Finally, the author finds that the outcome of the model changes very little for σ 's between 5 and 10.

5.3 Vertically linked industries and agglomeration

In the rest of this chapter, we will develop a Venables-type of economic geography model and estimate its parameters using data on trade between

American states.

As before, we use the structure that was introduced by Venables (1996a) and that was discussed in the previous chapters. In this model, connections between firms exist because of the use of each other's intermediates in production. We previously studied how the input-output relations between sectors influenced the possible equilibria of this model. In this chapter we assume, rather implausibly, that each producer uses *all* available products as intermediate inputs. In this case, we do not need to specify different sectors, so that we have a single industrial sector that works under Monopolistic Competition (MC) and IO-matrix is collapsed to the number 1. The single-sector assumption is an obvious shortcut with which we will live in this chapter, as it facilitates estimation; we will estimate a more general version in chapter 6.

Next to the industrial sector, whose products can be traded across regions, we assume the existence of a 'local products' sector, whose production is nontradeable. This sector comprises activities such as local services and the production of locally consumed goods. The local sector is perfectly competitive and uses a simple linear production technology where the only input is labor.

In the industrial sector, firms use labor and products from other industrial firms as inputs. Each firm i in region r produces a single variety with production function

$$Y_i = \theta_\alpha L_i^\alpha Q_i^{1-\alpha} - F$$

where $\theta_\alpha = \alpha^{-\alpha}(1-\alpha)^{\alpha-1}$ is a constant scaling term, L_i is the amount of labor employed and Q_i an aggregate of intermediate products used. Finally, a fixed cost F is paid in the company's own product. Adding a fixed cost excludes the solution where a firm would produce an infinite number of goods in infinitely small amounts to exploit the increasing returns to variety.

For the aggregation of industrial products into a single intermediate product we employ the usual CES function, so that

$$Q_i = Q(x_1, \dots, x_N) = \left[\sum_{j=1}^N x_j^{1-\frac{1}{\sigma}} \right]^{1/(1-\frac{1}{\sigma})}. \quad (5.2)$$

As discussed, the intermediate product Q is composed of products of all N firms in the economy. Again, it pays to be close to other firms for two reasons: the lower price of intermediates and the local demand for your own product (these are the forward and backward linkages of Hirschman 1958).

Given that each industrial firm i in region r faces exactly the same conditions, they will all demand the same amount of intermediate product,

which we denote Q_r . The price index associated with Q_r is

$$G_r = G(p_1^r, \dots, p_N^r) = \left[\sum_{j=1}^N (p_j^r)^{1-\sigma} \right]^{\frac{1}{1-\sigma}}, \quad (5.3)$$

$$= \left[\sum_{s=1}^R n_s (p_s^r)^{1-\sigma} \right]^{\frac{1}{1-\sigma}} \quad (5.4)$$

where p_j^r is the price of firm j 's products in region r . Given that all firms in the same region use the same price, we arrive at (5.4), with p_s^r the price of an industrial product from region s in region r ($1 \leq r, s \leq R$) and n_s the number of firms in s . Prices of the same product differ between regions, as there are different costs of transport. The exact nature of the transport costs is discussed below.

Consumers optimize a simple Cobb-Douglas utility function over local and tradeable goods. In region r , the function is

$$U(Z, x_1, \dots, x_N) = U(Z, Q(x_1, \dots, x_N)) = Q^{\mu_r} \cdot Z^{1-\mu_r} \quad (5.5)$$

Here, Z is the quantity of local goods consumed. Industrial products enter as multiples of the aggregate Q , defined in (5.2). This implies that consumers and producers have the same elasticity of substitution between products, equal to σ . Because of this, we can invoke the result that all firms in a region use the same, optimal price, for final and intermediate demand. The price is a markup over marginal costs:

$$p_r^* = \frac{\sigma}{\sigma - 1} w_r^\alpha G_r^{1-\alpha}. \quad (5.6)$$

We denote the number of people working in region r as L_r . Since a fraction $1 - \mu_r$ of the region's wage income³ goes to local producers (per formula 5.5), it follows that a fraction $1 - \mu_r$ of all workers are active in the local-goods sector. The remaining $\mu_r L_r$ people work in the industrial sector. From the size of the workforce we can compute the number of firms, given free entry and exit so that each firm makes a profit of zero. For then we must have that for any firm i , wholesale profits exactly compensate the fixed costs that were incurred, or

$$\begin{aligned} Y_i &= (\sigma - 1)F \Rightarrow \\ \sigma F &= \theta_\alpha L_i^\alpha Q_i^{1-\alpha} \\ &= \frac{L_i (w_r / G_r)^{1-\alpha}}{\alpha} \end{aligned}$$

³In this static model, wage income is the only income since we abstract from saving and capital.

where we used the definition of θ_α , the markup equation (5.6) and the fact that, after optimization, $Q_i = \frac{(1-\alpha)w_r}{\alpha G_r} L_i$. This gives us the optimal amount of labor used by firm i in region r , L_i^* . The number of firms in region r can then be computed by dividing the number of industrial workers in the region by L_i^* .

$$n_r = \frac{\mu_r L_r (w_r / G_r)^{1-\alpha}}{\alpha \sigma F} \quad (5.7)$$

The number of firms in a region varies as the local price index G_r changes.

Transport costs are incorporated in the model using the ‘iceberg assumption’, whereby transport charges are incurred in the product itself. The amount that needs to be shipped to get one unit of the product to arrive from location s in location r , T_{sr} , corresponds to the distance travelled as

$$T_{sr} = \exp(\tau d_{sr})$$

where τ is a positive parameter. Alternatively, one could consider T_{sr} a markup over the home price: for each region s , there holds $p_s^r = p_s^r T_{sr}$. We rewrite the price index in (5.4) as

$$G_r = \left[\sum_{s=1}^N n_s (p_s T_{sr})^{1-\sigma} \right]^{\frac{1}{1-\sigma}}. \quad (5.8)$$

Given parameters, wages, and numbers of workers L_r we can now solve the model in terms of prices, price indices and the equilibrium number of firms. These three sets of region-specific variables p_r , G_r and n_r can be stored in three vectors of length R . Each is defined in terms of the other: we have vector \mathbf{p} from equation (5.6), \mathbf{n} from equation (5.7), and \mathbf{G} from equation (5.8). Ideally, we would solve this system of $3 \times R$ equations analytically. However, this is not possible for reasons that were discussed in chapter 2, so we have to rely on numerical methods instead. We use an iterative routine in Matlab to find the values of \mathbf{p} , \mathbf{G} and \mathbf{n} that satisfy the equations. In general, this routine finds the equilibrium very quickly.

In this model, a region’s expenditures on industrial products comes from final and intermediate demand. In region r , the L_r consumers spend $E_r^f = \mu_r w_r L_r$ on industrial products from all over the economy. Firms in region r buy intermediates, spending an amount directly proportional to their total wage bill: $E_r^{\text{int}} = \mu_r L_r w_r ((1-\alpha)/\alpha)$ (this follows from the Cobb-Douglas production function). The total expenditure in region r is thus $E_r = E_r^f + E_r^{\text{int}} = \mu_r L_r w_r / \alpha$.

Depending on the parameters, this model can have several different equilibria. If the costs of transportation and the share of intermediate products in production are low, and the elasticity of substitution σ is high, then economic production will be distributed proportional to the population size. But if transport costs are high, intermediate products are important

and σ is low, then production can agglomerate in one or a few regions (For a full derivation of these results, see Venables 1996a). Which region gets the agglomeration is decided by initial conditions.

Region r 's expenditures on industrial goods, E_r , are spread over all industrial producers in the economy. Products from region s cost $p_s^s T_{sr}$ and the price index is as in (5.8). Standard Dixit-Stiglitz optimization leads to the familiar result that in region r , the demand for a product from region s is

$$D_{rs} = E_r (p_s^s T_{sr})^{-\sigma} G_r^{\sigma-1}. \quad (5.9)$$

Given that there are n_s firms in region s , each producing a unique differentiated product with the same price, the total demand in region r for products from region s is n_s times the expression in (5.9). To get the value of this stream of goods, we also multiply by the price⁴ which gives

$$X_{rs} = n_s (p_s^s)^{1-\sigma} T_{sr}^{1-\sigma} E_r G_r^{\sigma-1} \quad (5.10)$$

with X_{rs} the value of shipments from s to r .

Of course, equation (5.10) is reminiscent of the gravity equation. The term n_s is indicative of the economic size of the sending state, just as E_r is of the receiving state. The distance between the two is captured by T_{sr} . But the price index term G_r also fits in nicely with the literature on gravity-models. It serves as a proxy for what has been called *remoteness* in Wolf (1997), the property that two regions will trade more than the simple gravity model predicts if the two of them are close, and relatively far from all other regions.⁵ In this model, the two isolated states will have relatively high values of the price index G . This depresses their total trade, but relatively increases their bilateral trade. To see this, write the stream of goods in (5.10) as

$$\bar{D}_{rs} = n_s \frac{E_r}{G_r} \left(\frac{p_s^s T_{sr}}{G_r} \right)^{-\sigma}.$$

We see from the second factor that a high G_r causes the real trade spending E/G to be low. When this high G_r is caused by high values of T_{sr} as we assumed, this is not remedied by the fact that the price index enters with a positive power in the third term. However, if the value of the T_{sr} is especially low for a certain region s , trade between regions s and r is

⁴Due to the assumption of 'iceberg' transportation costs, to get the amount of goods in (5.9) to arrive in region r the producers in region s must actually ship T_{sr} times as much. This way, they account for the goods that 'melt' in transit. In many papers, this leads to an extra factor T_{sr} in this expression. However, iceberg transport costs are a convenient fiction and these extra goods are not observed in the data; it is therefore defensible to leave the extra T out, as we do in other chapters. Here, for consistency with other work, we maintain the extra T .

⁵Wolf uses as a measure for remoteness the ratio of bilateral distance to the average of the output-weighted mean distance to all other regions. This regressor is expected to have a negative sign.

relatively high. A specific case of two isolated states is examined below, on page 139. Related work that derives the notion of gravity and remoteness from a trade model based on CES-demand functions can be found in Anderson and van Wincoop (2003).⁶ There, price indices G_i are referred to as “multilateral trade resistance,” as they serve to measure the average impediments to trade for region i .

As we saw in formula (5.7) above, if we assume that each firm makes zero profits, each firm’s production is fixed at certain level \bar{Y} . Because production, or supply, must be equal to demand, this introduces a relationship between the factors that influence total demand (formula 5.9) and the factors that determine a firm’s price (formula 5.6):

$$\begin{aligned} \sum_{j=1}^N E_j G_j^{\sigma-1} T_{i,j}^{1-\sigma} &= \bar{Y}_i p_i^\sigma \\ &= \bar{Y}_i \left[\frac{\sigma}{\sigma-1} w_i^\alpha G_i^{1-\alpha} \right]^\sigma \end{aligned} \quad (5.11)$$

This formula shows that there exists a negative relationship between a region’s transport costs $T_{i,j}$ and its wage level w_i . Regions which are far away from large markets (and have a small market themselves) can be expected to have lower levels of wages compared to those close to the industrial core. This relationship forms the basis for several exercises, which study the model’s relevance by looking at the correlation between a region’s wage and its ‘closeness’ to other regions. As we saw in section 5.2.2 above, such a relationship is tested by Hanson (1998, 1999) for American counties and by Brakman et al. (2001) for German regions. Redding and Venables (2001) test the relationship using data on 101 countries worldwide, after they have approximated the term on the left hand side of this equation. We will take a closer look at this study below, where we use the same methodology on

⁶Anderson and van Wincoop (2003) do not explicitly model supply, but assume that each region produces one unique variety of goods. Demand takes a CES form over all goods and market equilibrium allows them to derive the gravity equation

$$X_{rs} = \frac{y_r y_s}{y_w} \left(\frac{T_{sr}}{G_s G_r} \right)^{1-\sigma}$$

where the notation is in terms of this chapter. Regional income y_s and y_r play the role of n_s and E_r in (5.10), normalized by world income y_w . The crucial difference between their model and the current model lies in the role of G_s , the price index in the sending state. In our model, a higher G_s implies higher costs of production, and thus a higher price: G_s appears in (5.10) mainly as a component of p_s^* , see (5.6). Thus, high values of G_s inhibit trade. In Anderson and van Wincoop, a higher G_s encourages trade as

Higher barriers faced by an exporter will lower the demand for its goods and therefore its supply price p_i . (p. 176)

This result would be hard to obtain with a production model in which the price of intermediate inputs plays a role in the price of final production.

our own data. This will be the subject of in section 5.5.1. First, we look at the available data.

5.4 Data

The data used in this estimation concern 51 US states and are described in detail in the appendix on page 144. We make use of the Commodity Flow Survey 1997, a dataset compiled by the US Bureau of Transportation Statistics. It is an estimate of inter- and intrastate trade based on a survey among firms that ship traded goods within the US, using a variety of modes. Our other two main datasets concern wage levels in each state and the distances between states⁷ and are discussed in the appendix. A fourth set, data on gross aggregate output per state, is used in this paragraph only. A map of the 48 continental states and their two-letter abbreviations is in figure 5.1 below.

Before estimating the model we specified above, we will use this data to estimate a simple gravity-equation. This simple model can serve as a benchmark with which we can compare the results. The model we estimate is

$$\log X_{rs} = \alpha_1 + \alpha_2 \log Y_r + \alpha_3 \log Y_s + \alpha_4 \log D_{rs} + A_{rs}\beta + \varepsilon_{rs}. \quad (5.12)$$

This equation explains the flow of trade (X_{rs}) by total production in the receiving and in the sending state (Y_r and Y_s , respectively), the distance between the two states (D_{rs}) and a set of k dummy variables contained in the $1 \times k$ vector A_{rs} . We present estimates with different A 's, where k varies between zero and two.

A possible dummy variable contained in A is the intra-state dummy, which is one if the sending and the receiving state are the same, and zero otherwise. Another possible dummy variable is the adjacency- or border-dummy, which is one if the sending and receiving state share a border, and zero if they do not.

The results are in Table 5.1. This simple gravity equation explains about 80 percent of the variation in trade flows. The elasticities to income α_2 and α_3 are almost equal and close to one. There is a substantial home bias in trade, measured by the intra-state dummy, and a positive effect of sharing a border. These two effects reinforce the negative coefficient on distance.

We compare our estimates with those in Wolf (1997) who uses an older and less complete version of the Commodity Flow Survey. Note that Wolf uses a form slightly different from (5.12), namely

$$\log X_{rs} = \alpha'_1 + \alpha'_2 \log(Y_{rs}^{\text{comb}}) + \alpha'_4 \log D_{rs} + A_{rs}\beta' + \varepsilon'_{rs}. \quad (5.13)$$

⁷Distances are computed as the length of a straight line between the weighted centers of two states, see appendix 5.A below.

Dep. Variable	Trade 97	Trade 97	Trade 97	Trade 93
Technique	OLS	OLS	OLS	OLS
Observations	2201	2201	2201	1030
Source	Own comp.	Own comp.	Own comp.	Wolf 1997
<i>Constant</i>	-8.843 [-24.993]	-10.014 [-29.262]	-11.354 [-31.929]	-8.06 [21.04]
$\log(Y_r)$	0.992 [52.774]	0.998 [56.164]	1.004 [57.931]	0.961* [66.23]
$\log(Y_s)$	0.995 [53.377]	1.003 [56.925]	1.011 [58.804]	
$\log(D_{rs})$	-1.171 [-50.931]	-1.029 [-43.890]	-0.864 [-31.454]	-0.956 [33.69]
<i>Intra state</i>		2.076 [16.195]	2.492 [19.073]	1.338** [11.68]
<i>Adjacency</i>			0.763 [10.849]	
R^2	0.787	0.810	0.820	0.867
SSR	1728	1544	1465	

* Wolf uses one regressor, $\log(Y_r Y_s)$, instead of two. ** Note that Wolf divides the state-income variable Y by two if the trade flow is within the state itself (*i.e.*, the receiving and sending state are the same). This increases the intra-trade parameter. In our estimates, Y is not changed.

Table 5.1: Gravity equations

When the sending and receiving state are different, $Y_{rs}^{\text{comb}} = Y_r Y_s$, but when the state trades within itself there holds

$$Y_{rs}^{\text{comb}} = \frac{1}{2} Y_r \frac{1}{2} Y_s$$

with $r = s$, so that the regressor in that case is equal to $2 \log Y_s - \log(4)$.⁸

It seems that the parameter estimates between the two exercises are reasonably close. The only significant change can be seen in the parameter of the intra-state dummy. Our estimates show a higher value, and would have

⁸Without the exception for intra-state trade, the use of Y^{comb} amounts to the restriction that $\alpha_2 = \alpha_3$. Now, because the regressor Y^{comb} has a positive coefficient and is decreased for intra-state observations, the coefficient on the intra-state dummy is biased upward.



Figure 5.1: The 48 contiguous states in the US and their two-letter abbreviations.

been even higher had we followed Wolf's practice of dividing the GSP variable by two for intra-state flows. So where Wolf (1997) finds a home-bias effect, the later version of the same data shows an *increase* in that effect.

5.5 Estimation

5.5.1 The Redding-Venables approach

Redding and Venables (2001) use a closely related version of the above model to estimate its parameters in two steps. For this, they use data on 101 countries, including the trade flows and distances between them, whether they share a border and the level of wages in each country, approximated by GDP per capita. This data is used to estimate formula (5.10) with panel data methods as a gravity equation. Then, using the projected values for the terms in (5.10), they test the relationship between average trade costs and the level of wages in a region that we saw in formula (5.11). We will go over these steps in turn.

First stage estimation: gravity

For their initial estimation, the authors rewrite the trade relationship in formula (5.10) as⁹

$$X_{rs} = \phi_s T_{sr}^{1-\sigma} \psi_r \quad (5.14)$$

⁹As discussed in footnote 4, the authors add an extra T_{sr} to the equation. We will follow their convention here.

where ϕ_s is called country s 's *supply capacity* and ψ_r is country r 's *market capacity*. Each of these two terms contains information on a country's trade characteristics that is the same towards all its trading partners. Market capacity $\psi_r = E_r G_r^{\sigma-1}$ reflects the total amount of imported goods absorbed by country r . It increases when the country spends more on imports, or when it is (on average) far away from its trading partners.¹⁰ Supply capacity $\phi_s = n_s p_s^{1-\sigma}$ varies with the number of firms in country s , and hence with its total production of tradeables.

Given the structure of formula (5.14), it is possible to estimate it using the fixed effects panel data method. We rewrite the equation as

$$\log(X_{rs}) = \delta_0 + \phi'_s \iota_s + \delta_1 \log(\text{dist}_{rs}) + \delta_2 \text{bord}_{rs} + \delta_3 \iota'_s \iota_r + \psi'_r \iota_r + u_{rs} \quad (5.15)$$

Once again, X_{rs} is the value of the flow of trade from region s to region r . The $N \times 1$ vector ι_i is filled with zeros, except at the i th position, where it is one. Thus, the $N \times 1$ vectors ϕ and ψ contain the supply and market capacities of all regions.

The dependency between distance and trade is captured by the δ -parameters. The first, δ_0 , is a scaling factor. Distance (in miles) has a coefficient of δ_1 , which we expect to be negative. The influence of the spatial characteristics of the two regions is further captured by two dummy-variables: bord_{rs} is one if the two regions r and s share a border¹¹ and the product $\iota'_s \iota_r$ is one only if the sending and receiving state are the same.

With data about the distances between regions, the size of their bilateral trade and whether they share a border, it is possible to estimate the parameters in relation (5.15). We shall do so for our data on the states in the US, described in section 5.4 above. For comparison, we also mention the outcomes of Redding and Venables (2001). They use data on 1994 bilateral trade flows between 101 countries. The distance between two countries is that between the capital cities. Trade within a country is not taken into account in their estimations, so the regressor $\iota'_s \iota_r$ is left out. In our dataset, data on trade within a state is available; we estimate both with and without it.

What can we expect *ex ante* about the differences between the two estimations? Given that the methodology is exactly the same, variations in outcomes must be caused by differences between the two datasets. Firstly, the dataset of Redding and Venables is larger by a factor of four; *ceteris paribus*, this leads to smaller estimation errors. However, their data pertains to the whole world and is probably more heterogeneous than that measured

¹⁰Countries with a small home market that are far away from trading partners will have a high value for G ; this means that they will be less daunted by high import prices, since *all* of their import comes from far away.

¹¹In the United States, some states share a border of size zero as their corners just touch each other. This is the case for Arizona and Colorado, for instance. In spite of this tangential relationship, the border-dummy is set to one for these pairs of states.

within the United States. For instance, the distance between two countries is likely to include a stretch of ocean, whereas this rarely happens between two US states. Given that trade over sea is more complicated, we expect higher trade costs in the World dataset. Also, two countries sharing a border is a more unlikely event than two states sharing one. This could make the effect of borders more significant in the World dataset. Finally, trade between countries may or may not be hampered by restrictions such as tariffs, or by cultural differences. Given the relative homogeneity of the US states, we expect less unexplained variation in the latter sample.

The results of the estimation are in table 5.2 on page 146. There are three estimations for both datasets, and two extra for the US dataset. We report the values for $\hat{\delta}_1$, $\hat{\delta}_2$ and $\hat{\delta}_3$, leaving the (large) vectors $\hat{\phi}$ and $\hat{\psi}$ out. These coefficients will be used later on, however.

In the first estimation (in the first two columns), the full sample is used. This includes pairs of regions for which no trade is recorded. For both datasets, this means that the actual trade between the two regions is probably very small. We substitute a zero for (the logarithm of) these unmeasured flows. We see that distance has the expected negative sign, whereas the border-dummy has a positive parameter. Both are highly significant. The coefficient of the variable $\iota'_s \iota_r$, called own_{rs} in the table, is also positive and significant. As expected, both distance and the occurrence of a border have a larger effect in the World dataset. The explained variance is about the same for both.

In the second estimation, pairs of regions between which no trade is recorded are taken out of the sample. This leads to smaller, but more significant coefficient estimates. For the World dataset, the R^2 does not increase; leaving out the zeros does not improve the performance of the model. The R^2 does increase, markedly, for the US dataset. This is caused by the fact that many unobserved pairs involve either Hawaii or Alaska, two states which turn out to be outliers in this dataset.

In the third estimation, we reintroduce the pairs with unobserved trade and treat them as left-censored observations. The model parameters are estimated using the Tobit method. This increases the coefficient on distance and decreases the border dummy. Standard errors are slightly worse, though.

The final two columns pertain only to the US dataset. In the fourth estimation, we use only contiguous states, eliminating Alaska and Hawaii from the sample. These two states suffer from many missing observations, whereas those that are available act as outliers. The District of Columbia is also struck from the sample, as the model also performs relatively badly for this region. This is probably due to its small size and atypical sectoral makeup. In the fifth estimation we eliminate the remaining 49 observations of in-state trade data. This hardly affects any parameters, showing that the use of an in-state dummy adequately captures the special nature of trade

within the same state.

When we compare the parameter estimates for trade within the United States with those for world trade, at first glance the results rather similar. All corresponding parameters have the same sign and the order of magnitude is the same for similar parameters. The differences do amount to several times the standard error, though: the effect of distance and the effect of a shared border are greater for world trade data. The explanatory power of the model is greater for US data, however. Partly, this can be explained by the absence of administrative and physical barriers in the US. Also, the data on world wages is a proxy (GDP per capita), giving rise to extra measurement error.

Second stage estimation: Wages

We keep the results of the previous exercise to conduct a second stage estimation. For this, Redding and Venables construct two new variables: *Market Access* of a region s is defined as

$$\begin{aligned} MA_s &= \sum_{r=1}^N E_r G_r^{\sigma-1} T_{r,s}^{1-\sigma} \\ &= \sum_{r=1}^N \phi_r T_{r,s}^{1-\sigma} \end{aligned} \quad (5.16)$$

and *Supplier Access* of region r as

$$\begin{aligned} SA_r &= \sum_{s=1}^N n_s (p_s T_{r,s})^{1-\sigma} \\ &= \sum_{s=1}^N \psi_r T_{r,s}^{1-\sigma}. \end{aligned} \quad (5.17)$$

The names of these variables suggest that they are not chosen at random. Market access is a weighted average of the expenditures on differentiated goods by the region's potential trading partners. The weights contain distance to the region with a negative sign and the relative isolation of the potential trading partner (as indexed by their price index G) with a positive sign. As such, the measure is reminiscent of the market potential function suggested by Harris (1954).

Supplier access in (5.17) is inversely proportional to the regional price index G_r , as defined in (5.4). It is an index of the ease with which firms in the region can get intermediate goods, and with which consumers can get final goods. The two variables defined above share two desirable traits: firstly, using the results from our first-stage estimation, we can compute

their values. Secondly, they are related to the level of wages in a region and thus offer a way to test the model.

Computing the values of MA_s and SA_r involves using the estimated values of ϕ and ψ that we obtained earlier, and our estimate of the costs of transport. We construct

$$\begin{aligned}\widehat{MA}_r &= \exp(\phi_r) \cdot \text{dist}_{r,r}^{\delta_1} \cdot \exp(\delta_3) + \\ &\quad \sum_{s \neq r} \exp(\phi_s) \cdot \text{dist}_{s,r}^{\delta_1} \cdot \exp(\text{bord}_{s,r})^{\delta_2} \\ &\equiv \text{DMA}_r + \text{FMA}_r\end{aligned}\quad (5.18)$$

and

$$\begin{aligned}\widehat{SA}_r &= \exp(\psi_r) \cdot \text{dist}_{r,r}^{\delta_1} \cdot \exp(\delta_3) + \\ &\quad \sum_{s \neq r} \exp(\psi_s) \cdot \text{dist}_{s,r}^{\delta_1} \cdot \exp(\text{bord}_{s,r})^{\delta_2} \\ &\equiv \text{DSA}_r + \text{FSA}_r.\end{aligned}\quad (5.19)$$

Notice that we used our estimate of $T_{r,s}^{1-\sigma}$ from the previous section, which uses measures of distance, a border- and an own-state-dummy. In these formulas, we implicitly defined four other *access*-variables by splitting off access to the own region from access to other regions. DMA and DSA are domestic market- and supplier access, and FMA and FSA their foreign equivalents. Separating these terms will allow us to test them separately, later on.

To see how these measures of access interact with the wage level, write equation (5.11) as

$$\alpha\sigma \log(w_i) = \zeta + \log(MA_i) + (1 - \alpha) \frac{\sigma}{\sigma - 1} \log(SA_i) + \epsilon_i \quad (5.20)$$

for a region i . Notice that both market and supplier access have a positive coefficient in this equation. Products from a region with low market access incur large transport costs before they reach their customers. As these products have to compete with other, cheaper products, this limits the wages that can be paid in their production. Similarly, low supplier access means that intermediate goods are expensive: this squeezes the value that can be added in a region from the other side.

We will estimate equation (5.20) using generated values for MA and SA. These are computed as in (5.18) and (5.19), using predicted values for ϕ and ψ . This procedure renders OLS standard errors unusable: the stochastic errors in the gravity equation (5.15) turn up in the predicted values of MA and SA, which affect the stochastic behavior of ϵ_i in (5.20), violating the assumptions that underlie standard OLS analysis.

To estimate the standard error in spite of these difficulties, bootstrap methods are available (see Efron and Tibshirani 1993, for instance). For the gravity equation, we construct a new sample of the same size by drawing random observations (each observation a flow of trade and its regressors) from the original sample. This sample is a bootstrap-replication, for which original observations may be absent, or appear more than once. From the bootstrap-replication we re-estimate the trade-equation (5.15) and use the outcome to generate \widehat{MA} and \widehat{SA} as usual, which together with observations on wage make up a sample for equation (5.20). We generate 200 samples this way, the conventional number of bootstrap-replications according to Efron and Tibshirani. Of each of these samples, we use the same procedure to generate 200 bootstrap-replications. Estimating equation (5.20) on the resulting data gives forty thousand estimates, from which the standard error of the regressors can be directly observed.¹²

Several other problems potentially plague this estimation. As Redding and Venables remark, a contemporaneous shock to a region that affects both the independent variable and the regressors could introduce a bias the results. To eliminate the possibility of contemporaneous shocks, we estimate using wages from 1999 with regressors from 1997. This does not eliminate another class of ‘third variables,’ a time-invariant region-specific effect that plays in both a region’s wage and in its market- and supplier access. To correct for this possibility, we report regressions on total access as well as ‘foreign access,’ as defined in (5.18) and (5.19). In the latter regressor, data from the own region does not play a role. Below, we will also add a number exogenous regressors that proxy for a region’s time-invariant attractiveness and may capture its effect.

To start, we have to select a first stage estimation from the previous paragraph with which to work. We select the one that gives the best fit, called US 4 in table 5.2. This estimate uses the sample of all contiguous states, with trade flows including those to the sending state itself. As it turns out, the Market Access and Supplier Access regressors are highly colinear; the correlation between the two series is 0.95. This means that estimating (5.20) directly would be problematic. We proceed by using just Market Access as a regressor. At the end of the paragraph we compare the results to those obtained with Supplier Access.

The results of the estimation are in table 5.3 on page 147. We report the estimates on our US dataset, as well as the results obtained in Redding and Venables (2001, table 2). A scatterplot of the first two regressions for the United States is in figures 5.2 and 5.3 on page 148. Each point in the plot represents a state, indicated by its two-letter abbreviation. The horizontal

¹²As it turns out, bootstrap standard errors lie between one and two times the (invalid) OLS-standard errors, indicating that the extra variability due to generated regressors is reasonably small. We report only bootstrap-standard errors.

axis in figure 5.2 gives predicted market access according to formula (5.18). On the vertical axis the log of that state's average annual wage is plotted. Figure 5.3 is similar, only this time the variable on the horizontal axis is foreign market access.

From the first two columns, we note that the relation between foreign market access and the level of wages is much weaker in our estimation than in the *World* dataset. Both the explained variation and the statistical significance of the coefficient are smaller. The coefficient does have the right sign, however. From the scatterplot in figure 5.3 we can learn about the reasons for this weak performance. There is a clear positive relationship between FMA and wages for small states, such as Delaware (DE) and Vermont (VT). However, there are a number of outliers that spoil the correlation. These outliers consist of large states, whose own market is not a part of foreign market access. Especially those that are surrounded by (economically) smaller states fall outside the usual relationship, e.g. California (CA) and Texas (TX). This makes sense: explaining the wage levels in California by its proximity to Nevada and Arizona is bound to be problematic, but New Jersey's wage levels certainly have something to do with its wealthy neighbors.

The fact that relatively large states disturb our measurements may be an explanation for the fact that this estimation works better for worldwide data, where the dominance of large states is perhaps less of an issue.¹³

These problems disappear when we use full market access (MA) as a regressor, in the third and fourth column. The explained variance is about the same as in the *World* dataset, as is the statistical significance. This points to a large role for domestic market access, which is confirmed by the final estimation in columns five and six. Even though both coefficients have the correct sign, DMA clearly trumps FMA as a regressor for wages.

There may be a problem with the use of full market access as a regressor, though. As local demand in a state is included in this variable, local shocks that affect productivity in a state show up in the regressors as well as in the dependent variable. This causes simultaneity bias in the estimation.

Another detrimental effect of including local market access can be seen in the last two rows of table 5.3. There, we report the results of Moran's *I* test on the residuals of the estimated wage equation. Moran's statistic tests for spatial autocorrelation (see Cliff and Ord 1973, van Oort 2002, chapter 4) using a weight matrix to indicate which regions are close to each other. We use the matrix B as the weighing matrix, in which entries are equal to one if the two states share a border.¹⁴ The diagonal of B consists of zeros. We

¹³According to the BLS (see appendix for data sources), at the end of 1997 California, Texas and New York together accounted for 25% of employment in the USA.

¹⁴The choice of the weight matrix is, to a degree, arbitrary and its impact should be measured. We have computed alternative statistics using a matrix B' where $b'_{ij} = \exp(-.001 \cdot \text{dist}_{ij})$ (with dist_{ij} the distance between states i and j) and found that their level of signifi-

have used the data in B before, to estimate the trade equation (5.15).

Moran's I statistic is computed as

$$I = \frac{N}{\iota' B \iota} \frac{\epsilon' B \epsilon}{\epsilon' \epsilon} \quad (5.21)$$

with N the number of observations, ι a $N \times 1$ vector of ones and ϵ the $N \times 1$ vector of errors. In table 5.3 we also report the place of each Moran's I in the distribution of this statistic (under the hypothesis of no spatial autocorrelation).¹⁵ All realizations of the statistic allow us to reject zero spatial autocorrelation at the 1,5% level, indicating that a high realization of the wage in one state makes a higher than expected wage in the bordering states more likely. However, the estimations which include local market access as a regressor show by far the most significant realizations of this statistic.

Are things any different when we use supplier access instead of market access as an explanatory variable? Our theoretical model tells us that SA and MA each determine part of the variation in wages, as can be seen in equation (5.20). However, we determined above that the pair of regressors suffers from severe multicollinearity and decided to include only measures of market access in the regression. By the same token, we could have decided to use only supplier access. The results of this estimation are in table 5.4.

Once again, we compare our results with those in Redding and Venables (2001). We see a similar pattern as in table 5.3: a regression using only foreign access gives a lower, and less significant, value of the coefficient and a lower R^2 compared to the World data set. Using a full measure of supplier access improves the estimation but leads to higher spatial autocorrelation in the residuals.

We will try to improve these estimations below by adding data on the exogenous amenities to productivity that characterize each state, as well as by employing instrumental variables in our estimation.

Exogenous amenities

When we estimate state-level wages as a function of market- and supplier access, we neglect all other factors that may also have a bearing on those wages. In as much as these factors correlate with our regressors, they can

cance was very close to the values obtained with B .

¹⁵The expectation of Moran's I is $-1/(N - 1)$, with N the number of observations. We bootstrap the distribution of I by generating 100,000 vectors ϵ^* , where each ϵ^* is a random permutation of ϵ (in the usual terminology of spatial autocorrelation, we use *nonfree sampling*). We compute the corresponding values of I , and indicate the percentage of outcomes higher than the recorded statistic. An asymptotic distribution for the statistic is known (Cliff and Ord 1973, chapter 2) but its small-sample behavior inspires more confidence in bootstrap methods (see Anselin and Florax 1995).

cause a bias in the estimation. It is easy to think of a situation in which this may happen.

In the first paragraph of this chapter we mentioned that so-called ‘first-nature’ causes of geographic concentration also play a role: the physical features of the area, its climate and natural infrastructure may have an effect on productivity. Imagine, for instance, that a predominantly warm climate opens up economic opportunities (*e.g.*, tourism) in a state. This may raise the general level of wages. If a number of neighboring states share the same climate, this ‘third factor’ will increase wages in all of them. Being close together, market- and supplier access for each of these states will probably be at a comparable level. Suppose it is lower than average; in that case, the unobserved regressor ‘climate’ causes a downward bias in our estimates.

In order to test the robustness of our initial estimates against the influence of third factors, this section presents the results of a number of regressions similar to those above, but including a number of possible third factors as regressors. We use the following exogenous amenities:

- **Climate.** In order to control for an exceptionally warm or cold climate we use two regressors, normal yearly heating degree days (*nrmhdd*) and normal yearly cooling degree days (*nrmcdd*). The former is defined as the cumulative number of (Fahrenheit) degrees in a year by which the mean temperature of each day falls below 65°F, the latter as the cumulative number of degrees in a year by which the mean temperature lies above 65°F.¹⁶ The idea is that an exceptionally warm or cold climate may account for differences in productivity. For reasons of scale, we divide these regressors by 1000 in the actual regression.
- **Geology.** Special economic opportunities may arise from the presence of precious minerals in a state. To proxy for these opportunities, we use the value of nonfuel mineral production per square kilometer in 1997, as reported by Smith (1997), in thousands of dollars.
- **Access to sea.** Finally, we include a dummy variable that indicates if there exists a deep sea port in the state. Access to sea may proxy for the possibility of international trade.

We expect heating- and cooling degree days, regressors that indicate an unpleasant climate, to have a negative impact on productivity. The presence of minerals is likely good for wages, as is the presence of a port. We first regress wages on these exogenous amenities alone, and then include our measures of market access. The results are in table 5.5.

¹⁶Somewhat counterintuitively, cooling degree days measure warmth and heating degree days measure coldness. An example may clarify: if the mean temperature in a state is 67°F all year long, the yearly cooling degree days are $(67 - 65) \times 365 = 730$ and the yearly heating degree days are zero.

From the first column of this table, we notice that the four exogenous regressors have the expected sign and succeed in explaining about half the variation in wages. However, Moran's I is rather high (higher than all but 2.8% of the distribution under H_0) and indicates that we may not have all region-specific exogenous amenities in our set of regressors. The inclusion of Foreign Market Access in the regression hardly changes the values of the earlier coefficients. However, the value of the coefficient for FMA is about a third of the earlier measure (table 5.3) and lies below one standard error. Explained variation hardly improves. Including FMA does improve Moran's statistic to a point where we are no longer able to reject the hypothesis of no spatial autocorrelation at the 5% level.

Things turn out differently when we include measures of (Domestic) Market Access. The coefficients of the exogenous regressors change substantially (more than one standard error in all cases) and Moran's statistic again increases to a significant level. This result again points to problems with the inclusion of Domestic Market access.

IV estimation

The estimations above may suffer from the occurrence of simultaneity bias, which occurs when the error term from an estimation is correlated with one (or more) regressors. In this matter our model is clearly the culprit, as it indeed allows the error terms to influence the market access regressors. We discuss how this happens and how we can correct for it. After that we assess the size of the problem.

The error terms in the regression imply that observed wages are, to a degree, inconsistent with our model, either because of measurement error or because of misspecification. Where we expect a wage w_i^* in state i , we actually find $w_i = w_i^* + \epsilon_i$. That w_i is the dependent variable in our estimation, but it also makes its way into the regressors; according to (5.6), prices are a function of the wage and via (5.8), those prices end up as an element of all the price indices $G_r, r = 1, \dots, N$. Our regressors, MA and SA, are again a function of prices and price indices (*cf.* formulas 5.16 and 5.17). This puts the error ϵ_i in the (supposedly) exogenous variables. The question is, whether the weight that ϵ_i receives in MA_i and SA_i is large enough to influence the estimation.

With this problem in mind we used two regressors above, MA and FMA, where the former excludes market data from the own state. The use of local market capacity in the regressor MA will probably introduce ϵ in MA with a large weight. Indeed, we find that the regressions where FMA is used instead of MA show lower spatial autocorrelation of the errors.

However, we can also eliminate ϵ from the regressors entirely if we employ instrumental variable estimation. This idea is used in Mion (2003), who takes a panel-approach on Italian data. Brakman et al. (2004) use it to

isolate the effect of one particular disturbance in a spatial growth process and Ciccone and Hall (1996) employ four “deep historical” instruments that proxy for the innate attractiveness of American states as places of residence.

For IV, we need instruments that correlate with the regressors MA and SA, but not with the errors ϵ . Once again following Redding and Venables (2001) we use distance from major economic centers as instruments, in particular the distance from New York City and from Los Angeles.¹⁷

The results are in table 5.6. The first two columns use only Market Access variables as regressors, and can be compared to the results in table 5.3. Note that the coefficient of FMA is comparable, while MA has a coefficient that is much lower than before. It appears that simultaneity bias plays an important role in the estimations which use this regressor.

The last two columns once again make use of the exogenous amenities that were introduced above, and can be compared to table 5.5. Here also, the coefficient for MA has fallen. Note that we can no longer reject the hypothesis of no spatial autocorrelation in the third column. On the other hand, the significance of Market Access as a regressor is tenuous.

Discussion

We have estimated a relationship that explains the levels of wages in the United States by the level of market access. The variable that indicates market access is itself a construct from the results of a regression, which resembles a gravity-type relationship. To construct the measures of access, heavy use was made of the theoretical model of economic geography.

The estimations mimic those of Redding and Venables (2001), but the results are less satisfying. To a certain extent, this can be explained by the nature of our dataset: it is smaller and possibly more dominated by large regions. However, the fact that we use data on US states also brings some advantages, which fail to realize. For our regressor, we are able to use actual recorded wages instead of a proxy.¹⁸ Also, institutions are bound to be more similar inside the USA than worldwide. This means that institutional differences (and, for that matter, international frictions such as tariffs) are no longer a factor. These differences were proxied for by distance, but supposedly less than perfectly. In spite of these advantages, the explanatory power of our model, especially when it relies on foreign market access, is less than that measured on a worldwide scale. The same hold when supplier access is used as a regressor.

¹⁷As usual, distance is measure from the (employment-weighted) center of the state so that New York and California each have positive distances to these economic centers.

¹⁸Redding and Venables (2001) use GDP per capita for their main estimations, although they do estimate the relation with wage data for a smaller sample.

Our initial estimations suffer from an omitted variable bias that results in spatial autocorrelation of the errors. We remedy this problem by introducing an extra set of regressors that proxy for exogenous qualities of each state, such as climate and infrastructure. Furthermore, we estimate using distance to economic centers as an instrument. This ensures that the active element in the Market Access variables is indeed the access to markets in other states. These estimations show that the explanatory power of the model is present, but limited.

A potential problem with the methodology used above is the fact that the estimated relationships are not necessarily consistent with the general equilibrium solution of the model. For instance: when we start off estimating the gravity equation in (5.14), we parametrize the relation in (5.10). The latter shows that each region's supply capacity is directly related to the number of firms and the price, both of which are in turn determined other variables in the model, as seen in (5.6) and (5.7). The same goes for market capacity. However, this relationship is not used in the procedure until much later: only when we regress regional wages on the access variables do we observe that in fact, the relationships of the model do not hold: if they did, the regression would have had to give us a perfect fit. The variables that were kept constant would not, had they been subjected to the rules of the model, have stayed so.

This leaves us with the question of how to interpret the findings in this section. On the one hand, we have shown that there exists a significant correlation between regional wages and access variables. This is an indication that the model has some explanatory power in our dataset. On the other hand, the less-than-perfect correlation between wages and access variables shows us that some of the relations inside the model are violated; with our current methodology, we have worked from one end of the model (the trade relationship) towards the other end, leaving all discrepancies to accumulate along the way.

There are many plausible reasons why, even if the real world were governed by this model, we could not hope for a perfect correlation in our final regression. Measurement error, for instance, or the imperfect approximation that we use for transport costs. It remains slightly unsatisfying, however, that the numbers that we use for our estimation are not necessarily an equilibrium outcome of the model. This is especially true in the class of economic geography models, where for certain parameters a distributed outcome is infeasible, and agglomeration the only stable solution.

This is why we estimate the same model in a different way in the next section. The procedure there takes the general equilibrium nature of the model seriously and allows all the relationships to hold. We will compare the outcomes of the two procedures and use them to judge the importance of keeping variables constant.

5.5.2 General equilibrium estimation

In most estimations of the economic geography model, part of the model is kept constant. This leads to outcomes that are interesting, but not model-consistent. For instance, in our estimation in the previous section, the variable n and by extension E are kept constant, even though they are functions of other variables in the model. In this section we will attempt to estimate parameters while maintaining general equilibrium.

The properties of our model are governed by a small number of key parameters, the most important of which are transport costs and the elasticity of substitution. As we saw in paragraph 5.2.3 above, it is not always possible to separately identify these two parameters from the data, if flows of trade are the only information. If both parameters affect these flows in a similar way, it is impossible to separate out their influence. However, when data on wages is also available, this difficulty can be overcome. We will show this in the following paragraph. After that, we will estimate the key parameters in our model.

Identification

Bröcker (1999) is not able to separately identify transport costs parameter τ and elasticity of substitution σ . This is intuitive: a higher elasticity of substitution means that goods have become less differentiated. Given that products from another region are relatively expensive because of transport costs, if they become less 'special', their consumption will decrease. Thus an increase in either transport costs or the elasticity of substitution will have the same effect, hence it is not possible to identify the parameters separately.

Consider the model of the previous section and assume for a moment that there are only two regions. The value of the flow of trade from region 2 to region 1 will then equal

$$D_{21} = \frac{n_2 p_2^{1-\sigma} (T_{12})^{1-\sigma} E_1}{n_1 p_1^{1-\sigma} + n_2 p_2^{1-\sigma} (T_{12})^{1-\sigma}} \quad (5.22)$$

(this is a version of equation 5.10). Now assume that $\mu = 1$, so there are only tradable goods, and $\alpha = 1$ so that the production process only uses labor as an input. We can then write $n_r = L_r / \sigma F$ and $p_r = w_r \cdot \sigma / (\sigma - 1)$. Equation (5.22) reduces to

$$D_{21} = \frac{w_1 L_1}{\frac{L_1}{L_2} \left(\frac{w_1}{w_2} \right)^{1-\sigma} (T_{12})^{\sigma-1} + 1} \quad (5.23)$$

The above equation shows that if there is no data available about the level of wages in the different regions, w_r , then separate identification of T_{rs} and

σ is impossible, as the variables only show up in a joint term. The analysis by Bröcker (1999) was carried out without data on wages, so that implicitly all wages were assumed equal. This explains why there were identification problems in that study.

If we revoke some of our assumptions, identifiability is less clear. For instance, if we allow the share of intermediate products $1 - \alpha$ to be positive, do we still need variation in wages for separate identification of the parameters? We can no longer tell from the simple expression for D_{21} , as we need to include the more complicated terms buried in the p 's and n 's of (5.22).

To find out about the possible separate identification of σ and τ we run the following simulation experiment: for a model with three regions, we generate 2000 random distances and labor supplies. This sample is split in two halves. For the first half, we generate random (different) wage levels for the three regions; for the second half, all wages are assumed equal to one.

For all 2000 sets of data we compute the equilibrium flows of trade as described above.¹⁹ Because there are three regions, this is matrix of nine flows (both intra- and inter-region). We rearrange these flows into a vector T .

It is possible to compute the derivative of this vector T to changes in σ and τ . We do this numerically, so that

$$T_{\sigma} \doteq \frac{T(\sigma + d\sigma) - T(\sigma)}{d\sigma}$$

where $d\sigma$ is a small number, on the order of $0.01 \cdot \sigma$. The operation for τ is similar. These two derivatives (two elements of \mathcal{R}^9) are compared by computing the angle ϕ between them. From linear algebra, there holds that

$$\cos(\phi) = \frac{\langle T_{\sigma}, T_{\tau} \rangle}{\sqrt{\langle T_{\sigma}, T_{\sigma} \rangle} \sqrt{\langle T_{\tau}, T_{\tau} \rangle}}$$

with $\langle \dots \rangle$ the inner product of two vectors. An angle close to zero means that there exists colinearity between the two derivatives, which could be a sign of identification trouble. Ideally, the angle between the two derivatives should be 90° , indicating that the two regressors are orthogonal and perfectly identified.

Of course the derivative is only a first-order approximation to the non-linear problem that we are trying to solve. However, the numerical derivative that we examine in this experiment will likely also be used by the solver-routine that computes the estimation. If we find closely correlated

¹⁹We use $\sigma = 6$, $\tau = 0.08$, $F = 1$, and $\alpha = 0.6$. These are fairly typical values. The results of the experiment are robust to variations in this set. When Z is a random variable with a standard normal distribution, we draw $w = 6 + 2X$, $L_s = 10 + 2X$. Each region is given coordinates (x, y) which both are draws from X . Distances are then computed by Pythagoras' theorem.

regressors in this case, it will certainly make estimation very difficult and lead to imprecise estimates.

The results of our simulation experiment are in figure 5.4. The two panels of the figure are histograms of the distribution of the angle between the two regressors. If we take as a rule of thumb that identification is feasible if the angle is larger than 20° , we see that most cases in the bottom panel (where wages vary) pose no problem. In the top panel, we see that identification of both σ and τ is a more rare event. In either case, identification is not clear cut—a bad realization of the regressors can throw a wrench in the works at any time.

Equilibrium Estimation

In this section, we will attempt to estimate the parameters of the model using a method in which full equilibrium is maintained. For this estimation, we use largely the same dataset as in section 5.5.1. The sources of this data are described in an appendix on page 144. As before, we take the 48 contiguous US states as our sample; of these states, we take as exogenous the matrix of distances between them, and a matrix of dummies indicating whether a border exists between two states. Our sample year is 1997: we use the amount of employment in each state and average state wages for that year as exogenous inputs.

The model is described in section 5.3. Given the exogenous variables above and values for its parameters, we can numerically compute a solution to the model in which equilibrium is attained on each market. That is, firms make zero profit, demand and supply are equalized for each product and prizes reflect a marked up average of the costs of labor and intermediate products. From this equilibrium we can compute a flow of trade between every pair of regions.

Our estimation procedure finds the parameters that make these flows of trade as close as possible to the observed data. It works as follows: we compute the outcome of the model as described above, given a set of values for the model's parameters. From this outcome, we take the matrix of trade flows and compare it to the matrix of actual flows of trade, available from our dataset. We then evaluate how close the two are (a discussion of how we measure closeness is below) and repeat the procedure for another set of parameters. This way, we search for the parameters that generate a flow of trade that is closest to the actual data. Though we are not able to exactly replicate the actual flow of trade in our model, we do find the model that comes closest while still being consistent.

There are a number of problems with the above procedure. First of all, our matrix of trade flows is incomplete: some of the state-pairs have an unknown flow of trade between them. We leave these empty observations out of the sample.

Secondly, the number of parameters that has to be specified is reasonable but too large to estimate all using this procedure.²⁰ In order to compute the equilibrium, we need to specify values for the elasticity of substitution σ , labor share α , industrial consumption share μ , firm fixed costs F and the two transport costs parameters τ , and the coefficient for the border dummy, δ . Transport costs are computed as

$$T_{i,j} = \exp(\tau \text{dist}_{i,j} + \delta \text{border}_{i,j}). \quad (5.24)$$

Our estimation consists of the numerical minimization of the objective function over all possible parameters. In order to cut down on computer time, we will formulate the model so that only three parameters have to be estimated.

First of all, we recognize that the fixed cost parameter F does not influence the outcome. The amount of fixed costs that is involved with starting a firm will influence the number of firms n linearly, as can be seen from (5.7). However, since fixed costs are the same in every region, the relative number of firms between regions stays the same and the pattern of trade is not affected. We set $F = 0.01$, a value that leads to numerically efficient values of n .

Next, we specify our objective function so that the absolute value of the flows of trade is not relevant. This will allow us to keep the share parameter μ out of the estimation. Finally, we get the share of labor in the production function from actual observation rather than this estimation. Table 5.7 gives the factor shares in production from 1997. We use our assumption that only labor and intermediate goods are used and compute labor's relative share from that, and set $\alpha = 0.605$.

We are left with three parameters to estimate: σ , τ and δ . The object of the estimation is to find the values for these parameters that give a model solution in which the simulated flow of trade is as close as possible to the observed flow of trade. We operationalize this criterion in the following way, dealing at once with missing observations and matters of scale: we start with the matrix of trade flows that is generated by the model, and which has the same dimensions as the flow of trade-matrix from the data (48×48). From this matrix, we set all entries for which the corresponding entry in the data-matrix is missing, to zero. In the data-matrix, these missing entries are also set to zero. We then erase all the diagonal entries, which are the within-state shipments. In the previous section, we found that these flows of trade behave differently from interstate flows, so that the dummy variable *own* was necessary. Leaving these observations out frees us from having to measure a parameter for the effect of the *own*-dummy, leaving us with only three parameters to estimate.

²⁰While making point-estimates would be possible, the Monte Carlo error analysis below would become rather involved.

Next, both matrices are normalized so that the row-totals are equal to one. Because the matrices are drawn up so that each row contains all the imports of a particular region, this converts the entries to *shares of observed import*. That is, we express each flow of trade as a fraction of the total measured imports of the receiving region. Both matrices have a number of zeroes in them, which indicate that some entries were erased or not observed; as we shall see, these zeroes do not hinder the estimation process.

We can now formulate the criterion for closeness between the two matrices. Our estimation will minimize the sum-of-squares criterion

$$C(\sigma, \tau, \delta) = \sum_{i=1 \dots 48} \sum_{j=1 \dots 48} (D_{i,j} - P_{i,j}(\sigma, \tau, \delta))^2, \quad (5.25)$$

where D is the data matrix of observed import shares and P the model's projection of the same matrix. Notice that missing entries in D and P , which have both been set to zero, do not add anything to C .

By minimizing the difference between the two matrices of shares, we are able to leave the consumption share parameter μ unspecified; its values does not influence the outcome of P as it affects both the single flows of trade and total imports in the same (multiplicative) way. Also, we bring balance between the data on imports from large states and data on small states; by using shares instead of dollar values, large states do not dominate the estimation.

The results of the estimation are in table 5.8. The values for the parameters in this table minimize, in terms of the criterion in (5.25), the difference between the projected import share matrix and the observed version. We have used an adapted Monte Carlo method to arrive at an estimate for the standard error of the coefficients. For this method, we use the projected import share matrix $P(\sigma^*, \tau^*, \delta^*)$ and the set of projection errors $[D - P(\sigma^*, \tau^*, \delta^*)]_{i,j}$ for observed pairs (i, j) .

In a normal Monte Carlo procedure, we would randomly draw errors from this set and add them to P to generate a new dependent variable. Using this new dependent variable, we would re-estimate the parameters in a replication of the original estimation procedure. The variation in a set of about 200 replications of parameters re-estimated in this way would be an indication of the standard error of the original estimate.

However, in this case we run into trouble using the above procedure: it is quite possible that some of the newly generated dependent variables contain a number of negative entries, where a large negative error was added to a small initial projection. Beside the fact that negative trade would never be observed in practice, our model can not hope to replicate this observation as its trade is positive by definition. One could choose to leave the negative observations out of the replicated sample, but this leads to another problem: because small observations have a much higher probability

of being left out, the new estimation is biased. This is observed in the fact that the average of the Monte Carlo replications is quite different from the original parameter estimate.

We therefore modify our Monte Carlo procedure as follows: firstly, we divide the set of projection errors by 40 before embarking on the original procedure. These smaller errors do not lead to negative dependent variables anymore. We observe the variation in our replicated parameter estimates in the usual way and then rescale it to take account of our initial attenuation of the errors. This is done by multiplying the observed standard errors by 40. It is these standard errors that are reported in table 5.8. We also plot the distribution of the Monte Carlo parameter estimates *before* correcting for the attenuated errors; these plots are in figures 5.5 through 5.7. From these figures, we notice that even though the estimation procedure is nonlinear, the distribution of the errors appears close to normal.

Our dependent variable in this estimation has been observed import shares. It is a natural variable to choose, as our model delivers these shares from the simple formula

$$\phi_{rs} = n_s (T_{rs} p_s^s)^{1-\sigma} G_r^{\sigma-1} \quad (5.26)$$

with ϕ_{rs} the share of region s goods in region r 's imports. In principle, however, we could also have looked at *export* shares. These shares, ψ_{rs} , follow from our model as well:

$$\psi_{rs} = \frac{E_r n_s (T_{rs} p_s^s)^{1-\sigma} G_r^{\sigma-1}}{\sum_{r=1}^N (E_r n_s (T_{rs} p_s^s)^{1-\sigma} G_r^{\sigma-1})}$$

but the expression is much more involved. Estimation on the basis of ψ_{rs} can be carried out through simple rescaling, however, if we use observed total exports and total imports for each region and use the restriction that these totals must hold for the projected trade matrix as well.²¹

The results of this second estimation are in the second column of table 5.8. Notice that while the parameter estimates appear to be robust to this transformation, the standard errors have increased. This can be explained as follows: in the first estimation we only use the model to derive import shares, which are a simple function of prices and transport costs (see formula 5.26). There, we do not use the model's information on the relative size of regions. In the second estimation, we have made an added assumption that involves the (economic) size of regions; this assumption is embodied in our transformation of import shares into export shares, and the added assumption has decreased the model's fit on the data. For our

²¹This way, we forego the use of expenditure variables E_r from our model. In a projection, the import shares are computed as in (5.26), and rows are rescaled to sum to total observed imports. Then, columns are rescaled to sum to one. This matrix is compared to the actual matrix of export shares.

analysis of counterfactuals below, we therefore use the parameters from the first estimation (import shares) only.

The table with results also reports the R^2 of the estimation. This number is defined, as usual, to be a measure of explained variance. We compute the total sum of squares TSS as sum of squares of $D - \bar{D}$. The matrix \bar{D} is constructed as a matrix of ones and zeros, where the entries indicate if an observation is present at that position in D . Each entry is then divided by the sum of its column. The matrix \bar{D} is our best guess for D if no data is used; it is a generalized intercept of the estimation. We also compute the residual sum of squares, RSS , as the total sum of squared errors $D - P$. Then, $R^2 = 1 - RSS/TSS$.

We would like to be able to compare the results of this estimation to the results of section 5.5.1. Both approaches try to fit a model in which the flow of trade is explained by costs of transport and the economic size of regions, a relationship which we wrote down in formula (5.10). The two approaches differ, however, on their specification of the transport cost function. The estimation in the previous section used formula (5.15), where the log of trade is proportional to the log of distance. Therefore, the parameters in table 5.2 give $\partial \log(X_{ij})/\partial \log(\text{dist}_{ij})$ and $\partial \log(X_{ij})/\partial \text{border}_{ij}$. In this section, we specified the transport cost function (5.24), which is used in formula (5.10). Combining these two tells us that the parameters in table 5.8 give $\partial \log(X_{ij})/((1 - \sigma)\partial \text{dist}_{ij})$ and $\partial \log(X_{ij})/((1 - \sigma)\partial \text{border}_{ij})$. In table 5.8, we therefore include comparable parameter estimates. For these estimates, the distance coefficient is $\hat{\tau}$ multiplied by $(1 - \hat{\sigma})$ and by the average distance used, which is 997 miles.²² The border coefficient is $\hat{\delta}$ multiplied by $(1 - \hat{\sigma})$.

When we look at these comparable parameters, we should compare them to the outcomes of column US4 in table 5.2, which was estimated on the same sample. We see that when import shares are used, trade is more responsive to distance in the equilibrium estimation. When export shares are used, trade is less responsive to distance. The border parameter is greater in both equilibrium estimations.

Remoteness

With the parameter estimates in, we can use our model to take another look at the remoteness-effect that was described on page 117. The effect tells us that two regions trade more than a gravity-model would predict if they are relatively close together, and relatively far from the other regions. This follows from the fact that the gravity model fails to take into account a region's

²²The latter multiplication is necessary because the use of distance in one, and log of distance in the other regression. It follows from the rule that $\partial \log(x) = \partial x/x$, but is only an approximation.

other possibilities for trade, when looking at one particular trade relationship. For instance, the prediction of the gravity model for trade between the Netherlands and Belgium is the same, whether these two countries are part of continental Europe or whether they are twin islands in the middle of the Pacific. Our model does account for other possibilities through the price index of manufactured goods G , which is high for remote regions.

We use our estimated parameter values (import shares) to compute the price indices G for the states of Oregon and Washington. They are 40% and 51% above the national average, respectively, indicating that these two neighboring states both occupy remote positions in the USA. If only distance and size matter, we can predict the trade from Washington to Oregon with the simple gravity equation (5.12). Using our estimation results from table 5.1, we find that predicted trade would be \$ 2125 million. The Venables-model predicts a much higher trade, as expected: the prediction is \$ 6392 million. This is due to the *remoteness*-effect that was discussed above. Both models grossly underestimate the actual trade, incidentally, which is \$ 10301 million.

5.6 Evaluating counterfactuals

In this section, we use the parameter estimates of paragraph 5.5.2 above in our model to compute the effect of two changes in the economy. First, we look at the spatial impact of a fall in wages in Illinois, a simulation exercise that is similar to an experiment conducted by Hanson (1999). We then use our model to compute the effect of a new interstate highway that causes a fall in transport costs.

5.6.1 A wage shock

In Hanson (1999), the author estimates the coefficients of a model where the level of wages in a region depends on economic activity in the surrounding regions, with distance reducing the influence. Hanson then studies the effects of a reduction in personal income in the state of Illinois, one of the regressors, on wages around the country.

In this section, we conduct an experiment that is similar in spirit. Since wages in our model are fixed and personal income is computed endogenously, we reverse the shock: wages in the state of Illinois are decreased by ten percent, and we observe the effects of this change on demand in every state. We pick Illinois for the same reasons as Hanson: its large economic size, which gives us a sizeable effect, and its central location.²³ Also, we can compare the range of this shock to that of the shock in income.

²³In the map in figure 5.1 on page 121, Illinois is indicated with IL.

How does a reduction in wages in one state affect the demand for industrial goods locally, and in other states? In the afflicted state, prices drop as they are a markup over, amongst other, wages (formula 5.6). The cheaper goods from this state reduce the price index of industrial goods in every other state, but more so if the other state is close to Illinois. This causes the other states to lower their prices as well (formula 5.6 again) which sets off another round of falling price indices around the country. This process continues until convergence.

In the state where wages are lowered, two other effects come into play. The number of firms in the state decreases (see formula 5.7, where w goes down more than G). And, with state spending in this model directly related to the level of wages, the demand that the state itself exerts drops with 10%. Both effects depress the demand for goods from this state, but they are offset by an increase in demand from other states. This increase is the result of the lower price of industrial goods.

The change in prices and price indices leads to a change in the demand for each state. The effect is greatest for the state in which wages went down, because the price change is greatest in this region. Other states face two opposite effects: because all prices are down, a wealth effect causes an increase in demand for all states. However, their prices have all increased relative to Illinois which deflects demand away from the other states. We noted that states close to Illinois saw the greatest fall in prices, but they also suffer the most from Illinois' lower import demand.

Table 5.9 summarizes the results of this simulation. We look at the effects on several variables in Illinois itself, the neighboring states, a group of states at 'average' distance and two faraway states.

The total demand for industrial goods is up in Illinois, and down everywhere else. For states close to Illinois, the fall in that state's imports plays a major role in the drop in demand. For states further away, the deflected demand due to lower Illinois prices is the main cause. This can be seen from the second column in table 5.9, which gives the change in demand when we keep the imports in Illinois constant, and just look at the effects of the changed prices. This makes a big difference for states close to Illinois, who were able to lower their prices due to the cheaper inputs from their neighbor.

The number of firms falls in Illinois, which is a result of our zero-profit condition (5.7). Everywhere else, this number increases due to the fall in costs of input. Once again, the neighbors see the biggest increase.

The price index of industrial goods is down in every state. In Illinois itself, the change is relatively small due to the lower number of firms in that state (leaving it with less varieties). For other states, the fall in G with constant wages leads to a fall in prices, as can be seen in column 6.

Our conclusion that only the state with lower wages gains in demand obscures the fact that some states do see an increase in demand from cer-

tain trading partners. For instance, the final column of table 5.9 shows the change in demand from Florida, a state that is relatively far away from the Midwest. We see that Florida increases its imports from (cheaper) Illinois sharply after the shock, offsetting the change with a decrease in imports from other states. However, it also starts importing more from the states *around* Illinois, who are also able to lower their prices. So, these neighboring states see an increase in demand from trading partners that are far away, but the net effect remains negative.

What is remarkable about the effects of the shock, is that they do not seem to decline with distance much. States that are far away see a change in demand that is not smaller than the change experienced by intermediate states. Neighbors have a slightly different experience, but in the end see a similar drop in total export demand. This stands in marked contrast with the results of Hanson (1999), who measured that the effects of a ten percent drop in expenditures in Illinois reached no further than 900 kilometers. We must note that a drop in expenditures in his model leaves prices unchanged. The change in prices was the driving force behind the effects in our model.

We also note that the stability of the model owes a great deal to the assumption that the number of workers in each state remains constant. Presumably, with the changes in wages and the entrance and exit of firms, people would enter or exit the labor market. We have abstracted from this effect, possibly underestimating the results of the shock.

5.6.2 A fall in transport costs

We now experiment with a fall in the costs of transport between two states. Once again our change takes place in the Midwest. Assume the construction of a new highway between Illinois and Indiana, two neighboring states, that reduces the costs of transport between the two states by fifty percent.²⁴ For each state, it also reduces the internal costs of transport by half. We proxy for costs of transport by distance, so in effect we reduce the three relevant distances in our model.

Lowering the costs of transport between the two states affects either's price index of industrial goods: for firms and consumers in Indiana, products from their own state and those from Illinois become less expensive. Firms and consumers in Illinois see their own products and those from Indiana become cheaper. This means that the price of inputs drops in both states, allowing firms to lower their prices and new firms to enter.

We do not model the effect that the new highway has on traffic passing *through* the two states, on its way from Ohio to Missouri, for instance, or the effect on traffic reaching either state from the outside. This means that

²⁴In the map in figure 5.1, Indiana is IN.

the only thing that the other states notice from the new road is the change in Illinois' and Indiana's prices. This change in turn lowers their average price for inputs and their prices for the final good, setting off another round of price drops. In the end, the new prices cause a shift in demand.

How exactly this shift plays out for each state can be read from table 5.10. We see that the two states which received the shock see an increase in demand. Their neighboring states are able to lower their prices, which would have earned them additional business, if it had not been for the competition from Illinois and Indiana. We see that the neighbors see little change in the demand from faraway Florida (column 6), but suffer an overall decline in demand. This decline is the result of Indiana and Illinois spending turning inward and to each other, away from the neighbors.

Intermediate states notice that their competitiveness *vis à vis* the Midwest has gone down, which causes demand to fall. The same holds for states that are far away, although we see that this time, the effect does become smaller with distance.

We note that the way we have modelled the effects of a new road is rather coarse: only the two affected states notice the change in their costs of transport. A more realistic modelling of the change would have taken the effects on other entries in the distance matrix into account. Such a computation is beyond the scope of this chapter, however, as it requires a complex model of transportation. In this section, we merely aim to illustrate the effects of a change in transport costs, and the mechanisms involved. In chapter 6 below, we do employ the results of a complete model of transport costs when evaluating the effects of a new railway link in the Netherlands.

5.7 Conclusions

In this chapter, we have estimated the parameters of an economic geography model in two ways: one method, previously employed by Redding and Venables (2001), assumes certain parts of the model constant and measures the correlation between market access and wage. A second method finds the model's general equilibrium solution that best reproduces the data. Both methods have been applied to a dataset that covers the US states in 1997. The dataset is described below, in appendix 5.A.

We find that the model gives a reasonable description of the trade between US states. The 'foreign' regressors in the first method performed a bit worse than they did on a worldwide dataset. The general equilibrium estimates gave reasonable values for the model's parameters, but their standard errors leave room for doubt.

When we simulate counterfactuals with our calibrated model, we see that the effects of a shock reverberate throughout the US. The calibrated model allows us to track the effects through space, which makes it a useful

tool for regional policy evaluation.

5.A Data

The dataset used in this paper concern the 51 US states in the year 1997. The complete set can be found on the internet, at <http://knaap.com/gdata>. Data was collected from a variety of sources. We list them here, together with a download address.

- **Gross State Product.** Supplied by the Bureau of Economic Analysis, US Department of Commerce. The June 7, 1999 edition of the current-dollar GSP estimates were used. Available at <http://www.bea.doc.gov/bea/regional/gsp/>.
- **Employment.** Total nonfarm employment per state, from the Bureau of Labor Statistics. Available at <http://146.142.4.24/cgi-bin/srgate>. Request series SASxx0000000001, where xx is the state number.
- **Wages.** Average annual pay for 1997, from the Bureau of Labor Statistics, December 15, 1999 edition. Available at <http://stats.bls.gov/news.release/annpay.t01.htm>.
- **Interstate flow of commodities.** Bureau of Transportation Statistics 1997 State-to-state commodity flows in millions of US\$. Available at <http://www.bts.gov/cfs/cfs97od.html>.
- **Distance between states.** Duncan Black kindly supplied a computer file with the latitude and longitude of each US county. I averaged these into state coordinates, weighing them with county employment. The distance between two states is then computed in miles using the great circle formula. For the distance within a state, I obtained the state area A_i and computed the quasi-radius as $\sqrt{A_i/\pi}$. This number approximates the average distance travelled within a state. State areas may be found at http://www.census.gov/population/censusdata/90den_stco.txt.
- **Factor shares.** The numbers in table 5.7 come from the BLS website at <http://www.bls.gov> and have ID numbers MPU300013 through MPU300017.
- **Weather data.** National Climatic Data Center, Asheville, NC. Tables can be accessed via <http://ols.nndc.noaa.gov/plolstore/plsql/olstore.prodspecific?prodnum=C00095-PUB-A0001>.

- **Mining.** Data from Smith (1997) available at http://minerals.er.usgs.gov/minerals/pubs/commodity/statistical_summary/871497.pdf

$\log(X_{r,s})$	World 1	US 1	World 2	US 2	World 3	US 3	US 4	US 5
Obs.	10100	2601	8079	2201	10100	2601	2091	2042
Year	1994	1997	1994	1997	1994	1997	1997	1997
Estimation	OLS	OLS	OLS	OLS	Tobit	Tobit	OLS	OLS
$\log(\text{dist}_{r,s})$	-1.538 [-0.041]	-1.181 [-0.056]	-1.353 [-0.032]	-1.044 [-0.025]	-1.738 [-0.043]	-1.330 [-0.063]	-0.983 [-0.024]	-0.987 [-0.023]
$\text{bord}_{r,s}$	0.976 [0.195]	0.774 [0.126]	1.042 [0.141]	0.492 [0.052]	0.917 [0.179]	0.658 [0.140]	0.554 [0.049]	0.554 [0.048]
$\text{own}_{r,s}$	-	2.462 [0.232]	-	2.210 [0.095]	-	2.335 [0.257]	2.232 [0.090]	-
R^2	0.789	0.779	0.786	0.921	-	-	0.924	0.921
$\log L$	-	-	-	-	-20306	-4422	-	-

World columns are from Table 1 in Redding and Venables (2001), US columns are own computations. Estimation 1 uses the full sample, including zeros. Estimation 2 uses a censored sample, from which the zeros have been eliminated. Estimation 3 again uses the full sample, taking care of the left-censored observations by using a Tobit estimation. Estimation 4 uses only the contiguous states, eliminating Hawaii and Alaska, as well as the District of Columbia. Estimation 5, finally, uses that sample without the within-state flows.

Table 5.2: Panel estimates for the gravity trade equation

$\log(w_r)$	World	US	World	US	World	US
Obs.	101	48	101	48	101	47
Year	1996	1999	1996	1999	1996	1999
$\log(\text{FMA}_r)$	0.476 [0.076]	0.133 [0.082]	-	-	0.316 [0.088]	0.066 [0.044]
$\log(\text{MA}_r)$	-	-	0.479 [0.063]	0.257 [0.029]	-	-
$\log(\text{DMA}_r)$	-	-	-	-	0.141 [0.059]	0.119 [0.014]
R^2	0.346	0.079	0.610	0.601	0.584	0.613
Moran's I		0.197		0.317		0.404
$1 - F(I)$		0.0138		0.0006		0.0000

World columns are from Table 2 in Redding and Venables (2001), *US* columns are own computations. The dependent variable in World columns is GDP per capita. Bootstrapped standard errors are in parentheses (200 replications). First stage estimation is Tobit for the World columns, US 4 (see table 5.2) for US columns. Moran's I is computed on the residuals of the estimation, using a matrix of border-dummies as a weighing matrix. On the line below is the position of the statistic in a bootstrapped distribution function (100,000 replications).

Table 5.3: Market Access and wage levels

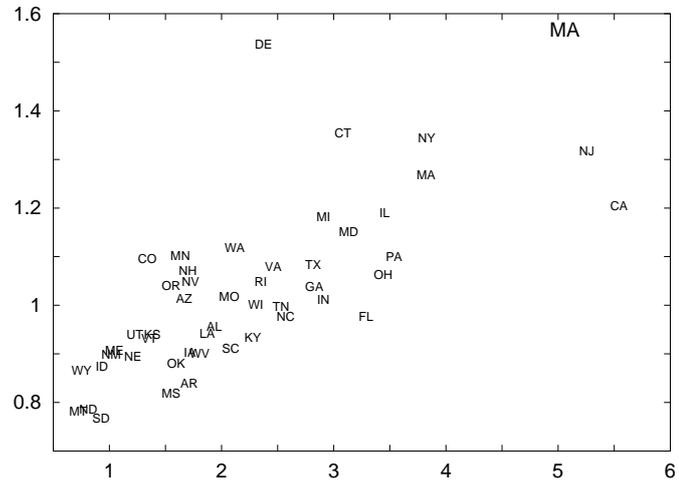


Figure 5.2: Predicted Market Access (horizontal, based on 1997 data) versus log wages (vertical, data from 1999, wages in ten thousands of dollars) for 48 states. MA regressors come from the US 4 estimation.

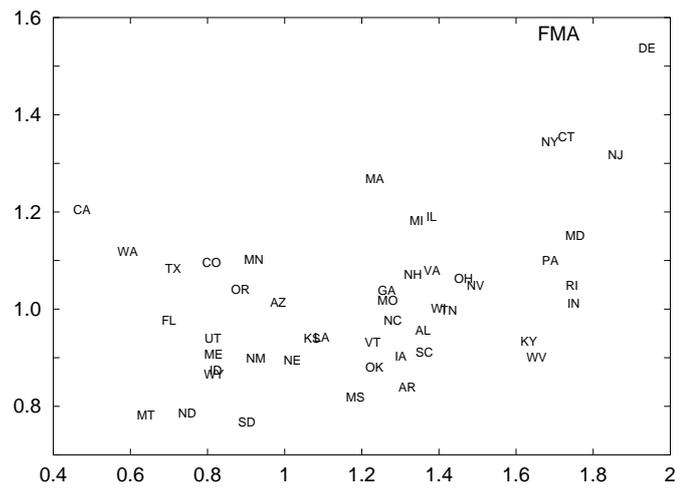


Figure 5.3: Predicted Foreign Market Access (horizontal) versus log wages (vertical) for 48 states. FMA regressors come from the US 4 estimation.

$\log(w_r)$	World	US	World	US
Obs.	101	48	101	48
Year	1996	1999	1996	1999
$\log(\text{FSA}_r)$	0.532 [0.114]	0.118 [0.082]	-	-
$\log(\text{SA}_r)$	-	-	0.345 [0.032]	0.229 [0.030]
R^2	0.377	0.075	0.687	0.542
Moran's I		0.217		0.322
$1 - F(I)$		0.0091		0.0006

World columns are from Table 9 in Redding and Venables (2001), *US* columns are own computations. See the note under table 5.3.

Table 5.4: Supplier Access and wage levels

$\log(w_r)$	US	US	US	US
Obs.	48	48	48	47
Year	1999	1999	1999	1999
$\log(\text{FMA}_r)$		0.042 [0.063]		0.044 [0.046]
$\log(\text{MA}_r)$			0.234 [0.041]	
$\log(\text{DMA}_r)$				0.112 [0.019]
<i>nrmcdd</i>	-0.103 [0.044]	-0.103 [0.049]	-0.049 [0.036]	-0.063 [0.032]
<i>nrmhdd</i>	-0.010 [0.016]	-0.010 [0.016]	0.013 [0.013]	0.009 [0.011]
minerals	0.022 [0.004]	0.021 [0.004]	0.006 [0.005]	0.005 [0.005]
port	0.130 [0.035]	0.130 [0.036]	0.070 [0.029]	0.044 [0.031]
R^2	0.545	0.551	0.756	0.776
Moran's I	0.164	0.128	0.230	0.205
$1 - F(I)$	0.0281	0.0586	0.0061	0.0127

Standard errors in parentheses. Except for the first column, these errors come from bootstrap methods (200 replications). First stage estimation for market access variables is US 4 (see table 5.2). Moran's I is computed on residuals, using a matrix of border-dummies. The position of the statistic in a bootstrapped distribution function is indicated below (100,000 replications).

Table 5.5: Exogenous amenities, Market Access and wage levels

$\log(w_r)$	US	US	US	US
Obs.	48	48	48	48
Year	1999	1999	1999	1999
$\log(\text{FMA}_r)$	0.169 [0.101]		0.076 [0.079]	
$\log(\text{MA}_r)$		0.232 [0.107]		0.144 [0.111]
<i>nrmcdd</i>			-0.103 [0.057]	-0.070 [0.046]
<i>nrmhdd</i>			-0.011 [0.019]	0.004 [0.017]
minerals			0.020 [0.005]	0.012 [0.009]
port			0.129 [0.036]	0.093 [0.037]
R^2	0.073	0.599	0.547	0.724
Moran's I	0.163	0.291	0.105	0.145
$1 - F(I)$	0.0299	0.0013	0.0908	0.0423

Instrumental variables estimation. In the first two columns, instruments are the distance from New York and the distance from Los Angeles. In the third and fourth column, the four exogenous regressors are added to the set. Standard errors come from bootstrap methods (200 replications). First stage estimation for market access variables is US 4 (see table 5.2). Moran's I is computed on residuals, using a matrix of border-dummies. The position of the statistic in a bootstrapped distribution function is indicated below (100,000 replications).

Table 5.6: Instrumental variables estimation

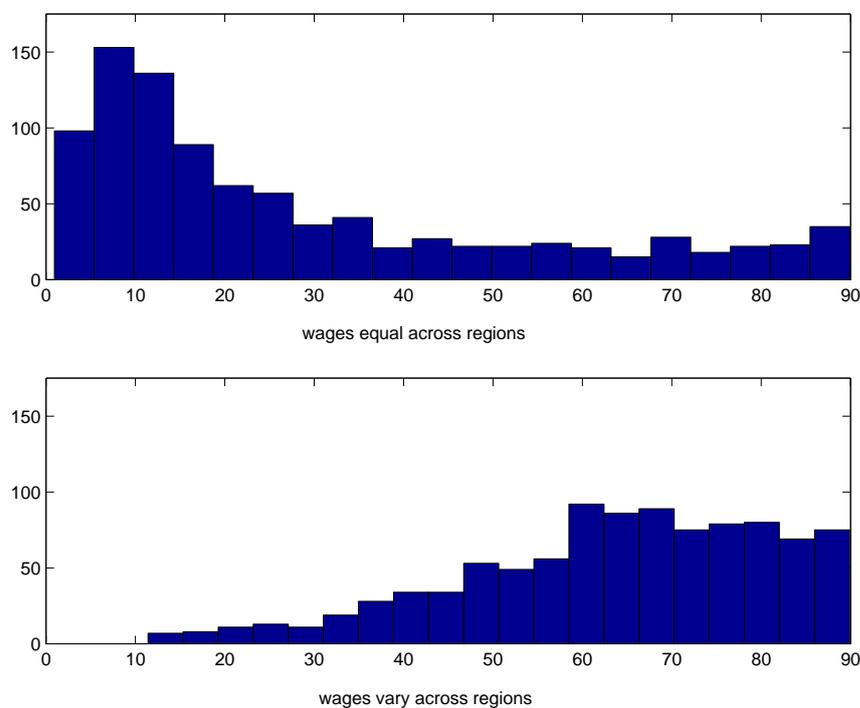


Figure 5.4: Distribution of the angle between the two vectors $dT/d\sigma$ and $dT/d\tau$ where T is the 9 by 1 vector of equilibrium trade flows between three regions. The top panel has wages fixed at one, the bottom panel has a random variation in wages between regions. An angle of 90° implies orthogonality of the two regressors, and an angle of 0° implies perfect collinearity.

	1997 factor shares		
	costs (\$ billions)	share (total)	share(relative)
Labor	934.3	37.8%	60.5%
Capital	516.7	20.9%	
Energy	63.8	2.6%	
Materials	610.5	24.7%	39.5%
Business Services	346.4	14.0%	

Data come from the Bureau of Labor Statistics and give the cost of different factors for the manufacturing sector in 1997. See also the appendix on data on page 144.

Table 5.7: Factor shares in US production, 1997

$P(\sigma, \tau, \delta)$	Import		Export	
Obs.	2043		2043	
Year	1997		1997	
σ	4.110	[1.066]	5.120	[2.499]
τ	$5.124 \cdot 10^{-4}$	$[2.199 \cdot 10^{-4}]$	$1.332 \cdot 10^{-4}$	$[8.957 \cdot 10^{-5}]$
δ	-0.235	[0.106]	-0.300	[0.179]
Parameters comparable to those in table 5.2.				
$\log(\text{dist}_{rs})$	-1.616		-0.547	
bord_{rs}	0.731		1.234	
R^2	0.725		0.780	

Estimation results from general equilibrium estimation as outlined in section 5.5.2. σ is the elasticity of substitution, τ and δ come from the transport cost function $T_{i,j} = \exp(\tau \text{dist}_{i,j} + \delta \text{border}_{i,j})$. Standard errors are Monte Carlo estimates based on 281 replications.

Table 5.8: General equilibrium estimation results

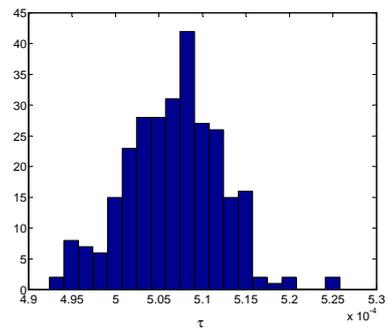
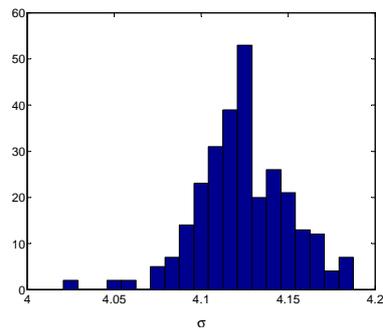


Figure 5.5: MC distribution of $\hat{\sigma}$, 'export' estimation, 1/40 errors

Figure 5.6: MC distribution of $\hat{\tau}$, 'export' estimation, 1/40 errors

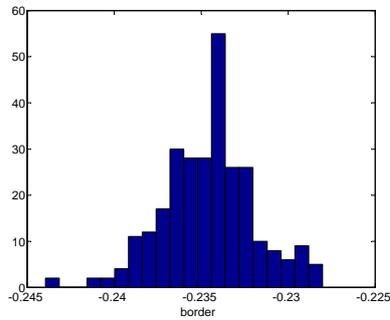


Figure 5.7: MC distribution of $\hat{\delta}$, 'export' estimation, 1/40 errors

State	Demand change	Change w/o E	Change in n	Change in G	Change in p	Demand from FL
IL	14.6%	15.4%	-3.8%	-0.8%	-6.4%	16.5%
<i>Neighbors</i>						
IN	-1.8%	-0.4%	0.4%	-1.1%	-0.4%	0.3%
IA	-1.7%	-0.4%	0.5%	-1.2%	-0.5%	0.4%
WI	-1.7%	-0.3%	0.5%	-1.3%	-0.5%	0.6%
MO	-1.6%	-0.5%	0.4%	-1.1%	-0.4%	0.2%
<i>Intermediate states</i>						
PA	-1.5%	-1.1%	0.2%	-0.5%	-0.2%	-0.7%
CO	-1.1%	-0.9%	0.2%	-0.6%	-0.2%	-0.5%
TX	-1.3%	-1.0%	0.2%	-0.6%	-0.2%	-0.6%
<i>Faraway states</i>						
CA	-0.9%	-0.8%	0.2%	-0.4%	-0.2%	-0.8%
ME	-1.5%	-1.1%	0.2%	-0.5%	-0.2%	-0.6%

The first column indicates the state. The change in demand that each state experiences is in the second column; the third column shows what the change would have been without the accompanying drop in Illinois expenditures. Changes in the number of firms n , price index G and price p are next. The final column gives the change in demand from Florida for each state.

Table 5.9: Effects of a 10% decrease in Illinois wages

State	Demand change	Change in n	Change in G	Change in p	Demand from FL
IL	2.3%	0.3%	-0.8%	-0.3%	1.09%
IN	2.4%	0.3%	-0.7%	-0.3%	1.06%
<i>Neighbors</i>					
IA	-0.4%	0.0%	-0.1%	0.0%	0.01%
WI	-0.4%	0.0%	-0.1%	0.0%	0.02%
MO	-0.4%	0.0%	-0.1%	0.0%	-0.01%
KY	-0.4%	0.0%	-0.1%	0.0%	0.00%
MI	-0.4%	0.0%	-0.1%	0.0%	-0.01%
OH	-0.3%	0.0%	-0.1%	0.0%	-0.03%
<i>Intermediate states</i>					
GA	-0.2%	0.0%	-0.1%	0.0%	-0.08%
CO	-0.2%	0.0%	-0.1%	0.0%	-0.07%
TX	-0.2%	0.0%	-0.1%	0.0%	-0.07%
<i>Faraway states</i>					
CA	-0.1%	0.0%	0.0%	0.0%	-0.10%
ME	-0.2%	0.1%	-0.1%	0.0%	-0.07%

The first column indicates the state. The change in demand that each state experiences is in the second column; Changes in the number of firms n , price index G and price p are next. The final column gives the change in demand from Florida for each state.

Table 5.10: Effects of a 50% decrease in of transport costs between Illinois and Indiana