

DOCUMENT DE TRAVAIL / WORKING PAPER

No. 2017-06



**The world is not yet flat: Transport costs
matter!**

**Kristian Behrens, W. Mark Brown, *et*
Théophile Bougna**

Octobre 2017

The world is not yet flat: Transport costs matter!

Kristian Behrens, Université du Québec à Montréal, Canada; National Research University Higher School of Economics, Russie; et CEPR

W. Mark Brown, Statistics Canada (EAD), Canada

Théophile Bougna, Université du Québec à Montréal, Canada; et The World Bank, Washington DC, É.-U.

Document de travail No. 2017-06

Octobre 2017

Département des Sciences Économiques
Université du Québec à Montréal
Case postale 8888,
Succ. Centre-Ville
Montréal, (Québec), H3C 3P8, Canada
Courriel : brisson.lorraine@uqam.ca
Site web : <http://economie.esg.uqam.ca>

Les documents de travail contiennent souvent des travaux préliminaires ou partiels et sont circulés pour encourager et stimuler les discussions. Toute citation et référence à ces documents devrait tenir compte de leur caractère provisoire. Les opinions exprimées dans les documents de travail sont ceux de leurs auteurs et ne reflètent pas nécessairement ceux du département des sciences économiques ou de l'ESG.

Copyright (2017): Kristian Behrens, W. Mark Brown, et Théophile Bougna. De courts extraits de texte peuvent être cités et reproduits sans permission explicite à condition que la source soit référencée de manière appropriée.

The world is not yet flat: Transport costs matter!

Kristian Behrens* W. Mark Brown[†] Théophile Bougna[‡]

October 13, 2017

Abstract

Using micro-level commodity flow data and micro-geographic plant-level data, we construct industry-specific ad valorem trucking rate series and measures of geographic concentration to provide evidence on the relationship between transport costs and agglomeration. We find that low transport cost industries display significantly more geographic concentration in the cross-sectional dimension, and that falling transport costs agglomerate industries in the panel dimension. The effects are large: the fall in trucking rates between 1992 and 2008 implied a 20% increase in geographic concentration on average, all else equal.

Keywords: Ad valorem transport costs; geographic concentration; trucking rates; manufacturing industries.

JEL classification: R12; C23; L60.

**Corresponding author:* Université du Québec à Montréal, Canada; National Research University Higher School of Economics, Russia; CIRPÉE, Canada; and CEPR, UK. E-mail: behrens.kristian@uqam.ca

[†]Statistics Canada, Economic Analysis Division (EAD); Chief, Regional and Urban Economic Analysis. E-mail: mark.brown@canada.ca

[‡]Dept. of Economics, Université du Québec à Montréal, Canada; and DECRG, The World Bank, Washington DC. E-mail: tbougna@worldbank.org

1 Introduction

Transport and communication costs fell precipitously during the last century, leading many observers to posit that the world has ‘become flat’, that locational differences no longer matter. With the help of cheap transport and communication, business can be done almost anywhere, or so the story goes, leaving policymakers with the impression that we have entered ‘brave new frictionless’ world. But, if this were true, the costs of transporting goods should no longer influence firms’ location choices and, thereby, the spatial structure of economic activity. Why is then the tendency for economic activity to cluster in space still strong? Why do many industries still exhibit strong geographic patterns, including new entrants that should a priori face few locational constraints?

We address this seeming contradiction head-on by identifying the causal effect of transport costs on the geographic concentration of industries in Canada, both in the cross-sectional and the panel dimensions. Focussing on trucking—the main transport mode for freight in North America—we find that low transport cost industries display significantly more geographic concentration than high transport cost industries.¹ Decreasing transport costs also tend to agglomerate industries—especially at small spatial scales—and increase regional specialization. Our qualitative results, which are in line with those of [Krugman’s \(1991\)](#) celebrated ‘core-periphery’ model, hold up to a large variety of robustness checks and to instrumental variables estimations that deal with potential endogeneity concerns. Furthermore, the quantitative effects are large. In our preferred cross-sectional specification, an industry with twice as high ad valorem transport costs than another industry is on average 5.13% less geographically concentrated at 50 kilometers distance. In terms of changes over time, the fall in trucking rates between

1992 and 2008 implied a 20% increase in the geographic concentration of the average manufacturing industry, all else equal. In a nutshell, the world is not yet flat: Transport costs matter!

Assessing empirically the effect of transport costs on the geographic concentration of industries is important for several reasons. First, despite their fundamental theoretical role in spatial modeling, little is still known empirically on how transport costs drive the geographic structure of industries, especially at the regional level. Second, among the possible determinants of clustering, transport costs have been studied the least, much less than the ‘Marshallian’ forces such as input-output links, labor market pooling, and knowledge spillovers (see [Rosenthal and Strange, 2004](#); [Combes and Gobillon, 2015](#)). We thus have little quantitative evidence on the impact of those costs on geographic patterns. Third, changes in transport costs driven by, e.g., infrastructure investments, bear on the local composition of economic activity and inform us on how that composition may change (e.g., [Duranton, Morrow, and Turner, 2014](#)). These changes affect regional exposure to international trade shocks that have direct repercussions in local labor markets, a topic of great policy importance (e.g., [Autor, Dorn, and Hanson, 2013](#)).

Assessing empirically the effect of transport costs on the geographic concentration of industries is also a difficult task. First, we need fine measures of geographic concentration to look at the cross section, and enough time-series variation in those measures to look at changes. In this paper, we construct—for the first time to our knowledge—a long panel of continuous measures of geographic concentration, computed from micro-geographic plant-level data using the [Duranton and Overman \(2005\)](#) approach. Our micro-geographic data exhibit enough time-series variation so that they can be meaningfully used in a panel context.

Second, we need detailed measures of industry-specific transport costs and how they relate to output prices. We devote substantial effort to the calculation of *domestic ad valorem trucking rates* for 257 manufacturing industries. Most of the literature has looked at infrastructure as a proxy for transport costs (see [Redding and Turner, 2015](#)). By contrast, we build our trucking rates directly from micro-data files on truck shipments from Canada and the export values of goods. These ad valorem trucking rates can be extended to a 20 years time series using industry-specific output price indices.

Last, we need to deal with the possible endogeneity of ad valorem transport costs. For example, if the productivity gains from geographic concentration are passed on to consumers in the form of lower prices—which increases ad valorem trucking rates—the causality may run from agglomeration to transport costs and not the other way round. We deal with this issue by purging our ad valorem trucking costs of productivity effects. We also use ‘binning instruments’ in the cross section and U.S. industry price indices in the panel to construct external instruments for our trucking rate series. As predicted by theory, there is an upward bias in the OLS estimates of the transport cost coefficients, which reinforces our baseline results. Low transport cost industries are significantly more concentrated than high transport cost industries, and falling transport costs drive more geographic concentration.

The remainder of the paper is structured as follows. In **Section 2**, we briefly review the theoretical and empirical literatures on the effects of changes in transport costs on geographic concentration. We also develop a parsimonious conceptual framework of endogenously determined transport costs. **Section 3** documents our data, explains the calculation of our key variables, and provides a first look at geographic concentration and transport costs. In **Section 4**, we explain our empirical approach and discusses

various identification concerns. Our empirical results are summarized in **Section 5**, and **Section 6** concludes. We relegate many additional results and technical details to an extensive set of appendices.

2 Conceptual framework

2.1 Transport costs and geographic concentration: Theory

The geographic concentration of economic activity is determined by the balance between *agglomeration and dispersion forces*. In most regional models, agglomeration forces arise from a combination of transport costs and firm-level increasing returns to scale ([Krugman, 1991](#)), or from vertical buyer-supplier links ([Krugman and Venables, 1995](#)). Geographic concentration expands the size of the local market, thereby attracting firms that can serve a larger share of their final and intermediate demands—and purchase a larger share of their supplies—locally at lower cost. The cost savings are naturally larger the more expensive it is to ship goods. The strength of the agglomeration force thus *depends on the level of transport costs*: if transporting goods is cheap enough, it vanishes since firms no longer need to be close to final or intermediate customers and suppliers. All else equal, lower transport costs make economic activity more footloose and hence weaken the agglomeration forces.

How the equilibrium geographic concentration changes with transport costs depends on the interplay between agglomeration and dispersion forces. Different models of geographic concentration rely on different dispersion forces. In [Krugman’s \(1991\)](#) model, there is a spatially dispersed immobile demand that industries must serve. If transport costs are high, this immobile demand exerts a strong pull so that economic activity

disperses to follow it. The strength of this dispersion force *depends on the level of transport costs*. If the latter are low enough, the immobile demand can be served at virtually no cost from every location and thus no longer influences geographic patterns. Hence, in the Krugman model, both agglomeration and dispersion forces get weaker as transport costs fall. Yet, the dispersion force vanishes more quickly than the agglomeration forces, so that agglomeration of economic activity occurs for low levels of transport costs (see, e.g., [Fujita, Krugman, and Venables, 1999](#)). In [Helpman's \(1998\)](#) version of the model, the dispersed immobile demand is replaced with a fixed stock of regional housing. As economic activity concentrates in a location the price of housing increases, because of its inelastic supply, which tends to disperse economic activity since workers are better off when they can relax housing costs. Contrary to the Krugman model, the strength of this dispersion force *does not depend on the level of transport costs*. Hence, in the Helpman model, only the agglomeration force gets weaker as transport costs fall, so that dispersion of economic activity occurs for low levels of transport costs.²

2.2 Transport costs and geographic concentration: Evidence

While models of geographic concentration have proven useful to better understand spatial equilibrium patterns, different models generate different equilibria because they rely on different agglomeration and dispersion forces. These forces operate at varying spatial scales, and which one is the most relevant largely depends on the phenomenon one wants to analyze. While Helpman-type dispersion forces due to urban costs seem more relevant at the city level, Krugman-type dispersion forces due to a geographically dispersed demand seem more relevant at larger regional or national scales. The question of how transport costs—and changes in these costs—affect geographic patterns thus

becomes essentially an empirical one and depends on the spatial scale of the analysis.

Starting with the city level, the empirical literature has used carefully designed identification strategies to establish the causal links between infrastructure and city structure, composition, and growth. It has convincingly shown that more infrastructure, a proxy for decreasing transport costs for people and goods, causes: (i) population and economic activity to disperse within cities or to suburbanize (see, e.g., [Baum-Snow, 2007](#); [Garcia-López, Holl, and Viladecans-Marsal, 2015](#); [Baum-Snow, Brandt, Henderson, Turner, and Zhang, 2017](#)); (ii) cities to grow ([Duranton and Turner, 2012](#)); and (iii) cities to specialize in ‘heavy industries’, i.e., industries with a low value-to-weight ratio ([Duranton et al., 2014](#)). In a nutshell, lower transport costs tend to disperse economic activity within cities, in line with the Helpman model, but tend to concentrate specific industries across cities, in line with the Krugman model.

While the city-level evidence is rich, little is known on whether it translates to larger regional or national scales and how it applies to individual industries. This is firstly due to the fact that there is generally much less empirical work at those scales. Furthermore, much of the existing work is embedded in structural models that feature standard Krugman-type agglomeration forces. We have little evidence about the effect of transport costs on the regional organization of economic activity that does not come through the filter of a structural model that has the answer somewhat ‘baked-in’.³ We also have little evidence for individual industries. Since most agglomeration models have a macroeconomic flavor and focus on ‘aggregate economic activity’, they are silent on how transport costs affect the geographic concentration of individual industries. This is problematic since even when the overall geographic concentration of economic activity decreases (e.g., spreads across cities), individual industries can still agglomerate

(e.g., within cities) in response to changing transport costs. Ideally, we need to understand how transport costs affect the absolute geographic concentration of industries (i.e., spatial agglomeration), as well as their relative geographic concentration (i.e., regional specialization, conditional on overall concentration).

To link absolute and relative geographic patterns with industry-specific transport costs requires good estimates of the latter. Yet, there are surprisingly few empirical studies on geographic concentration that exploit direct measures of transport costs. Extant studies have mostly looked at changes in ‘market access’ broadly defined, stemming either from changes in international borders (e.g., the fall of the Iron Curtain; [Redding and Sturm, 2008](#); [Brülhart, Carrère, and Trionfetti, 2012](#)) or large-scale infrastructure investments (e.g., [Chandra and Thompson, 2000](#); [Baum-Snow, 2007](#); [Michaels, 2008](#); and [Duranton et al., 2014](#)). This literature relies on either natural experiments (e.g., the fall of the Iron Curtain) or historical instruments (e.g., planned interstate highway system) to obtain credible identification at the expense of indirect transport cost measures. One problem with market access or infrastructure is that they do not provide industry-specific variation in transport costs, thereby complicating an analysis of how individual industries tend to agglomerate or disperse and of how regional specialization changes.⁴ Another problem with infrastructure is that it does not exploit information on the value of the goods shipped. Yet, it is known since at least [Alchian and Allen \(1964\)](#) that transport costs are especially relevant in relation to the prices of the goods shipped. There is not much difference in shipping gold or gravel border-to-border across the city, but the difference will be crucial when shipping them coast-to-coast across the country. Hence, both per-unit transport costs and the prices of the goods shipped are required to investigate how *ad valorem* transport costs, and changes therein, affect the equilibrium

geographic concentration of industries. To our knowledge, this has not been done.

Our aim is to estimate industry-specific measures of ad valorem transport costs and to relate them to the absolute and relative geographic concentration of industries. To do so, we now lay out a simple framework that is useful for conceptualizing transport costs and for pinpointing various endogeneity concerns.

2.3 Modeling transport costs

Models of geographic concentration usually subsume transport costs by an *exogenous* parameter. Yet, in reality, transport costs are prices that are set to clear markets and as such reflect supply and demand conditions. While the assumption of exogenous transport costs is useful in some contexts, it masks a number of endogeneity concerns that are important to address in empirical work, especially when using ad valorem transport costs. To guide our subsequent analysis, we develop in Appendix A a simple two-region model based on [Behrens and Picard \(2011\)](#) and [Behrens and Brown \(2017\)](#) where ad valorem transport costs are endogenously determined by the interplay of manufacturing shippers (demand) and competitive carriers (supply). We show that the ad valorem transport costs for industry i between regions r and s can be expressed as

$$\tau_{rs}^i = \frac{1}{1 + M^{-1/\sigma_i}} \left[1 + \frac{\sigma_i}{\sigma_i - 1} \left(\frac{m_s^i}{p_r^{i,\text{prod}}} + \frac{2\gamma(\mathbf{Y}^c, d_{rs})}{p_r^{i,\text{prod}}} \right) \right], \quad (1)$$

which depend on the producer price $p_r^{i,\text{prod}}$ and the demand elasticity σ_i of industry i ; on the relative market size M of region r to region s ; on the cost m_s^i of producers in industry i in region s ; and on the carriers' cost function $\gamma(\mathbf{Y}^c, d_{rs})$. The latter depends on the distance d_{rs} of a one-way trip and a vector \mathbf{Y}^c of carrier- and commodity-specific factors (carrier's productivity, diesel prices, commodity-specific packaging and handling etc.)

Expression (1) highlights several key features. First, freight rates are heterogeneous along many dimensions. They depend on the type c of commodity shipped (e.g., dry bulk, liquid bulk, container), the industry i of the product that is shipped (which determines demand conditions that the shippers face), the distance d_{rs} shipped, shippers' production costs m_s^i (and characteristics that correlate with those costs), carriers' productivity as per Y^c , and the spatial distribution M of demand. Controlling for all those dimensions is important when estimating freight rates.

Second, freight rates are endogenous: *they are prices that are set to clear markets and thus reflect supply and demand conditions*. Even if freight rates are largely determined by suppliers' costs in a competitive market, these costs are endogenous to the spatial structure of the economy. For example, imbalances in the geographic distribution of economic activity create imbalances in shipping patterns and influence freight rates via backhaul problems and density economies. Freight rates also depend on factor costs and on the distance shipped, both of which are endogenous to the geographic structure of the economy. The key message is that freight rates and the spatial distribution of economic activity are jointly determined in equilibrium. Dealing with that simultaneity is key to assess the causal effect of transport costs on geographic concentration.

Last, and most importantly, as shown by (1) the importance of transport costs also depends on the value $p_r^{i,\text{prod}}$ of the goods being shipped. When prices are high, transport costs become less important for firms and consumers compared to the value of the goods: it is more profitable to ship gold (expensive or high-quality goods) than gravel (cheap or low-quality goods), all else equal. Changes in prices also affect the importance of transport costs. It may, e.g., be profitable to ship crude oil over long distances when a barrel costs 100\$ but not when it costs 30\$, thus affecting regional production

patterns.⁵ This dependence of ad valorem transport costs on producer prices has two important consequences. On the negative side, there is a simultaneity problem. Geographic concentration in region r may increase productivity because of agglomeration economies or the spatial sorting of firms along productivity (see [Combes and Gobillon, 2015](#)). In turn, a higher manufacturing productivity maps into lower prices $p_r^{i,\text{prod}}$ for manufactured goods and thus affects the importance of freight rates—and possibly the spatial organization of the economy. On the positive side, fluctuations in output prices provide a source of variation that allows to understand the importance of transport costs for geographic concentration independently of changes in unit transport costs. Denote the latter by t_{rs} . As shown in Appendix A,

$$\tau_{rs} = 1 + \frac{\sigma_i - 1}{\sigma_i} \left(t_{rs} / p_r^{i,\text{prod}} \right). \quad (2)$$

Hence, *changes in ad valorem transport costs depend on changes in the unit transport costs t_{rs} relative to the producer price $p_r^{i,\text{prod}}$* . Fluctuations in prices have a direct effect on ad valorem transport costs τ_{rs} . This point is important since prices can fluctuate substantially over medium time horizons (e.g., natural resources such as crude oil), even if per unit transport costs t_{rs} do not change much. In any case, prices introduce an industry-specific component into transport costs, and that component is important to analyze the relationship between transport costs and geographic concentration.

3 Data, measurement, and descriptive evidence

Our analysis requires two key pieces of industry-level information: (i) measures of (absolute and relative) geographic concentration; and (ii) measures of ad valorem transport costs. We now discuss our data and measures, and take a first look at the evidence.

3.1 Data sources

Our primary data source for geographic concentration is the Annual Survey of Manufacturers (ASM) Longitudinal Microdata file from 1992 to 2009. This confidential dataset contains between 32,000 and 53,000 manufacturing plants per year, covering 257 NAICS 6-digit industries. For every plant we have information on: its primary 6-digit NAICS code (the codes are consistent over the whole period); its year of establishment; its total employment; whether or not it is an exporter in selected years; its sales; the number of non-production and production workers, and the hours they worked; its ownership status and whether it belongs to a multiunit firm; its intramural research and development expenditures; and its 6-digit postal code. The latter, when combined with the Postal Code Conversion files (PCCF), allow us to effectively geo-locate the plants using latitude and longitude coordinates of postal code centroids, which are spatially very fine-grained in Canada (see Figure 6 in the appendix). We use the latitude and longitude information to estimate our measures of geographic concentration. We use the remaining information to construct various controls related to industry structure by aggregating to the industry level. Additional industry-level information from the KLEMS database is used to construct proxies for natural advantage and industries' labor force composition. A summary and descriptive statistics is provided in Table 5 in Appendix B, as well as more detailed information on the PCCFs and the sampling frame of the ASM.

We turn next to the data required to estimate ad valorem transport costs $(\tau_{rs} - 1)$ at the industry level. We need information on the revenue to carriers (t_{rs}) and on the producer unit price of goods $(p_r^{i,\text{prod}})$. Statistics Canada's Trucking Commodity Origin-Destination Survey (TCOD) includes both domestic and cross border shipments and covers the period from 1994 to 2009. It provides most of the information required for the

measurement of transport costs. Although it reports revenues to carriers on a shipment basis, it does not report the value of goods shipped. The latter is estimated by multiplying the tonnage of the commodity shipped—reported by the TCOD—by the commodity’s value per tonne estimated using an ‘experimental export trade file’ produced in 2008 (see [Brown, 2015](#), for details). We leave it to Section 3.3 to describe how these data are used to construct a panel of industry-level ad valorem transport rates.

3.2 Geographic concentration

We measure the geographic concentration of industries using the [Duranton and Overman \(2005](#); henceforth, DO) K -densities. These measures are independent of any arbitrary spatial division of the economy and comparable across industries and time.⁶

We estimate the K -density (probability density function, PDF) of the distribution of bilateral distances between n plants in an industry as follows:

$$\hat{K}(d) = \frac{1}{hn(n-1)/2} \sum_{i=1}^{n-1} \sum_{j=i+1}^n f\left(\frac{d - d_{ij}}{h}\right), \quad (3)$$

where d is the distance at which the K -density is evaluated; h is Silverman’s optimal bandwidth; and f is a Gaussian kernel. The distance d_{ij} between plants i and j is computed using the great circle formula. Note that (3) does not weight plants by any measure of size. Rather than using a plant-count based measure, we can also compute the geographic concentration of employment or sales in an industry. This can be done by adding appropriate employment or sales weights to (3) as follows:

$$\hat{K}_W(d) = \frac{1}{h \sum_{i=1}^{n-1} \sum_{j=i+1}^n (e_i + e_j)} \sum_{i=1}^{n-1} \sum_{j=i+1}^n (e_i + e_j) f\left(\frac{d - d_{ij}}{h}\right), \quad (4)$$

where e_i and e_j are the value of employment or sales of plants i and j , respectively.⁷ The weighted K -density thus describes the distribution of bilateral distances between either

employees or dollars of sales in a given industry.

We measure *absolute geographic concentration* of industries by looking at their location patterns up to some distance d . This information can be obtained from the cumulative distribution (CDF) of the K -densities:

$$\text{CDF}(d) = \int_0^d \hat{K}(i) di \quad \text{and} \quad \text{CDF}_W(d) = \int_0^d \hat{K}_W(i) di. \quad (5)$$

The cumulative (5) at distance d provides a measure of the share of plants (or of employees or sales, in the weighted case) in an industry that are located at most at distance d from each other. For example, a value of 0.18 at 50 kilometers for for ‘Motor Vehicle Metal Stamping’ in 1999 (see panel (B) of Table 8 in Appendix D) means that 18% of plant pairs in that industry were located less than 50 kilometers from each other. Larger values of the cumulative K -densities are associated with more geographic concentration of the industry. However, this measure does not tell us anything about specialization, i.e., about the relative concentration of the industry compared to manufacturing in general.

We measure *relative geographic concentration* by comparing the K -density PDFs (3) or (4) with an appropriately defined benchmark distribution. We follow the procedure developed by [Duranton and Overman \(2005\)](#) to construct such a benchmark. The idea is to use Monte Carlo simulations, where plants in an industry are randomly reshuffled 1,000 times across all locations containing manufacturing plants to compute counterfactual distributions of K -densities. The distribution of these K -densities is then used to derive upper and lower bounds, $\overline{K}(d)$ and $\underline{K}(d)$, of confidence bands at every distance d . The latter can be used for statistical testing of the significance of geographic concentration patterns: if the K -density lies above the upper bound of the confidence band ($\hat{K}(d) > \overline{K}(d)$), the industry is significantly more concentrated than manufacturing in general at distance d ; and if the K -density lies below the lower bound of the confidence

band ($\hat{K}(d) < \underline{K}(d)$), the industry is significantly more dispersed at distance d . We implement this approach to assess the significance of geographic concentration using a 90% confidence band. We further use it to construct two measures of relative concentration. First, we create a dummy variable that takes value one if the industry is significantly localized, and zero otherwise. Second, we create a variable that is the difference between the K -density and the upper bound of the confidence band or zero otherwise:

$$\hat{\Gamma}(d) = \int_0^d \max \left\{ \hat{K}(i) - \overline{K}(i), 0 \right\} di. \quad (6)$$

This variable measures the excess agglomeration of an industry compared to manufacturing in general and provides a natural measure for relative geographic concentration.

We provide detailed results of our K -density estimations in Table 6 of Appendix D. The evidence points to a significant decrease in the absolute geographic concentration of manufacturing industries in Canada over the last 20 years, no matter whether that concentration is measured in terms of plant counts, employment, or sales. Furthermore, albeit less pronounced, there were also changes in the relative geographic concentration of industries, i.e., patterns of geographic specialization have evolved, with slightly less specialization overall. Table 1 provides summary statistics of our K -density estimates. As expected, there is substantial variation between industries in their degree of geographic concentration. There is also substantial time-series variation, especially at short distances. Although the bulk of the variation in the K -densities is cross sectional, the rapid churning of plants also provides substantial temporal variation, especially at close proximity. Similar patterns hold for our relative concentration measures. This variation helps with the identification of the effect of changes in the importance of transport costs on changes in geographic concentration using the panel dimension of our data. Understanding those changes is the objective of the remainder of this paper.

Table 1: Summary statistics for different geographic concentration measures and transport costs.

Variable	Industry detail	Mean	Standard deviation		
			Overall	Between	Within
Duranton-Overman K -density cumulative (CDF) at 10 km	6-digit	0.015	0.031	0.023	0.021
Duranton-Overman K -density cumulative (CDF) at 50 km	6-digit	0.065	0.060	0.050	0.034
Duranton-Overman K -density cumulative (CDF) at 100 km	6-digit	0.120	0.085	0.074	0.042
‘Significant concentration’ dummy	6-digit	0.352	0.478	0.384	0.284
Excess concentration Γ_i	6-digit	0.029	0.080	0.069	0.041
Ad valorem trucking rates as share of the value of the goods shipped	6-digit/ L -level	0.034	0.035	0.030	0.005

Notes: Based on the sample that we use in our regression analysis, which includes $4,369 = 257 \times 17$ industry-year observations.

The standard deviation is decomposed into between and within components, which measure the cross-sectional and the time-series variation, respectively. The ‘Ad valorem trucking rates as share of the value of the goods shipped’ is for an average load. They are estimated using 6-digit level detail in the cross-section, and the industry price indices are at the L -level. The ‘Significant concentration’ dummy is a variable that takes value one if industry i is significantly geographically concentrated in year t , and zero otherwise. The ‘Excess concentration Γ_i ’ variable is defined in (6). We restrict these two variables to the 1,802 industry-year pairs for industries that are at least once significantly concentrated over the 1992–2008 period. Additional information regarding our data sources and the construction of our key variables are provided in Section 3.1 and in Appendix B.

3.3 Transport costs

Our second key ingredient is an industry-specific measure of ad valorem transport costs. Contrary to most existing studies, we use a *direct measure* constructed from detailed micro-data files on shipments within Canada. To estimate ad valorem rates, as noted above, we require information on freight rates and the unit price of the goods shipped.

We first model freight rates using shipment (waybill) data from the TCOD. We assume carrier m sets freight rates for shipment l such that both fixed and variable (linehaul) costs are just covered: $t_{m,l} = \alpha + \beta d_l$, where α is the fixed price component, β is the rate per kilometer, and d_l is the distance shipped. In the context of our conceptual framework in Appendix A, $t_{m,l}/w_l$ is an estimate of the per unit freight rate (t_{rs}) of a shipment, where units are measured by weight in tonnes (w). Firms may also price on a per tonne-km basis and this is taken into account by assuming firms set prices based on an unknown average tonnage w^* shipped, which implies that the rate is $t_{m,l} = \alpha + (\beta/w^*)d_l w^*$. This provides a flexible functional form that permits firms to price on a per tonne-km or per km basis. If firms price using the latter, for loads less (greater) than w^* the price per tonne-km will be scale upward (downward). This is captured by the following function:

$$t_{m,l} = \alpha + \left[\frac{\beta}{w^*} + \phi(w^* - w_l) \right] d_l w_l = \alpha + \left(\frac{\beta}{w^*} + \phi w^* \right) d_l w_l - \phi d_l w_l^2, \quad (7)$$

where w_l is the observed tonnage shipped and $\phi(w^* - w_l)$ is the scaling factor. Factoring out the known tonnage w_l results in an estimable function that allows firms to price using either rule or some hybrid of the two. Equation (7) is estimated from the TCOD across three types of carriers—truck-load, less-than-truck load, and specialized—for which variable and fixed costs are expected to vary due to differences in technology

or business model employed.⁸ We also allow fixed and linehaul costs to vary across time in a flexible way. Additional controls include distance and its square to take into account backhaul effects on prices, a diesel price index (that is highly correlated with rates through time), and commodity-carrier fixed effects to control for time costs and the quality/nature of the transportation service.⁹ We use (7) to predict freight rates $\hat{t}_{c,\xi}^y$ by carrier type ξ and commodity c in year y using the average tonnage of a shipment.

The predicted rates $\hat{t}_{c,\xi}^y$ are converted to ad valorem rates—expressed as a proportion of the value of the good shipped—by using the value of shipments by commodity. Since that value is not reported by the TCOB, it is estimated from the ‘experimental export trade file’ produced in 2008 (see [Brown, 2015](#), for details). Let \hat{v}_c^{2008} denote the value of the average tonnage of a shipment of commodity c in 2008, which will serve as our estimate of ‘producer unit prices’, $p_r^{i,\text{prod}}$. Weighting the $\hat{t}_{c,\xi}^y$ across carrier types ξ , using as weights the value of the goods shipped by each carrier type, the ad valorem estimate at the commodity level in 2008 is $\hat{\tau}_c^{2008} - 1 = \hat{t}_c^{2008} / \hat{v}_c^{2008}$. Finally, using an industry-commodity concordance, the ad valorem transport rates for commodities are aggregated to an industry basis ($\hat{\tau}_i^{2008} - 1$) using the value of commodities shipped as weights.

As stated above, the value per ton estimates allow us to construct our ad valorem trucking costs for 2008 only. To generate an industry time series, yearly trucking industry price indices ($\mathbb{P}_{\text{trans}}^t$) and manufacturing industry price indices (\mathbb{P}_i^t) from Statistics Canada’s KLEMS database are used to project the ad valorem rates backwards and forwards in time, thereby creating an industry-specific transportation rate time series:

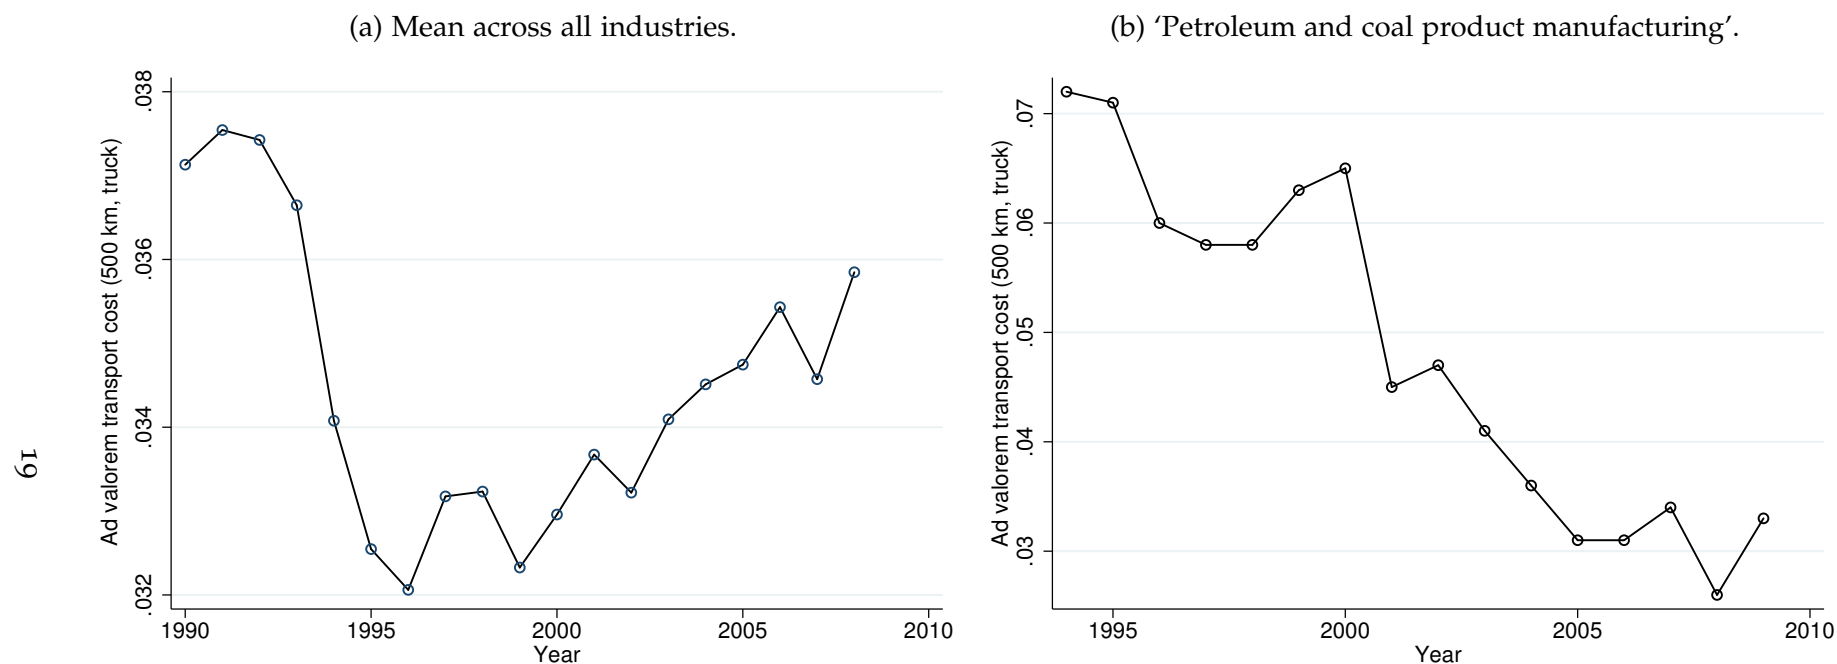
$$\hat{\tau}_i^t - 1 = \frac{\mathbb{P}_{\text{trans}}^t}{\mathbb{P}_i^t} \left(\hat{\tau}_i^{2008} - 1 \right). \quad (8)$$

Cross-sectional variation in industry ad valorem transport rates will depend on: (i) variation in the value per tonne of the good, which can vary by orders of magnitude across

commodities; (ii) average tonnage shipped, which affects both the value of the shipment and the transport rate; and (iii) the nature and level of service provided, which affects rates and will vary, for instance, by carrier type. We provide a full set of 4-digit industry freight rates in Table 12 in Appendix D. The time-series variation comes from the changes in the industry- and transport price indices. As explained above, these capture relevant changes in the importance of transport costs to suppliers and customers. We return to this important point in more detail later.

Panel (a) of Figure 1 depicts the year-on-year changes in the (unweighted 6-digit level) industry average ad valorem trucking costs. As can be seen from that figure, transport costs are first decreasing—due essentially to decreasing labor costs at constant fuel prices—and then increasing—due essentially to increasing fuel prices at constant labor costs. They range from about 3.8% of the value of the shipments in the early nineties, to about 3.2% in the mid-nineties, with an average value of 3.4% (see Table 1 for summary statistics on our estimated ad valorem trucking rates). These figures are fairly close to the average ad valorem rates of 4.6% reported by [Glaeser and Kohlhase \(2004, p.206\)](#) using more aggregated 2002 U.S. data. As in their case, there is significant cross-industry variation in our data. Between 1990 and 2008, the average rate of the ten most expensive-to-ship industries is between 12.2%–14.3%, while that of the ten cheapest-to-ship industries is between 0.34%–0.40% (see Tables 11 and 12 in Appendix D). As expected, the highest ad valorem transport costs are for industries with low value-to-weight ratios (e.g., cement and gypsum product manufacturing and breweries), with an average rate across the top 10 industries in 2008 of 14%. The lowest ad valorem transport costs are for industries with high value-to-weight ratios (e.g., computer and peripheral equipment manufacturing, and medical equipment and supplies manufacturing), with

Figure 1: Changes in ad valorem trucking costs in Canadian manufacturing.



Notes: Panel (a) depicts the unweighted average across 6-digit NAICS industries. These correspond to the estimates that we use in the remainder of this paper. Panel (b) depicts an example estimated for an industry at the 4-digit level. For confidentiality reasons, we cannot disclose detailed 6-digit estimates (see Appendix D.4 for additional details). We report summary results at the 6-digit level in Table 11 and an additional full set of 4-digit estimates in Table 12 in Appendix D.

an average rate across the bottom 10 industries in 2008 of barely 0.39%.

Table 1 shows that, although the bulk of the variation in trucking rates is cross-sectional, we also have time-series variation due to changes in prices. This is further illustrated by panel (b) of Figure 1 for ‘Petroleum and coal product mfg.’ (NAICS 3241). As can be seen from that figure, the ad valorem trucking costs in that industry fell from more than 7% in 1994 to less than 3% during the ramp up to the 2008 spike in oil prices. These changes show that the effects of trucking costs are likely to crucially depend on industry prices and on how those prices change over time. Large changes over time provide variation that will be useful to identify how transport costs affect geographic patterns in the panel dimension; while large differences across industries will be useful to understand whether high or low transport cost industries are more or less agglomerated in the cross-sectional dimension. We now estimate these effects.

4 Empirical approach

We provide both cross-sectional and panel estimates of the effect of ad valorem transport costs on the geographic concentration of industries. There are two reasons for providing both types of estimates. First, the cross-sectional and the panel estimates answer two different questions. While the cross section tells us whether high transport cost industries are more or less geographically concentrated in a given spatial equilibrium, the panel evidence provides ‘comparative statics’ of whether falling transport costs tend to agglomerate or disperse industries between different spatial equilibria. While related and equally interesting, these are two different questions that require different specifications to be answered. Second, as explained in Section 3.3, the cross-sectional variation in our

ad valorem transport costs comes from estimates of unit costs from the TCOB and from unit prices of goods using the special exporter files in 2008; whereas the time-series variation stems from changes in industry price indices. While this variation is sizable, and should therefore allow us to assess the effects of changes in transport costs on changes in geographic concentration, some may worry that we are picking up special effects due to the use of price indices. Because we find robust results between the 2008 and the pooled cross sections and the panel—both for absolute and relative geographic concentration—our results are unlikely to depend on the way we construct our measures.

4.1 Specification

In what follows, we estimate different versions of the following model, both in its (pooled) cross-sectional and panel versions:

$$\text{CDF}_{i,t}(d) = (\tau_{i,t} - 1)\beta_\tau + \mathbf{X}_{i,t}\beta_X + \alpha_t + \mu_i + \varepsilon_{i,t}, \quad (9)$$

where $\text{CDF}_{i,t}(d)$ is the K -density CDF for industry i in year t at distance d (either unweighted or weighted); $(\tau_{i,t} - 1)$ is our measure of ad valorem transport costs (8) of industry i in year t ; $\mathbf{X}_{i,t}$ is a vector of time-varying industry controls; α_t and μ_i are year and (in the panel) industry fixed effects, respectively; and $\varepsilon_{i,t}$ is an i.i.d. error term. As Figure 8 in Appendix D shows, the distributions of $\text{CDF}_{i,t}(d)$ and $(\tau_{i,t} - 1)$ are both right-skewed and look relatively normal once log-transformed. Hence, we apply a log transformation to all variables in our estimations, except for trade shares which we keep in levels, to obtain a distribution of error terms that is closer to a normal distribution.

In the panel version of (9), we include industry and year fixed effects. The former soak up unobserved time-invariant industry characteristics that can map into siz-

able cross-sectional differences in geographic concentration patterns (see Table 1, which shows that much of the variance in the K -densities is cross sectional). The latter control for general trends that affect the geographic concentration of industries, like improvements in informations and communications technologies that could have made economic activity more footloose over our study period. Turning to our controls $\mathbf{X}_{i,t}$, we first construct two proxies, one for the proximity to customers and suppliers, and the other for exposure to international trade. The former is important since changes in the transport costs of one industry may induce changes in the location patterns of vertically linked industries (see, e.g., [Fujita, Krugman, and Venables, 1999](#), for a model). The latter is important since the theoretical literature has shown that international trade costs interact with domestic transport costs to affect geographic concentration patterns (see [Brühlhart, 2011](#), for a review). We proxy access to customers and suppliers using our microgeographic data to construct input-output share weighted distance measures. These measures capture how close an industry is to other vertically linked industries from which it buys or to which it sells. We measure industry-level trade exposure (exports and imports), broken down by broad country groups—NAFTA, OECD excluding NAFTA, and low-cost countries. Appendices B.5 and B.6 provide details, descriptive statistics, and discuss a number of additional concerns related to these measures.

The urban economics literature has substantiated the existence of other agglomeration forces that are independent of the costs of transporting goods but depend on the costs of transporting people and ideas. For example, firms benefit from localized knowledge spillovers and local pools of specialized labor if they locate close to one another. When transport costs are low enough, firms no longer need to be close to their customers and suppliers, which can lead to more geographic concentration of specific industries to

exploit those agglomeration forces.¹⁰ Industries also display different agglomeration patterns, depending on characteristics linked to industry structure (Rosenthal and Strange, 2003). We provide more details on our controls related to agglomeration forces and industry structure in Appendix B.4.

Our coefficient of interest, β_τ , captures whether high ad valorem transport cost industries are more or less geographically concentrated in a given spatial equilibrium (in the cross section); or whether industries with falling transport costs experience more or less geographic concentration between two equilibria (in the panel). Observe that, as explained in Section 2.1, the sign of the coefficient β_τ also provides information about whether the Krugman or the Helpman model offers a better description of the geographic concentration process: if $\beta_\tau < 0$, industries with lower transport costs or industries that see their transport costs fall concentrate more geographically, as in Krugman (1991); whereas if $\beta_\tau > 0$, industries with lower transport costs or industries that see their transport costs fall disperse more geographically, as in Helpman (1998).

4.2 Identification

For β_τ to capture the causal effect of transport costs on geographic concentration, we need to address a number of identification problems. A first problem is due to agglomeration economies: the geographic concentration of an industry may decrease producer prices, which affects ad valorem transport costs. A second related problem arises because geographic concentration leads to imbalances in shipping patterns, and the latter increase transport costs due to ‘backhaul’ of empty trucks (for evidence see, e.g., Jonkeren, Demirel, van Ommeren, and Rietveld, 2009; and Tanaka and Tsubota, 2017). To summarize, our ad valorem transport costs (2) are potentially endogenous to the geographic

concentration of an industry via both t_{rs} and $p_r^{i,\text{prod}}$. Thus, the OLS estimate of β_τ is likely to be upward biased.¹¹ We discuss these problems more formally in Appendix A.

In the cross section, the choice of instruments is unfortunately limited. We rely on ‘binning instruments’, i.e., we use the rank of the cross-sectional ad valorem rates as an instrument, with either quintile bins or tertile bins. The underlying idea is that the potential endogeneity bias in the ad valorem rates is less likely to change the ranks of industries in the distribution than the magnitude of transport costs, and even less likely to push transport costs across bins of the ranking.

In the panel regressions, we have more options to deal with the above-mentioned problems. In what follows, we use the following strategies. First, we clear out the productivity effects—one presumed source of endogeneity—on prices by regressing our transport cost series (8) on industry multi-factor productivity indices from the KLEMS database. We then use the residual from that regression as a proxy for the transportation cost series. By definition, that residual is orthogonal to any productivity-driven price changes that could stem from the changing geographic concentration of industries and affect $\mathbb{P}_{i,t}$ in (8). We refer to these as the ‘ad valorem transport cost residuals’ (AVTCR). This strategy does not deal directly with the transportation price index $\mathbb{P}_{\text{trans},t}$.

Second, we use U.S. industry price indices as external instruments to construct our transport costs series. The underlying idea is the following. Assume that the geographic concentration of an industry increases over time because of unobserved factors. This increasing geographic concentration then raises ad valorem transport costs via decreases in producer prices and increases in trucking rates. Provided that the U.S. changes are not driven by the same unobserved factors than in Canada, but that the U.S. series $\mathbb{P}_{\text{trans},t}^{\text{US}}/\mathbb{P}_{i,t}^{\text{US}}$ are correlated with the Canadian series $\mathbb{P}_{\text{trans},t}/\mathbb{P}_{i,t}$, this yields valid instru-

ments for the Canadian transport cost series. One potential problem arises if the geographic concentration of an industry in Canada directly affects the productivity—and thus the price indices—in the U.S. While we cannot completely rule out this possibility, it does not strike us as very plausible: Canada is ten times smaller than the U.S., so that changes in the geographic distribution of Canadian industries are unlikely to drive changes in the U.S. industry price indices. As an additional check, we run IV regressions that exclude industries (e.g., automobile) with high NAFTA trade shares.

Last, we also use internal instruments in the estimation of (9) using the method of [Lewbel \(2012\)](#) that exploits heteroscedasticity and variance-covariance restrictions to obtain identification with 2SLS when some variables are endogenous and when external instruments are either weak or not available. See Appendix C for details. This approach deals with endogeneity concerns of our trade exposure and input-output distance measures (see Appendices B.5 and B.6 for a discussion). We continue to use the external U.S. instrument for transport costs. We view this approach as an additional robustness check on top of the two strategies that rely on ‘filtering’ and external instruments.

5 Results

5.1 Spatial equilibrium: Cross-sectional evidence

Table 2 shows our cross-sectional estimates of specification (9), which provide evidence on whether high or low transport cost industries are more or less geographically concentrated in a given spatial equilibrium. We pool all years of our data to increase the number of observations and include year fixed effects. Results for the individual cross section in 2008—the year that provides our purely cross-sectional variation—are pro-

vided in Table 13 of Appendix E. The results are qualitatively similar and robust across the pooled cross section and the individual cross section, so we report only the former.

Model (X1) reports our basic cross-sectional estimates without any controls. As shown, the coefficient on ad valorem trucking costs is negative and highly significant. In words, low transport cost industries are on average geographically more concentrated than high transport cost industries. We then progressively add in **Models (X2)** to **(X5)** our controls for international trade exposure, input-output links, and other industry characteristics. Starting with **(X2)**, import and export exposure do not significantly correlate with geographic equilibrium patterns in the cross section. The estimated coefficients (not reported) are almost all insignificant. In **(X3)**, we add our input-output distances.¹² The estimated coefficients on both variables are negative and highly significant. Industries that locate close to their suppliers and customers—i.e., small values of those distances—tend to be more geographically concentrated in the cross section. Note that the coefficient on the ad valorem transport costs decreases substantially when including our measures for access to suppliers and customers. This shows that controlling for equilibrium cross-industry structure is important and that the dispersion of high transport cost industries in equilibrium is partly explained by the dispersion of supplier and client industries. However, the coefficient on the transport costs remains negative and significant. **Model (X4)** shows that the joint inclusion of both trade and input-output controls does not significantly change our estimates. Last, **Model (X5)** adds various industry controls that may influence the geographic concentration of industries. Our results are robust to the inclusion of these controls, although the coefficient on transport cost drops in absolute terms a bit and is less precisely estimated. In our preferred specification, **(X5)**, if an industry has 100% higher ad valorem transport costs than another industry, it

Table 2: Estimation results for specification (9) in the pooled cross section.

	(X1)	(X2)	(X3)	(X4)	(X5)	(X6)	(X7)		(X8)		(X9)
	Base	Trade	IO-links	Trade&IO	Controls	Purged	First stage	IV-2SLS Q5	First stage	IV-2SLS Q3	Lewbel
AVTC	-0.266 ^a	-0.250 ^a	-0.085 ^a	-0.098 ^a	-0.076 ^b						
	(0.042)	(0.043)	(0.027)	(0.030)	(0.034)						
AVTC residual						-0.076 ^b		-0.120 ^a		-0.121 ^a	-0.087 ^a
						(0.034)		(0.039)		(0.041)	(0.034)
AVTC residual, binning instrument							0.504 ^a		0.762 ^a		
							(0.0123)		(0.030)		
Input distance			-0.257 ^a	-0.264 ^a	-0.323 ^a	-0.323 ^a	-0.106	-0.322 ^a	-0.171	-0.322 ^a	-0.311 ^a
			(0.084)	(0.084)	(0.071)	(0.071)	(0.057)	(0.071)	(0.072)	(0.071)	(0.109)
Output distance			-0.416 ^a	-0.422 ^a	-0.391 ^a	-0.391 ^a	0.163	-0.393 ^a	0.217	-0.393 ^a	-0.466 ^a
			(0.084)	(0.085)	(0.069)	(0.069)	(0.055)	(0.069)	(0.064)	(0.070)	(0.090)
Industry controls included	No	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Trade shares included	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations (NAICS \times years)	4,369	4,369	4,369	4,369	4,369	4,369	4,369	4,369	4,369	4,369	4,369
R^2	0.149	0.202	0.652	0.654	0.725	0.725	0.880		0.823		0.327
F test of excluded instruments								1,667.73		661.70	

Notes: ^a, ^b and ^c denote coefficients significant at the 1%, 5%, and 10% levels, respectively. Standard errors in parentheses are clustered at the industry level. The dependent variable is the count-based Duranton-Overman K -density CDF at 50 kilometers distance. We have 17 years and 257 industries. All regressions include year dummies. Our measures of input and output distances, as well as average minimum distance, are computed using $N = 5$ (see Appendices B.5 and B.6 for details). A constant term is included but not reported. (X5)–(X9) include the following industry controls: Total industry employment; Firm Herfindahl index (employment based); Mean plant size; Share of plants affiliated with multiplant firms; Share of plants controlled by foreign firms; Natural resource share of inputs; Energy share of inputs; Share of hours worked by all workers with post-secondary education; In-house R&D share of sales. ‘AVTC residual’ denotes the residual of the regression of ad valorem trucking costs on industry multi-factor productivity. (X6) reports bootstrapped standard errors. (X7) and (X8) use the rank-bin of the cross-sectional rates as an instrument, with quintile bins (Q5) in the former and tertile bins (Q3) in the latter. (X9) follows [Lewbel \(2012\)](#) to instrument input-output distances and trade shares (see Appendix C for details). We still use quintile bins (Q5) as external instrument for transport costs.

is on average $100\% \times (2^{-0.076} - 1) = -5.13\%$ less concentrated at 50 kilometers distance.

As discussed in Sections 2.3 and 4.2, our measure of ad valorem transport costs is potentially endogenous. Table 2 reports a number of estimates that address this problem. Starting with **Model (X6)**, we use the residual transport cost obtained from a regression of that cost on industry multi-factor productivities.¹³ The results are virtually the same as in **(X5)**. In **Models (X7)** and **(X8)**, we report IV-2SLS results where we instrument our ad valorem transport costs by their rankings in the cross section, using either their quintile bins (in **(X7)**) or tercile bins (in **(X8)**). In line with our expectations, the coefficient on transport costs indeed significantly drops when instrumenting using ranking bins, whereas the standard errors hardly change. This suggests that our results are robust to controlling for endogeneity and that doing so is important in the cross section. Furthermore, the effects of transport costs on geographic concentration are sizable. In **Model (X8)**, industries with 100% higher ad valorem transport costs are on average $100\% \times (2^{-0.121} - 1) = -8.04\%$ less concentrated at 50 kilometers distance. Finally, **Model (X9)** addresses remaining endogeneity concerns that could affect the trade variables and the input-output distance measures, while keeping the quintile-bin instrument for transport costs. Our results are again robust to this instrumentation strategy.

5.2 Change in spatial equilibrium: Panel evidence

We now estimate the panel version of equation (9), which provides answers to the comparative statics question of how changes in transport costs change the geographic concentration of industries between two spatial equilibria. Table 3 summarizes our results.

Model (P1) reports our basic panel estimates without any controls. As can be seen, the coefficient on the ad valorem trucking costs is negative and significant. In words,

Table 3: Estimation results for specification (g) in the panel version.

	(P1)	(P2)	(P3)	(P4)	(P5)	(P6)	(P7)		(P8)	(P9)
Variables	Base	Trade	IO-links	Trade&IO	Controls	Base	First stage	IV-2SLS	Lewbel version 1	Lewbel version 2
AVTC	-0.337 ^b (0.155)	-0.263 ^b (0.122)	-0.250 ^b (0.098)	-0.183 ^b (0.078)	-0.208 ^b (0.088)					
AVTC residual						-0.261 ^a (0.078)		-0.393 ^a (0.096)	-0.197 ^b (0.093)	-0.218 ^b (0.092)
AVTC U.S. instrument							0.485 ^a (0.111)			
Asian share of imports		-1.639 ^a (0.413)		-1.309 ^a (0.379)	-1.132 ^a (0.380)	-1.118 ^a (0.383)	-0.056 (0.107)	-1.095 ^a (0.381)	-1.592 ^a (0.533)	-1.618 ^a (0.501)
Input distance			-0.359 ^a (0.064)	-0.343 ^a (0.059)	-0.361 ^a (0.055)	-0.358 ^a (0.055)	0.035 ^c (0.020)	-0.356 ^a (0.055)	-0.143 ^c (0.075)	-0.224 ^a (0.075)
Output distance			-0.262 ^a (0.046)	-0.290 ^a (0.044)	-0.313 ^a (0.042)	-0.318 ^a (0.043)	-0.011 (0.015)	-0.322 ^a (0.042)	-0.386 ^a (0.083)	-0.360 ^a (0.087)
Industry controls included	No	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Trade shares included	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations (NAICS \times years)	4,369	4,369	4,369	4,369	4,369	4,369	4,369	4,369	4,369	4,369
R^2	0.054	0.096	0.442	0.473	0.516	0.518	0.628		0.319	0.330
F test of excluded instruments								19.07		

Notes: ^a, ^b and ^c denote coefficients significant at the 1%, 5%, and 10% levels, respectively. Standard errors in parentheses are clustered at the industry level. The dependent variable is the count-based Duranton-Overman K -density CDF at 50 kilometers distance. We have 17 years and 257 industries. Our measures of input and output distances, as well as average minimum distance, are computed using $N = 5$ (see Appendices B.5 and B.6 for details). A constant term is included in all regressions but not reported. (P5)–(P11) include the following industry controls: Total industry employment; Firm Herfindahl index (employment based); Mean plant size; Share of plants affiliated with multiplant firms; Share of plants controlled by foreign firms; Natural resource share of inputs; Energy share of inputs; Share of hours worked by all workers with post-secondary education; In-house R&D share of sales. ‘AVTC residual’ denotes the residual of the regression of AVTC on industry multi-factor productivity. (P6) reports bootstrapped standard errors. (P7) instruments the ‘AVTC residual’ using transport costs constructed with U.S. price indices. (P8) and (P9) follow Lewbel (2012) to instrument input-output distances and trade shares. In (P8), only a subset of the trade shares is instrumented, while all trade shares are instrumented in (P9).

falling ad valorem trucking costs within industries are associated with their geographic concentration. We then progressively add in **Models (P2) to (P5)** our controls for international trade exposure, input-output links, and other industry characteristics. Starting with **(P2)**, rising import shares are across-the-board associated with falling geographic concentration. The (non-OECD) Asian share of imports—a proxy for low-wage countries—has the largest estimated coefficient in absolute value and is the most statistically significant (the other import shares, though not reported, are negative and significant too). One explanation for the dispersive effect of import competition is that firms become more footloose as they source a larger share of their intermediates from abroad and no longer rely on localized domestic suppliers. Another explanation is that import competition leads to substantial exit of plants in geographic clusters, which reduces geographic concentration (see, e.g., [Holmes and Stevens, 2014](#)).¹⁴ In **Model (P3)**, we add our input-output distances (and our minimum distance control). The estimated coefficients on the input and output distance measures are negative and highly significant: industries tend to follow their suppliers and customers, i.e., industries where potential suppliers or clients disperse tend to also disperse. The coefficient on trucking costs changes only slightly when including our measures for access to suppliers and customers. **Model (P4)** shows that the joint inclusion of both trade and input-output controls does not significantly change our baseline estimates. Although the coefficient on transport cost drops in absolute terms, it remains negative and highly significant. Last, **(P5)** adds our industry controls and our results are again robust.

How large are the effects of changes in transport costs on changes in geographic concentration? First, in our preferred specification **(P5)**, if transport cost in an industry increase by 100%, then geographic concentration falls by $100\% \times (2^{-0.208} - 1) = -13.43\%$

at 50 kilometers distance. Clearly, these are large effects. Second, we can compute the predicted change in the CDFs by holding the ad valorem trucking costs constant at their 1992 values, while still allowing the other variables to change through time. The observed change in the cross-industry average CDF between 1992 and 2008 at 50 kilometers is -23.37%. Holding transport costs fixed at their 1992 levels, the change would have been -28.36%. Thus, had ad valorem trucking costs not fallen between 1992 and 2008, the average geographic concentration of industries would have fallen by about 5 percentage points more (about 20% of the overall change).¹⁵ These are sizable effects.

As in the cross section, Table 3 also reports estimations that address the potential endogeneity of our transport costs. **Model (P6)** uses the residual transport cost obtained from a first-stage regression of that cost on industry multi-factor productivities. The coefficient on transport costs becomes more negative when using the productivity-purged residual (compare **(P5)** to **(P6)**). This is in line with our expectations that agglomeration effects that reduce producer prices are likely to bias the coefficient upwards (towards zero in this case). In what follows, we systematically use the residual measure of ad valorem trucking costs in all of our regressions.

Although the residual transport cost is purged from productivity effects, endogeneity concerns linked to, e.g., backhaul, remain. Hence, we run some instrumental variables regressions to check the validity of our results. **Model (P7)** summarizes our IV-2SLS results where we instrument the ad valorem trucking rate residual by replacing the Canadian price indices with their U.S. counterparts. As can be seen, the instrument is strong, with a first-stage F -test value of 19.07 and a first-stage R^2 of 0.63. The instrumented coefficient is substantially more negative than the coefficient for the residual ad valorem trucking rate. OLS estimates are therefore likely to underestimate the impact of changes

in ad valorem transport costs on the geographic concentration of industries.¹⁶

Last, **Models (P8)** and **(P9)** address remaining endogeneity concerns that may affect the trade variables and the input-output distance measures. We again use the [Lewbel \(2012\)](#) estimator with internal instruments for the input-output distances and the trade shares.¹⁷ The excluded external instrument is the U.S. price-based ad valorem trucking cost residual, as before. As Table 3 shows, the instrumented coefficient on the Asian share of imports increases (as do those on most of the other unreported trade shares), while both the magnitude of transport costs and of the input and output distances decreases slightly. However, these variables remain significant and their magnitude is in the same ballpark than in the case of OLS, thus showing that our results are robust.

5.3 Extensions and robustness

We now report two extensions: (i) the relative geographic concentration of industries; and (ii) the spatial extent of the effects that we estimate. We relegate many additional robustness checks to Appendix E. In particular, we show there how coefficients vary across industries that trade internationally using different transport modes (truck vs ship) and how location patterns depend on import entry points.

Starting with specialization patterns, we first look at the geographic concentration of industries controlling for the overall geographic concentration of manufacturing. There are two reasons for doing so. First, we want to understand whether and how individual industries can concentrate, irrespective of what happens to the overall distribution of economic activity. Doing so provides answers to the question of whether falling transport costs favor regional specialization. Second, one may be worried that our previous results just pick up the downwards trend in Canadian manufacturing and are, therefore,

just about the ‘deindustrialization of southern Ontario’. We will show that our results are not driven by this downwards trend: even controlling for the fact that manufacturing has been dispersing in general—and shrinking in importance—we find that high transport cost industries are less geographically concentrated than manufacturing generally, and that decreasing transport costs lead to regional specialization. Our results are thus not specific to an economy in which manufacturing has been declining in relative terms, since absolute and relative concentration patterns are similar.¹⁸

We use two different measures of relative geographic concentration. First, we create an indicator variable $y_{i,t}$ that takes value one if industry i is significantly concentrated in year t , and zero otherwise (see Section 3.2 for details). We include in our regressions only industries that were at least once significantly concentrated over our study period, and drop the remaining ones. We report linear probability results for both the cross section and the panel.¹⁹ Second, we use a measure of the ‘excess concentration’ of an industry (see expression (6) in Section 3.2). Again, we restrict our sample to industries that were at least once significantly concentrated over our study period. Descriptive statistics for both variables are reported in Table 1.

Table 4 shows our results, both for the cross section ((**X10**) and (**X11**)) and the panel ((**P10**) and (**P11**)). As shown, all coefficients on transport costs are negative and significant, with the exception of (**P11**) where only the point estimate is negative. The coefficient on the dummy variable for significant geographic concentration is large and precisely estimated in both the panel and the cross section, which shows that high transport cost industries tend to be more dispersed than manufacturing in general, and that decreasing transport costs tend to concentrate industries geographically more strongly than manufacturing in general. The results using the ‘strength of agglomeration’ vari-

Table 4: Relative geographic concentration.

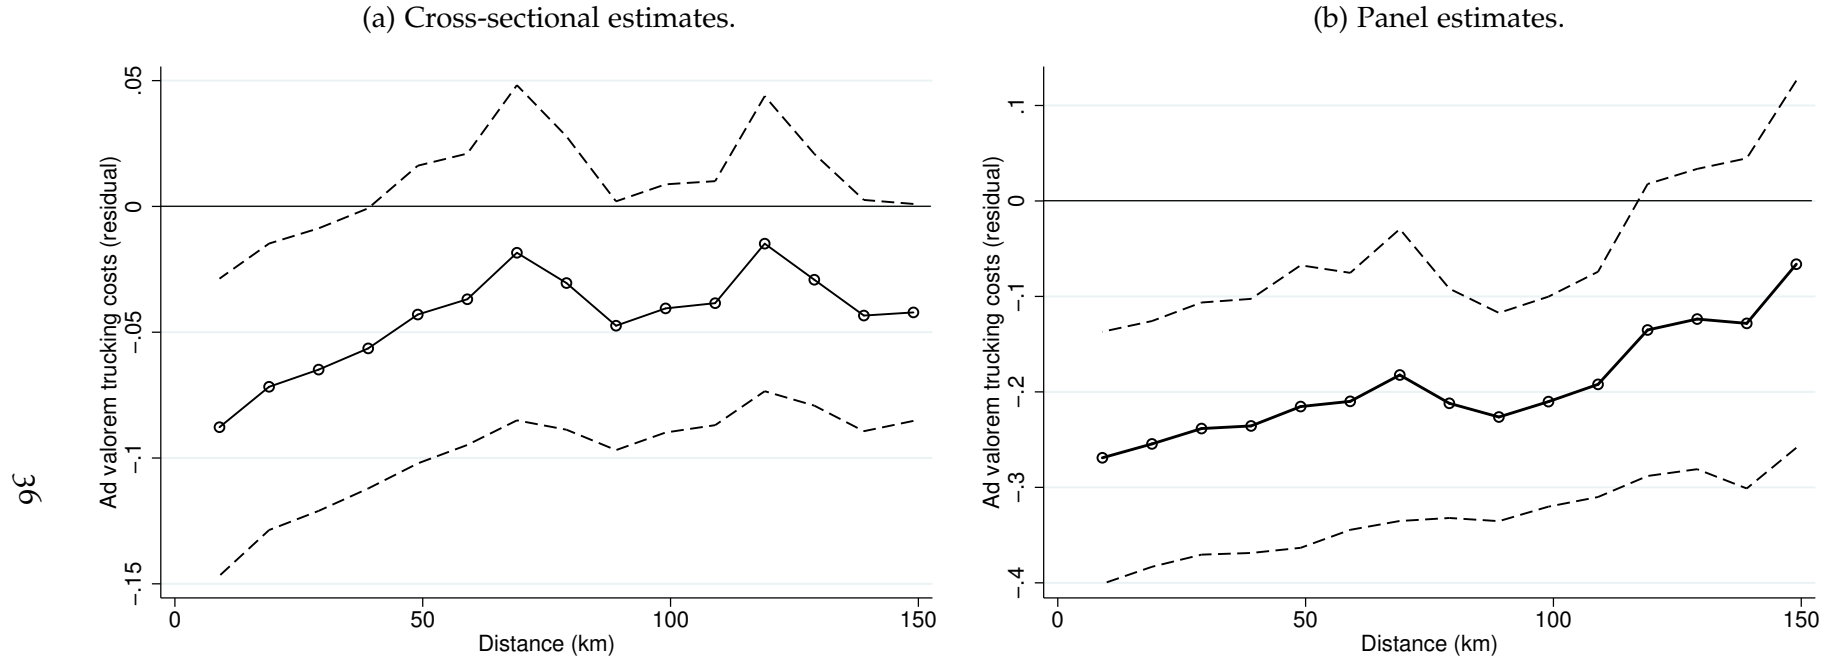
	(X10)	(P10)	(X11)	(P11)
Variables	Significant concentration measure	Excess concentration measure Γ_i		
AVTC residual	-0.124 ^a (0.040)	-0.611 ^a (0.150)	-0.025 ^b (0.010)	-0.038 (0.025)
Asian share of imports	-0.500 ^c (0.333)	-0.042 (0.151)	-0.167 ^c (0.253)	-0.246 (0.099)
Input distance	-0.254 ^a (0.084)	-0.412 ^a (0.021)	-0.094 ^a (0.092)	-0.043 ^b (0.027)
Output distance	-0.065 (0.089)	-0.210 ^b (0.014)	-0.003 (0.085)	-0.029 ^b (0.016)
Industry dummies	No	Yes	No	Yes
Observations (NAICS \times years)	1,802	1,802	1,802	1,802
R^2	0.351	0.209	0.437	0.238

Notes: ^a, ^b and ^c denote coefficients significant at the 1%, 5% and 10% levels, respectively. Standard errors are clustered at the industry level and given in parentheses. The dependent variable is a dummy for ‘significant concentration’ in **Models (X10)** and **(P10)**, and the excess concentration measure Γ_i in **Models (X11)** and **(P11)**. We retain only industries that are significantly geographically concentrated at least for one year during 1992–2008. Our measures of input and output distances, as well as average minimum distance, are computed using $N = 5$ (see Appendices B.5 and B.6 for details). A constant term is included in all regressions but not reported. All models include year dummies, the input-output distances, the minimum distance, and the following industry controls: Total industry employment; Firm Herfindahl index (employment based); Mean plant size; Share of plants affiliated with multiplant firms; Share of plants controlled by foreign firms; Natural resource share of inputs; Energy share of inputs; Share of hours worked by all workers with post-secondary education; In-house R&D share of sales. All models include all trade shares.

able are similar, albeit weaker and less precisely estimated. Overall, the results linking transport costs to relative geographic concentration are qualitatively similar to the results linking them to absolute geographic concentration, thus suggesting that they are not driven by the relative decline in Canadian manufacturing and that agglomeration and specialization patterns both respond in similar ways to levels of, and changes in, ad valorem transport costs.

Turning next to the strength with which changes in transport costs operate at different spatial scales, we estimate the marginal effects of ad valorem transport costs for 10 kilometer ‘distance bands’. To do so, we change the distance d at which the K -density CDF is evaluated in our regressions and look at the additional increment in coefficients by 10 kilometer windows: $\Delta\gamma_i(d_1, d_2) = \text{CDF}_i(d_2) - \text{CDF}_i(d_1)$, where $d_2 - d_1 = 10$ kilometers. Figure 2 and from Table 16 in Appendix E show that the strongest incremental effects of transport costs on geographic concentration occur at short distances. There are basically no additional effects beyond about 100 kilometers, with a shorter distance in the cross section compared to the panel. The largest coefficients (in absolute value) and statistically most significant results occur in the distance bands between 10 and 30 kilometers in the cross section, and 10 to 100 kilometers in the panel. These results suggest that many of the agglomeration mechanisms linked to transportation operate at the scale of metropolitan areas, either by influencing within-metro patterns or between-metro specialization (see [Duranton et al., 2014](#)). At longer distances—beyond about 100–200 kilometers—other factors drive the clustering of firms and the incremental effect of transport costs on geographic patterns vanishes. In a nutshell, differences in transport costs map into differences in location patterns especially at small geographic scales, whereas the patterns are less affected at larger geographic scales.

Figure 2: Estimated 'AVTC residual' coefficients by distance.



Notes: Panels (a) and (b) report the marginal effects of ad valorem transport costs for 10 kilometer 'distance bands'. The variables included are the same as in **models (X6) and (P6)** but the dependent variable is defined as $\Delta\gamma_i(d_1, d_2) = \text{CDF}_i(d_2) - \text{CDF}_i(d_1)$, where $d_2 - d_1 = 10$ kilometers. The dashed lines are the 90% confidence intervals. We begin with $d_2 = 10$ kilometers and subsequently increase d_1 and d_2 by 10 kilometer steps up to $d_2 = 150$ kilometers. we limit the plot to a range of 150 kilometers since all coefficients are statistically zero after that threshold.

6 Conclusion

We use micro-level commodity flow data and micro-geographic plant-level data to build continuous measures of geographic concentration and industry-level ad valorem trucking rates. Using those measures, we provide cross-sectional and panel evidence on the causal relationship between transport costs and geographic concentration for Canadian manufacturing industries between 1992 and 2008. Our answer to the question whether ‘the world is flat’ is an emphatic, not yet! The key message of our findings is that the degree of geographic concentration of industries differs systematically with transport costs in the cross section, and that changes in the geographic concentration of industries due to changes in transport costs are sizable. Low transport cost industries are significantly more geographically concentrated, and falling transport costs lead to more agglomeration and regional specialization. These findings are in line with [Krugman’s \(1991\)](#) model of economic geography and they survive a battery of robustness checks, including extensive efforts to address inherent endogeneity issues that plague such estimations. We should also add that, to the best of our knowledge, this is the first instance where direct industry-level measures of ad valorem transport costs are used to assess their effects on the geographic concentration of industries.

The lessons for researchers from this work are twofold. The first is that it is difficult to contemplate investigating industry location (or co-location) without taking transport costs explicitly into account. In a nutshell, investing in better measures of transport costs is important and likely to pay substantial dividends. The second is that it is equally difficult to consider the effects of transport costs in isolation. Their general equilibrium effects on input-output links and competition, and more generally their endogenous na-

ture as market prices, have to be grappled with. This involves challenges—both theoretical and empirical—with large investments required for both. While we believe we have made some strides developing the necessary empirics, theoretical work that provides full-blown analytical results on the interaction between transport costs and location is still called for and needed.

The lesson for policy makers is simple: small changes in transport costs—e.g., due to infrastructure projects or simply fluctuations in output prices—still impact the economic geography of industries. Contrary to what seems a received wisdom in many policy circles, the world is not yet a flat featureless plain. Even small changes in ad valorem transport costs—combined with historically low levels of these costs—can strongly affect geography because firms compete globally and their slim profit margins depend on locational advantage. In the end, the debate surrounding the ‘flat world’ is a classical instance of the fallacy consisting in equating ‘low’ with ‘unimportant’.²⁰

Acknowledgements. We thank the editor Gordon Hanson, two anonymous referees, Gabriel Ahlfeldt, John Baldwin, Antoine Bonnet, Pierre-Philippe Combes, Gilles Duranton, Wulong Gu, Julien Martin, Se-il Mun, Yasusada Murata, Frédéric Robert-Nicoud, Mathieu Parenti, Eugenia Shevtsova, Peter Warda, and conference and seminar participants in many places for helpful comments. Tim Pendergast provided excellent technical assistance. Financial support from Statistics Canada’s ‘Trade Cost Analytical Projects Initiative’ (API) is gratefully acknowledged. This project was funded by the Russian Academic Excellence Project ‘5-100’. Behrens and Bougna gratefully acknowledge financial support from the CRC Program of SSHRC, Canada, and from CIRPÉE. This work was carried out while Behrens and Bougna were Deemed Employees of Statistics Canada. The views expressed in this paper are those of the authors and not those of institutions they are

associated with. This paper has been screened to ensure that no confidential data are revealed.

Notes

¹According to the Canadian Trucking Alliance, 90% of final consumed goods are delivered by truck and approximately two-thirds of Canada-U.S. trade is moved by truck.

²Other dispersive forces, such as idiosyncratic taste differences for regional amenities (e.g., [Murata, 2003](#)) or nontraded goods yield the same result. Some recent contributions combine different dispersion forces like urban costs and idiosyncratic taste differences for regional amenities (see, e.g., [Behrens, Murata, Mion, and Suedekum, 2017](#)). Another strand of agglomeration models relies on the [Eaton and Kortum \(2002\)](#) framework. Although conceptually different, that framework relies on the same agglomeration forces as [Krugman \(1991\)](#) and [Helpman \(1998\)](#). When combined with different dispersion forces—e.g., a limited housing supply and idiosyncratic taste differences—or a limited land supply ([Donaldson and Hornbeck, 2016](#)) it makes the same predictions as to how transport costs shape geographic concentration.

³We thank a referee for pointing out this important issue.

⁴Even the scarce studies that use more direct measures of transport costs typically do not have industry-level variation in those costs (e.g., Generalized Transport Costs in [Combes and Lafourcade, 2005](#); or road distances interacted with oil prices in [Storeygard, 2016](#)). [Duranton et al. \(2014\)](#) combine changes in infrastructure at the city level with value-to-weight ratios—a proxy for transport costs—to show how cities' industrial specialization patterns change with infrastructure investments.

⁵Using the elimination of a transport subsidy for grain exports in the Canadian prairies in 1995 as a natural experiment, [Ferguson and Olfert \(2016\)](#) and [Brown, Ferguson, and Viju \(2017\)](#) find strong evidence of this effect. Farms further away from export ports that received a greater subsidy responded to its elimination by shifting more of their production of low value per tonne wheat to high value per tonne canola than farms located closer to ports that received a lower subsidy.

⁶Like any scalar measure of geographic concentration the K -densities are global in nature: for each industry-year-distance triple, they provide a single measure of concentration. We can hence not talk about 'local' concentration, e.g., in the south-eastern part of Ontario or Quebec only. As Figure 6 in the appendix shows, manufacturing is geographically concentrated, and the measure picks up that concentration. How-

ever, it is silent about concentration in specific areas.

⁷The additive weighting scheme we use gives less weight to pairs of large plants and more weight to pairs of smaller plants (see [Behrens and Bougna, 2015](#)). Using a multiplicative scheme implies that the results can be strongly driven by a few very large plants. Canada is known to have a very skewed firm-size distribution, which may be problematic.

⁸Truckload and specialized carriers typically move loads between points, while less-than-truckload carriers ship multiple consignments between distribution centers. Specialized carriers use specialized forms of equipment (e.g., tank trailers), while truckload and less-than-truckload carriers do not.

⁹In order to take into account the time costs of transportation that will be, at least partially, embedded in the transport prices (which would capture quality of service for time-dependent trips) and the nature of the transport service that will vary across commodities and carriers, commodity-carrier fixed effects are included. See [Brown \(2015\)](#) for a more detailed discussion of the model and the data.

¹⁰See [Akamatsu, Mori, Osawa, and Takayama \(2017\)](#) for a discussion of models where falling transport costs simultaneously cause agglomeration at a large spatial scale and dispersion at a small spatial scale.

¹¹Geographically concentrated industries may ship their output over different distances than less agglomerated industries. While this can be taken into account in (7) by predicting ad valorem rates over a fixed distance, across commodities there was a 94% correlation between the fixed (500 kilometers) and variable distance-based estimates. Given that correlation, and that our ‘binning’ instrument should account for any remaining endogeneity in the cross section, we utilize the variable distance-based estimates. We also prefer these estimates because many goods are not shipped over a given fixed distance (e.g., cement is typically not shipped 500 kilometers), which may also bias the estimates. Note that our panel estimates are, by construction, unaffected by this choice.

¹²Whenever we add these variables, we also add a ‘minimum distance’ control. See Appendix B.6.

¹³When using the ‘AVTC residual’ as our explanatory variable, we should bootstrap the standard errors to control for the presence of an estimated regressor. We do this for **Models (X6)** and **(P6)** in Tables 2 and 3, and it makes virtually no difference. We thus report non-bootstrapped standard errors (yet clustered by industry) in all other specifications that use this variable.

¹⁴We cannot disentangle the impact of exit, entry, or relocation on the spatial structure. However, we control for the size of the industry, which at least partly picks up entry and exit dynamics. Note

that relocations are quite rare and should have little impact on our results. We are also silent on the contribution of transportation costs to the creation of economic activity, although this is important to consider (see Redding and Turner, 2015). As a tentative check, we ran panel regressions of industry size (employment) on our measures of ad valorem transport costs. The coefficients were insignificant.

¹⁵We can repeat this exercise for import shares. Holding all import shares fixed at their 1992 level, the change in the CDF would have been -14.63%. In words, had imports remained at their 1992 levels, the geographic concentration would have fallen by about 9 percentage points (i.e., 60%) less.

¹⁶We also ran (P7) excluding the top 10% of industries by trade shares (exports and imports) with NAFTA to minimize remaining endogeneity concerns. The instrumented coefficient on 'AVTC residual' drops to -0.445 (standard error 0.111), whereas the other coefficients remain fairly stable.

¹⁷Since there is an insignificant correlation between the OECD export share and the squared residuals, we did not include it. We substituted instead the NAFTA import share because it is consistently significant in the baseline set of models and it meets the criteria for being internally instrumented (see Appendix C).

¹⁸Canada's manufacturing sector has been shrinking, with total employment falling from 1,814.5 thousand workers in 1992 to 1,694.8 thousand workers in 2016, a 7.06% decrease (see Table 10 in Appendix D). Because of the growth in service industries, manufacturing fell more substantially as a share of overall employment, from 14.25% in 1992 to 9.37% in 2016.

¹⁹Probit results are very similar. However, it is well known that the probit estimator is not consistent when individual fixed effects are included because of the incidental parameters problem.

²⁰See, e.g., the 'kaleidoscopic comparative advantage' debate in international trade (Jagdish Bhagwati, "Why the world is not flat", 2010; available at <http://www.worldaffairsjournal.org/blog/jagdish-bhagwati/why-world-not-flat>). Last accessed on July 11, 2016.

References

- [1] Akamatsu, Takashi, Tomoya Mori, Minoru Osawa, and Yuki Takayama. 2017. "Spatial scale of agglomeration and dispersion: Theoretical foundations and empirical implications." Paper #80689, Munich Personal RePEc Archive.

- [2] Alchian, Armen A., and William R. Allen. 1964. *University Economics*. Belmont, CA: Wadsworth.
- [3] Autor, David H., David Dorn, and Gordon H. Hanson. 2013. "The China syndrome: Local labor market effects of import competition in the United States." *American Economic Review* 103(6): 2121–2168.
- [4] Baum-Snow, Nathaniel. 2007. "Did highways cause suburbanization?" *Quarterly Journal of Economics* 122(2): 775–805.
- [5] Baum-Snow, Nathaniel, Loren Brandt, J. Vernon Henderson, Matthew A. Turner, and Qinghua Zhang. 2017. "Roads, railroads, and decentralization of Chinese cities." *Review of Economics and Statistics* 99(3): 435–448.
- [6] Behrens, Kristian, and Théophile Bougna. 2015. "An anatomy of the geographical concentration of Canadian manufacturing industries." *Regional Science and Urban Economics* 51(C): 47–69.
- [7] Behrens, Kristian, and W. Mark Brown. 2017. "Transport costs, trade, and geographic concentration: Evidence from Canada" In: Bruce A. Blonigen and Wesley W. Wilson (eds.), *Handbook in International Trade and Transportation*. Palgrave. *Forthcoming*.
- [8] Behrens, Kristian, Giordano Mion, Yasusada Murata, and Jens Suedekum. 2017. "Spatial frictions." *Journal of Urban Economics* 97(A): 40–70.
- [9] Behrens, Kristian, and Pierre M. Picard. 2011. "Transportation, freight rates, and economic geography." *Journal of International Economics* 85(2): 280–291.

- [10] Brown, W. Mark. 2015. "How much thicker is the Canada-U.S. border? The cost of crossing the border by truck in the pre- and post-9/11 eras." *Research in Transportation and Business Management* 16****(XX)****: 50–66.
- [11] Brown, W. Mark, Shon M. Ferguson, and Cria Viju. 2017. "Agricultural trade reform, reallocation and technical change: Evidence from the Canadian prairies." NBER Working Paper #23857, National Bureau of Economic Research, Cambridge, MA.
- [12] Brülhart, Marius. 2011. "The spatial effects of trade openness: A survey." *Review of World Economics* 147**(1)**: 59–83.
- [13] Brülhart, Marius, Céline Carrère, and Federico Trionfetti. 2012. "How wages and employment adjust to trade liberalization: Quasi-experimental evidence from Austria." *Journal of International Economics* 86**(1)**: 68–81.
- [14] Chandra, Amitabh, and Eric Thompson. 2000. "Does public infrastructure affect economic activity? Evidence from the rural interstate highway system." *Regional Science and Urban Economics* 30**(4)**: 457–490.
- [15] Combes, Pierre-Philippe, and Laurent Gobillon. 2015. "The empirics of agglomeration economies." In: G. Duranton, J.V. Henderson, and W.C. Strange (eds.), *Handbook of Regional and Urban Economics*, vol. 5. North-Holland: Elsevier, pp.247–348.
- [16] Combes, Pierre-Philippe, and Miren Lafourcade. 2005. "Transport costs: Measures, determinants, and regional policy implications for France." *Journal of Economic Geography* 5**(3)**: 319–349.
- [17] Donaldson, Dave, and Richard Hornbeck. 2016. "Railroads and American economic growth: A 'market access' approach," *Quarterly Journal of Economics* 131**(2)**: 799–858.

- [18] Duranton, Gilles, Peter M. Morrow, and Matthew A. Turner. 2014. "Roads and trade: Evidence from the US." *Review of Economic Studies* 81(2): 681–724.
- [19] Duranton, Gilles, and Henry G. Overman. 2005. "Testing for localisation using micro-geographic data." *Review of Economic Studies* 72(4): 1077–1106.
- [20] Duranton, Gilles, and Matthew A. Turner. 2012. "Urban growth and transportation." *Review of Economic Studies* 79(4): 1407–1440.
- [21] Eaton, Jonathan and Samuel Kortum. 2002. "Technology, geography, and trade." *Econometrica* 70(5): 1741–1779.
- [22] Ellison, Glenn D., Edward L. Glaeser, and William R. Kerr. 2010. "What causes industry agglomeration? Evidence from coagglomeration patterns." *American Economic Review* 100(3): 1195–1213.
- [23] Ferguson, Shon M., and M. Rose Olfert. 2016. "Competitive pressure and technology adoption: Evidence from a policy reform in western Canada." *American Journal of Agricultural Economics* 98(1): 422–446.
- [24] Fujita, Masahisa, Paul R. Krugman, and Anthony J. Venables. 1999. *The Spatial Economy: Cities, Regions and International Trade*. MIT Press, Cambridge, MA.
- [25] Glaeser, Edward L., and Janet E. Kohlhase. 2004. "Cities, regions and the decline of transport costs." *Papers in Regional Science* 83(1): 197–228.
- [26] Garcia-López, Miquel-Àngel, Adelheid Holl, and Elisabet Viladecans-Marsal. 2015. "Suburbanization and the highways: When the Romans, the Bourbons and the first cars still shape Spanish cities." *Journal of Urban Economics* 85**(XX)**: 52–67.

- [27] Helpman, Elhanan. 1998. "The size of regions." In: D. Pines, E. Sadka, I. Zilcha (eds.), *Topics in Public Economics. Theoretical and Empirical Analysis*, Cambridge University Press, pp. 33–54.
- [28] Holmes, Thomas J., and John J. Stevens. 2014. "An alternative theory of the plant size distribution, with geography and intra- and international trade." *Journal of Political Economy* 122(2): 369–421.
- [29] Jonkeren, Olaf, Erhan Demirel, Jos van Ommeren, and Piet Rietveld. 2009. "Endogenous transport prices and trade imbalances." *Journal of Economic Geography* 11(3): 509–527.
- [30] Krugman, Paul R. 1991. "Increasing returns and economic geography." *Journal of Political Economy* 99(3): 483–499.
- [31] Krugman, Paul R., and Anthony J. Venables. 1995. "Globalization and the inequality of nations." *Quarterly Journal of Economics* 110(4): 857–880.
- [32] Lewbel, Arthur. 2012. "Using heteroscedasticity to identify and estimate mismeasured and endogenous regressor models." *Journal of Business and Economic Statistics* 30(1): 67–80.
- [33] Michaels, Guy. 2008. "The effect of trade on the demand for skill: Evidence from the Interstate Highway System." *Review of Economics and Statistics* 90(4): 683–701.
- [34] Murata, Yasusada. 2003. "Product diversity, taste heterogeneity, and geographic distribution of economic activities: market vs. non-market interactions." *Journal of Urban Economics* 53(1): 126–144.

- [35] Redding, Stephen J., and Daniel M. Sturm. 2008. "The costs of remoteness: Evidence from German division and reunification." *American Economic Review* 98(5): 1766–1797.
- [36] Redding, Stephen J., and Matthew A. Turner. 2015. "Transportation costs and the spatial organization of economic activity." In: G. Duranton, J.V. Henderson, and W.C. Strange (eds.), *Handbook of Regional and Urban Economics*, vol. 5. North-Holland: Elsevier, pp.1339–1398.
- [37] Rosenthal, Stuart S., and William C. Strange. 2004. "Evidence on the nature and sources of agglomeration economies." In: J.V. Henderson and J.-F. Thisse (eds.), *Handbook of Regional and Urban Economics*, vol. 4. North-Holland: Elsevier, pp.2119–2171.
- [38] Rosenthal, Stuart S., and William C. Strange. 2003. "Geography, industrial organization, and agglomeration." *Review of Economics and Statistics* 85(2): 377–393.
- [39] Storeygard, Adam. 2016. "Farther on down the road: Transport costs, trade, and urban growth in sub-Saharan Africa." *Review of Economic Studies* 83(3): 1263–1295.
- [40] Tanaka, Kiyoyasu, and Kenmei Tsubota. 2017. "Directional imbalance in freight rates: Evidence from Japanese inter-prefectural data." *Journal of Economic Geography* 17(1): 217–232.

‘The world is not yet flat: Transport costs matter!’

— Appendix material —

This set of appendices is structured as follows. **Appendix A** presents a simple framework of endogenously determined ad valorem transport costs and discusses its usefulness to highlight endogeneity problems. **Appendix B** documents additional data sources, provides complementary information, and explains the construction of our controls for trade exposure and input-output links. We also discuss potential endogeneity concerns related to these controls. **Appendix C** provides details on the [Lewbel \(2012\)](#) estimator and its implementation. **Appendix D** contains additional tables and results for geographic concentration patterns and transport costs. Last, **Appendix E** contains a large number of additional estimation results and a battery of robustness checks.

Appendix A. Endogenous transport costs

We develop a simple framework to model transport costs and to highlight possible identification issues. We are parsimonious in introducing the super- and subscripts of the model. In particular, we alleviate notation by suppressing subscripts when possible. Our aim is to guide the empirical analysis, not to provide a full-fledged model with all the bells and whistles.

A.1. A simple framework

Following [Behrens and Picard \(2011\)](#) and [Behrens and Brown \(2017\)](#), we consider a spatial economy with two regions, r and s . There are M_r manufacturers (shippers) in region r , and M_s manufacturers in region s . Without loss of generality, we assume that $M \equiv M_r/M_s \geq 1$, i.e., r is the larger region. Shipping goods requires the services of freight carriers who charge a *per unit freight rate* t_{rs} to ship commodities from region r to region s . In what follows, we focus on trucking as the shipping mode which is a highly competitive sector in Canada.²¹ We hence assume perfect competition between carriers who operate under constant returns to scale. We also abstract from market power that manufacturers may exercise over the transport sector and assume that they take freight rates as given.

Denote by m_r the marginal cost of a manufacturer in region r . He faces demand $Q_{rs}(m_r) \equiv Q(p_{rs}(m_r))$ in market s when quoting a delivered price denoted by $p_{rs}(m_r) \equiv p(m_r(\mathbf{X}_r^i), t_{rs})$, where \mathbf{X}_r^i is a vector of industry-region specific covariates like local factor prices or agglomeration economies that affect production costs. A carrier who transports

merchandise from r to s (the fronthaul part of the trip) needs to return from s to r (the backhaul part of the trip), whether the truck is fully loaded or not. The carrier will thus also transport merchandise from s to r conditional on its demand $Q(p_{sr}(m_s))$ at the price $p_{sr}(m_s) \equiv p(m_s(\mathbf{X}_s^i), t_{sr})$.²² Total demand for transport services from r to s and from s to r , conditional on t_{rs} and t_{sr} , is given by

$$D_{rs}(t_{rs}) = M_r Q(p_{rs}(m_r(\mathbf{X}_r^i), t_{rs})) \quad \text{and} \quad D_{sr}(t_{sr}) = M_s Q(p(m_s(\mathbf{X}_s^i), t_{sr})). \quad (\text{A.1})$$

Since carriers need to return to their point of departure, they face a logistics problem: they must commit to the capacity required by the largest demand on a return trip. Taking into account this backhaul problem, the carriers' profits are given by:

$$\pi(t_{rs}, t_{sr}) = S_{rs}t_{rs} + S_{sr}t_{sr} - 2\gamma(\mathbf{Y}^c, d_{rs}) \max\{S_{rs}, S_{sr}\}, \quad (\text{A.2})$$

where S_{rs} denotes the supply of transport services from r to s , and where $2\gamma(\mathbf{Y}^c, d_{rs})$ is the carriers' cost of a return trip that they must commit to. The function $\gamma(\cdot)$ depends on the distance d_{rs} of a one-way trip, and on a vector \mathbf{Y}^c of carrier- and commodity-specific factors like the carrier's productivity, diesel prices, wages, and the type c of commodity being transported.

A competitive equilibrium in the transport market is given by non-negative freight rates, t_{rs} and t_{sr} , and supplies, S_{rs} and S_{sr} , of transport services such that: (i) the carriers' supply profit-maximizing quantities of transport services, taking freight rates, goods prices, and the shippers' demand schedules as given; (ii) demand for transport services equals supply in each direction, i.e., $S_{rs} = D_{rs}$ and $S_{sr} = D_{sr}$; and (iii) carriers' profits (A.2) are maximized and equal to zero because of free entry. Using expression (A.2), profit maximization implies that if $S_{rs} > S_{sr}$, then $t_{rs} = 2\gamma(\mathbf{Y}^c, d_{rs})$ and $t_{sr} = 0$. The reverse holds if $S_{rs} < S_{sr}$ (see [Behrens and Picard, 2011](#), for details). Hence, $t_{rs} > 0$ and $t_{sr} > 0$ requires that $S_{rs} = S_{sr}$ and that $t_{rs} + t_{sr} = 2\gamma(\mathbf{Y}^c, d_{rs})$. Put differently, transport costs in both directions are strictly positive if and only if freight rates adjust to balance flows: $D_{rs}(t_{rs}) = D_{sr}(t_{sr})$. When $D_{rs}(2\gamma) > D_{sr}(0)$, the difference in demand for transport services is so large that freight rates on the backhaul part of the trip effectively fall to zero. Zero freight rates are of course an extreme case in our simple model, but it captures the idea that carriers are willing to transport at steep discounts in the direction of substantial excess capacity.

To derive simple expressions for freight rates and *ad valorem* transport costs, assume that manufacturers are monopolistically competitive and face constant elasticity (CES) demand schedules. Their profit-maximizing prices on the fronthaul and the backhaul parts of the trip, conditional on freight rates, are then given by:

$$p_{rs}^i = \frac{\sigma_i}{\sigma_i - 1} (m_r^i + t_{rs}), \quad \text{and} \quad p_{sr}^i = \frac{\sigma_i}{\sigma_i - 1} (m_s^i + t_{sr}), \quad (\text{A.3})$$

where $m_r^i \equiv m_r(\mathbf{X}_r^i)$ and $m_s^i \equiv m_s(\mathbf{X}_s^i)$ to alleviate notation; and where σ_i denotes the industry-specific (constant) price elasticity of demand the manufacturers' face.²³

Given CES demands, we have $Q_{rs}^i = A \cdot (p_{rs}^i)^{-\sigma_i}$ and $Q_{sr}^i = A \cdot (p_{sr}^i)^{-\sigma_i}$, where A is a demand shifter.²⁴ To balance flows in both directions thus requires that

$$M^{-1/\sigma_i} [(m_r^i + t_{rs})] = m_s^i + 2\gamma(\mathbf{Y}^c, d_{rs}) - t_{rs}, \quad (\text{A.4})$$

where we have used expressions (A.1), (A.3), and the condition $t_{rs} + t_{sr} = 2\gamma(\mathbf{Y}^c, d_{rs})$. Solving (A.4), we obtain the fronthaul freight rate:

$$t_{rs} = \frac{1}{1 + M^{-1/\sigma_i}} [m_s^i - M^{-1/\sigma_i} m_r^i + 2\gamma(\mathbf{Y}^c, d_{rs})]. \quad (\text{A.5})$$

The backhaul freight rate can be recovered from the zero profit conditions $t_{rs} + t_{sr} = 2\gamma(\mathbf{Y}^c, d_{rs})$. The ad valorem rate is given by $\tau_{rs} = 1 + t_{rs}/m_r^i = 1 + \frac{\sigma_i - 1}{\sigma_i} t_{rs}/p_r^{i,\text{prod}}$, where $p_r^{i,\text{prod}} = [\sigma_i/(\sigma_i - 1)]m_r$ is the producer price of the good. We thus obtain:

$$\tau_{rs} = \frac{1}{1 + M^{-\frac{1}{\sigma_i}}} \left[1 + \frac{m_s^i}{m_r^i} + \frac{2\gamma(\mathbf{Y}^c, d_{rs})}{m_r^i} \right] = \frac{1}{1 + M^{-\frac{1}{\sigma_i}}} \left[1 + \frac{\sigma_i}{\sigma_i - 1} \left(\frac{m_s^i}{p_r^{i,\text{prod}}} + \frac{2\gamma(\mathbf{Y}^c, d_{rs})}{p_r^{i,\text{prod}}} \right) \right], \quad (\text{A.6})$$

which is a key object in our empirical analysis and which is useful to highlight possible identification issues that we need to take care of.

A.2. Identification issues

For β_τ in our empirical analysis to capture the causal effect of transport costs on geographic concentration, we need to address a number of identification problems. The foregoing framework is useful to understand these problems, which are essentially driven by M , m_s/m_r , and $\gamma(\mathbf{Y}^c, d_{rs})$.

A first problem is due to agglomeration economies. Viewed through the lense of our conceptual framework, if $1/m_r > 1/m_s$ — i.e., firms in region s are more productive than firms in region r — and assuming that $M = 1$, we have $t_{rs} > \gamma(\mathbf{Y}^c, d_{rs}) > t_{sr}$: freight rates are larger on the fronthaul trip than on the backhaul trip, because of regional productivity differences that map into regional producer price differences that affect the flows of goods. The empirical literature has substantiated a significant causal effect of regional size, M_r , on productivity, $1/m_r$.²⁵ Geographic concentration may hence affect transport costs: (i) either directly, by decreasing carriers' costs $\gamma(\mathbf{Y}^c, d_{rs})$ through density economies (see [Mori and Nishikimi, 2002](#); [Tanaka and Tsubota, 2017](#)); (ii) or through market interactions, by increasing demand for transport services (via productivity gains), which increases output of goods to be shipped; (iii) or indirectly by increasing

productivity of manufacturers, which maps into lower prices and, therefore, larger ad valorem rates (see equation (A.6)).

A second related problem arises because geographic concentration leads to imbalances in shipping patterns, and the latter increase transport costs due to standard logistics problems like ‘backhaul’ of empty trucks (see, e.g., [Jonkeren, Demirel, van Ommen, and Rietveld, 2009](#), and [Tanaka and Tsubota, 2017](#), for empirical evidence). In terms of our conceptual framework, if markets are symmetric—equal size $M = 1$ and productivity $m_r = m_s$ —then $t_{rs} = t_{sr} = \gamma(\mathbf{Y}^c, d_{rs})$. In that case, freight rates are symmetric and equal to the marginal cost of the carriers on each leg of the trip. However, when $M > 1$, i.e., $M^{-1/\sigma_i} < 1$, then market r is larger than market s , so that there is more demand for transport services from r to s than the reverse. Assume further that $m_r = m_s$. Then we have from (1) that $t_{rs} > \gamma(\mathbf{Y}^c, d_{rs}) > t_{sr}$. Freight rates are higher on the fronthaul trip than on the backhaul trip, because of imbalances in shipping patterns due to market size. These imbalances force carriers to slash rates on the backhaul in order to fill their trucks. Yet, by doing so they change firms’ locational incentives, giving a transport cost advantage to the smaller region (see [Behrens and Picard, 2011](#), for a theoretical analysis). Hence, the geographic concentration of industries, by decreasing M , directly affects transport costs via the carriers’ backhaul problem.²⁶

To summarize, $\tau_{i,t}$ in (9) is potentially endogenous to the geographic concentration of an industry, with stronger concentration increasing transport costs due to a combination of rising freight prices t_{rs} and lower industry output prices $p_r^{i,\text{prod}}$. Thus, the OLS estimate of β_τ is likely to be upward biased in our empirical model.

Appendix B. Additional information on data, variables, and descriptives

B.1. Changes in the ASM sampling frame

The survey frame of the ASM has evolved over time. Early in the period, it was relatively stable with, on average, about 32,000 plants per sample year. The sample of plants was restricted to those that report employment and have sales in excess of \$30,000. Also, aggregate records were excluded. These records represent multiple (typically small) plants without latitudes and longitudes. In 2000, however, the number of plants in the survey increased substantially as the ASM moved from its own frame to Statistics Canada’s centralized Business Register, increasing the sample to an average of 53,000 plants. In 2004, the number of plants in the frame was once again restricted, with many of the small plants excluded, or included in aggregate records. With this in place, the

sample returned to near previous levels, averaging about 33,000 plants between 2004 and 2009. We find in our analysis that the expanded survey scope in the early 2000s had little effect on aggregate trends. Our analysis also deals with the change in the sample frame through the inclusion of year dummies.

B.2. Additional information on the PCCF

The Postal Code Conversion Files associate each postal code with different Standard Geographical Classifications (sgc) that are used for reporting census data in Canada. We match firm-level postal code information with geographic coordinates from the PCCF. Postal codes are less fine grained in predominantly rural areas, but the kernel smoothing of our geographic concentration measures takes care of these variations (see [Duranton and Overman, 2005](#), for additional details). Figure 6 provides an illustration of the granularity of our data, zooming onto the eastern part of Ontario and Quebec.

B.3. U.S. industry price indices

The U.S. price indices at the NAICS 6-digit level are from the NBER-CES Manufacturing Productivity Database (<http://nber.org/data/nberces5809.html>).

B.4. Controls for industry structure, other agglomeration effects, and natural advantages

Various organizational industry characteristics—industry size, the mean plant size, the size distribution of plants in the industry, the presence of multi-unit firms, or foreign ownership—are likely to affect geographic concentration (see, e.g., [Rosenthal and Strange, 2003](#)). We use plant-level information from the ASM to construct controls for: (i) industry size and structure (‘total industry employment’; ‘mean plant size’; ‘Herfindahl index of enterprise-level employment concentration’);²⁷ and (ii) the industry’s ownership structure (‘share of plants controlled by multiunit firms’; ‘share of foreign-controlled plants’). All variables are constructed from plant-level data by aggregating to the industry level. Table 5 summarizes the descriptive statistics for these variables.

We also use a proxy related to workers’ educational attainment to control for labor market aspects and industries’ labor force composition. More specifically, we use the share of hours worked by all workers with post-secondary education in the total number of hours worked in the industry. This measure is constructed by aggregating plant-level data.²⁸

Table 5: Summary statistics for the control variables and the instrument.

Variable names and descriptions	Industry detail	Mean	Standard deviation		
			Overall	Between	Within
Share of industry imports from Asian countries (excluding OECD members)	NAICS6	0.120	0.183	0.172	0.062
Share of imports from OECD member countries (excluding U.S. and Mexico)	NAICS6	0.157	0.141	0.131	0.053
Share of imports from NAFTA countries (U.S. and Mexico)	NAICS6	0.662	0.273	0.263	0.074
Share of industry exports to Asian countries (excluding OECD members)	NAICS6	0.029	0.058	0.047	0.035
Share of exports to OECD member countries (excluding U.S. and Mexico)	NAICS6	0.086	0.101	0.085	0.054
Share of exports to NAFTA countries (U.S. and Mexico)	NAICS6	0.833	0.198	0.184	0.073
Industry mean of avg. distance to a dollar of inputs from 5 nearest plants (km)	NAICS6	241	111	95	57
Industry mean of avg. distance to ship a dollar of output to 5 nearest plants (km)	NAICS6	243	124	103	69.1
Minimum average distance to 5×257 closest plants	NAICS6	64.7	44.5	42.1	14.4
Share of inputs from natural resource-based industries	L-level	0.113	0.171	0.170	0.026
Sectoral energy inputs as a share of total sector output	L-level	0.032	0.046	0.045	0.013
Total industry employment	NAICS6	7,038	8,060	7,858.11	1,856
Herfindahl index of enterprise-level employment concentration	NAICS6	0.101	0.097	0.092	0.032
Mean plant size	NAICS6	73.7	145	139	41.8
Share of plants controlled by multi-plant firms	NAICS6	0.212	0.193	0.183	0.061
Share of foreign controlled plants	NAICS6	0.153	0.157	0.146	0.059
Share of hours worked by all workers with post-secondary education	NAICS6	0.401	0.082	0.071	0.041
Intramural research and development expenditures as a share of industry sales	L-level	0.011	0.039	0.027	0.005
Minimum distance from major container ports (km)	NAICS6	414	110	103	48
Eastern share of plants	NAICS6	0.749	0.133	0.124	0.047
Ad valorem trucking costs (instrument based on U.S. price indices)	NAICS6	0.038	0.034	0.034	0.006

Notes: All descriptive statistics are for the sample that we use in the regression analysis, which includes 4,369 observations covering 257 industries and 17 years. The standard deviation is decomposed into a between and a within component, which measure the cross sectional and the time series variation, respectively. Some industry-level data—especially in the KLEMS database—are available at the L-level only, which is the finest level of data for public release in Canada (between the NAICS 3- and 4-digit levels of aggregation).

Finally, additional industry-level information from the KLEMS database is used to construct proxies for knowledge spillovers and natural advantage. We control for knowledge spillovers using as a proxy an industry’s research and development (R&D) intensity, i.e., the ratio of R&D expenditure in an industry to its total output (see [Rosenthal and Strange, 2001](#)). Furthermore, industries may also concentrate geographically because of localized natural advantages (see, e.g., [Ellison and Glaeser, 1999](#); and [Ellison et al. 2010](#)). We control for industries’ reliance on natural advantages using the share of inputs from natural resource-based industries, and the sectoral energy inputs as a share of total sector output.

B.5. International trade exposure

Data sources. The industry-level trade data come from Innovation, Science and Economic Development Canada’s Online Trade Database and cover the years 1992 to 2009. The dataset reports imports and exports at the NAICS 6-digit level by province and by country of origin and destination. We aggregate the data across provinces and compute the shares of exports and imports that go to or originate from a set of country groups:

Asian countries (excluding OECD), OECD countries (excluding NAFTA), and NAFTA countries. Since the trade data is only available from 1992 on, whereas the KLEMS data is only available until 2008, we restrict our sample to the 1992–2008 period in all estimations to maintain comparability of results.

Construction. We use detailed yearly data on imports and exports by industry and country of origin and destination to control for industries’ import and export exposure (the ratio of industry imports or exports to industry sales). To somewhat disentangle the different effects that depend on whether trade is in intermediates or final goods (on which we have unfortunately no information in our data), and on whether trade is ‘North-North’ or ‘North-South’, we break these measures down by countries of origin: low-cost Asian countries; OECD countries; and NAFTA countries. This break down is also meaningful to distinguish between dominant modes of transportation (seaborne for Asia, and land borne for NAFTA), an aspect we look at in Section 5.3 of the paper.

Figure 3: Changes in import- and export trade values (left), and import shares (right).



The left panel of Figure 3 depicts the changes in the average import and export values by industry over our study period. The right panel provides a snapshot of how import and export shares change across broad groups of trading partners. As one can see from Figure 3, the volume of international trade has increased—up to the great trade collapse of 2008—and there has been an increasing re-orientation of trade towards Asian countries (especially for imports, which have overtaken the OECD non-NAFTA share of imports).

Endogeneity concerns. The geographic concentration of plants increases productivity and, therefore, may increase the propensity of an industry to export or to import. For example, the agglomeration of an industry may reduce prices, which makes import

penetration harder. In that case, the dispersion of an industry may be associated with increasing imports since productivity falls. Also, the agglomeration of an industry may be associated with rising exports due to productivity gains—although the productivity gains reduce unit export values, the total value of exports may still increase. Since external instruments for import and export exposure of industries are difficult to find, we deal with potential endogeneity issues of trade exposure using the Lewbel (2012) estimator that relies on internal instruments.²⁹

B.6. Input-output distances

Our proxies for input and output links are of the reduced-form type and not structural (unlike, e.g., the ‘structural supplier and market access’ in [Redding and Venables, 2004](#)). As pointed out by [Combes and Gobillon \(2015, p.274\)](#), there is generally no satisfying solution to control for supplier and market access in empirical estimation. Either we use a structural model, which requires many assumptions and has its own limitations, or we use reduced-form proxies that aim at capturing those interactions. We choose the latter approach.

Data sources. We use the L -level national input-output tables from Statistics Canada at buyers’ prices to construct our plant-level proxies for the importance of input and output linkages. These tables—which constitute the finest sectoral public release—feature 42 sectors at an aggregation level somewhere in between the NAICS 3- and 4-digit levels. We keep the manufacturing portion only and break them down to the 6-digit level based on industries’ weights in terms of sales. For each industry, i , we allocate total inputs purchased or outputs sold in the L -level matrix to the corresponding NAICS 6-digit sectors. We allocate total sales to each subsector in proportion to that sector’s sales in the total sales to obtain a 257×257 matrix of NAICS 6-digit inputs and outputs, which we use to construct the linkages.³⁰

Construction. We construct novel proxies for distance to inputs and outputs that make use of the micro-geographic nature of our data. Consider a plant ℓ active in industry $\Omega(\ell)$. Let Ω denote the set of industries and Ω_i the set of plants in industry i . Let $k_i(r, \ell)$ denote the r -th closest industry- i plant to plant ℓ . Our micro-geographic measures of input- and output linkages are constructed as weighted averages as follows:

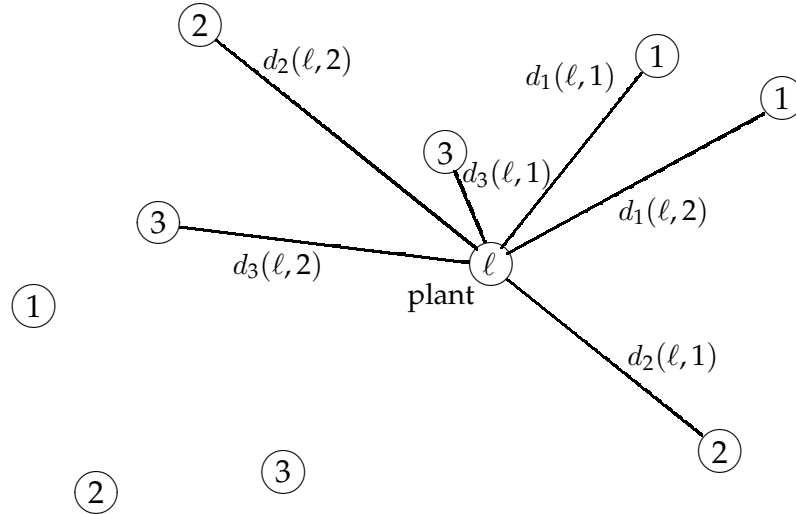
$$\mathcal{I}\text{dist}(\ell) = \sum_{i \in \Omega \setminus \Omega(\ell)} \omega_{\Omega(\ell), i}^{\text{in}} \times \frac{1}{N} \sum_{j=1}^N d(\ell, k_i(j, \ell)), \quad (\text{B.1})$$

for inputs, and

$$\mathcal{O}\text{dist}(\ell) = \sum_{i \in \Omega \setminus \Omega(\ell)} \omega_{\Omega(\ell),i}^{\text{out}} \times \frac{1}{N} \sum_{j=1}^N d(\ell, k_i(j, \ell)), \quad (\text{B.2})$$

for outputs, where $d(\cdot, \cdot)$ is the great circle distance between the plants' postal code centroids, and where $\omega_{\Omega(\ell),i}^{\text{in}}$ and $\omega_{\Omega(\ell),i}^{\text{out}}$ are sectoral input- and output shares computed using the 6-digit versions of the L -level input-output tables described in the foregoing. We exclude within-sector transactions where $\Omega(\ell) = i$ as those may be capturing all sorts of intra-sectoral agglomeration economies that are conducive to clustering but not correlated with input-output linkages. Figure 4 illustrates the construction of the input- and output links (B.1) and (B.2) for the case where $N = 2$ and with three industries (labeled 1, 2, and 3).

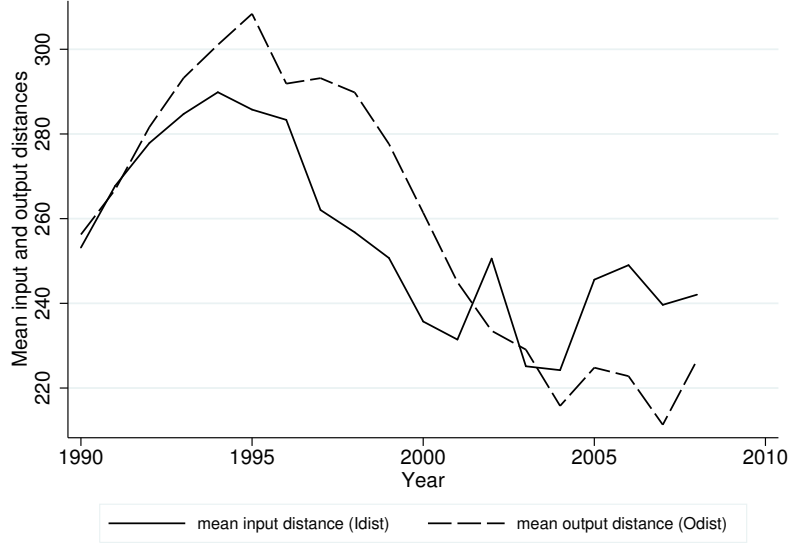
Figure 4: Constructing input-output distances and ‘minimum distance’ measures.



Observe that since $\sum_i \omega_{\Omega(\ell),i}^{\text{in}} = \sum_i \omega_{\Omega(\ell),i}^{\text{out}} = 1$ by construction, we can interpret $\mathcal{I}\text{dist}(\ell)$ as the minimum average distance of plant ℓ to a dollar of inputs from its N closest manufacturing suppliers. Analogously, $\mathcal{O}\text{dist}(\ell)$ is the minimum average distance plant ℓ has to ship a dollar of outputs to its N closest (industrial) customers.³¹ The larger are $\mathcal{I}\text{dist}(\ell)$ or $\mathcal{O}\text{dist}(\ell)$, the worse are plant ℓ 's input or output linkages – it is, on average, further away from a dollar of intermediate inputs or a dollar of demand emanating from the other industries. Note that our input and output linkages make use of *plant-level* location information, but only of *national* input and output shares. The latter is due to the fact that we do not directly observe input-output relationships at the plant level. Yet, given this, our procedure has the advantage to sidestep obvious problems of endogeneity of those plant-level input-output relationships. Furthermore, our

input-output measures are computed across all industries except the one the plant belongs to. Thus, our measures capture finely the whole cross-industry location patterns, but do not pick up industrial localization of the sector itself since it is excluded from the computation.

Figure 5: Changes in average input-output distances, 1990–2009.



We compute (B.1) and (B.2) for all years and for all plants, using the $N = 3, 5, 7$, and 10 nearest plants in each industry. We then average them across plants in each industry i and each year to get an industry-year specific measure of both input and output distances:

$$\mathcal{O}dist_i = \frac{1}{|\Omega_i|} \sum_{\ell \in \Omega_i} \mathcal{O}dist(\ell) \quad \text{and} \quad \mathcal{I}dist_i = \frac{1}{|\Omega_i|} \sum_{\ell \in \Omega_i} \mathcal{I}dist(\ell), \quad (\text{B.3})$$

where $|\Omega_i|$ denotes the number of plants in industry i . As expected, (B.1) and (B.2) are strongly correlated. Yet, despite that correlation we can include them simultaneously into our regressions and still identify their effect on geographic concentration.

Figure 5 depicts the time-series changes in the (unweighted) average input and output measures across all industries. As one can see, in 2000 for example, plants were on average located about 235 kilometers from a dollar of inputs, and had to ship a dollar of their output on average over a distance of 260 kilometers. Note that time-series changes in the input- and output-distance measures may reflect three things: (i) entry or exit of potential suppliers; (ii) changes in the geographic location of input suppliers and/or clients; and (iii) changes in the input-output coefficients, i.e., the technological relationships. We cannot disentangle the sources (i) and (ii) in our analysis, but, as explained in

the paper, entry and exit are more important than relocation when looking at plant-level data.

As can be seen from Figure 5, average input and output distances have fallen over the 1990–2009 period in Canada, from about 260 kilometers to about 240 kilometers. One may wonder how this result is compatible with our finding that industries have geographically dispersed, as documented in Section 3.2. To understand that result, one needs to keep in mind that the geographic dispersion we document in Section 3.2 is for *within-industry concentration*, whereas the measures of input-output distances are for *between-industry concentration*. Starting from a situation where industries are spatially segregated would yield a large value of within-industry geographic concentration, and large between-industry distances. As industries progressively disperse, the within-industry measure falls, whereas the between-industry distance can fall too if there is more ‘mixing’ of industries. In a nutshell, if there is less segregation and more mixing between industries, the geographic concentration of industries would fall, but their distance to input suppliers and clients can decrease too. Hence, the two findings are not incompatible.

Note, finally, that one potential problem with the measures in (B.3) is that they are mechanically smaller in denser areas. To control for this fact, we also compute a ‘minimum distance measure’, i.e., the distance of plant ℓ from the $M = N \times 257$ closest plants, regardless of their industry. Including that measure into our regressions then controls for the overall plant density in a location, which implies that our input-output linkage measures pick up the effect of being closer to a dollar of inputs or outputs conditional on the overall density of the area the plant is located in. Formally, we compute for each plant ℓ the following measure:

$$\mathcal{M}\text{dist}(\ell) = \frac{1}{M} \sum_{j=1}^M d(\ell, k_{\setminus \Omega(\ell)}(j, \ell)), \quad (\text{B.4})$$

where $d(\ell, k_{\setminus \Omega(\ell)}(j, \ell))$ denotes the distance to the j th closest plant in *any industry but* $\Omega(\ell)$. We then average this measure across all plants in the same industry as before and include it as an additional control into our regressions.

Endogeneity concerns. Our measures of input-output linkages are, by construction, reasonably exogenous to the spatial structure of a specific industry. First, observe that we compute those measures using national input-output shares instead of plant-level input-output shares. Hence, we do not pick up spuriously large values for inputs or outputs—due to substitution effects—when plants are located in close proximity to plants in related industries. Second, we exclude the own industry from the computation, so that the measures only pick up cross-industry links and not the geographic

concentration of the industry itself (which is on the left-hand side of our regressions).³² Last, for each plant, the input and output distance is computed using *all other 256 industries in Canadian manufacturing*. For the geographic concentration of one industry to drive the input-output linkage measure, that industry would need to substantially affect the whole location patterns of most other related industries, which strikes us as fairly unlikely (though we cannot completely rule out this possibility). Although the input- and output-measures should be reasonably exogenous, we deal with potential endogeneity issues of input-output links using the [Lewbel \(2012\)](#) estimator that relies on internal instruments. External instruments are hard to find for these measures.

Appendix C. Applying the Lewbel (2012) method

To apply the [Lewbel \(2012\)](#) procedure, we need to verify two conditions: heteroscedasticity and correlation. First, we regress the potentially endogeneous variables (input and output distances, as well as trade shares) on all other exogeneous variables of the model. We then predict the residuals of that regression and run a standard heteroscedasticity test. We need to reject the homoscedasticity assumption for the Lewbel method to be applicable. In our case, we strongly reject the null hypothesis of homoscedasticity for all series of residuals (the p -value is zero in all tests). Second, we take the square of the predict residuals from the foregoing regression, and check the correlation between the dependent variable of the regression (input distances, or output distances, or the different trade shares) and those squared residuals. The correlation needs to be ‘strong’ and statistically strongly significant for the instruments to work properly. In our case, this condition holds true for the input and output distances, and for all import shares: the correlation of the squared residuals with the variable itself is significant at 1% in all cases. It is 0.141 for transportation costs, -0.081 for input distances, -0.089 for output distances, 0.130 for the Asian share of imports, and -0.079 for the NAFTA share of imports. We find no statistically significant correlation for the export shares.

Since the two conditions—heteroscedasticity of the residuals and correlation of the squared residuals with the variable—are met in our case, we can apply the Lewbel estimator. Since fixed effects cannot be included in the estimation (see `ivreg2h` in Stata), we de-mean all variables by industry first. The exogeneous variables are partialled-out for the Lewbel estimator and so their coefficients are not reported. Since we have an exogeneous instrument for transportation costs, we apply the Lewbel estimator only to deal with potential endogeneity concerns of trade shares and input-output distances.

Appendix D. Additional tables and results for geographic concentration and ad valorem transport costs

We first provide additional results for geographic concentration patterns, especially for employment- or sales-weighted measures. We then provide additional descriptive evidence for ad valorem transport costs.

D.1. Geographic concentration

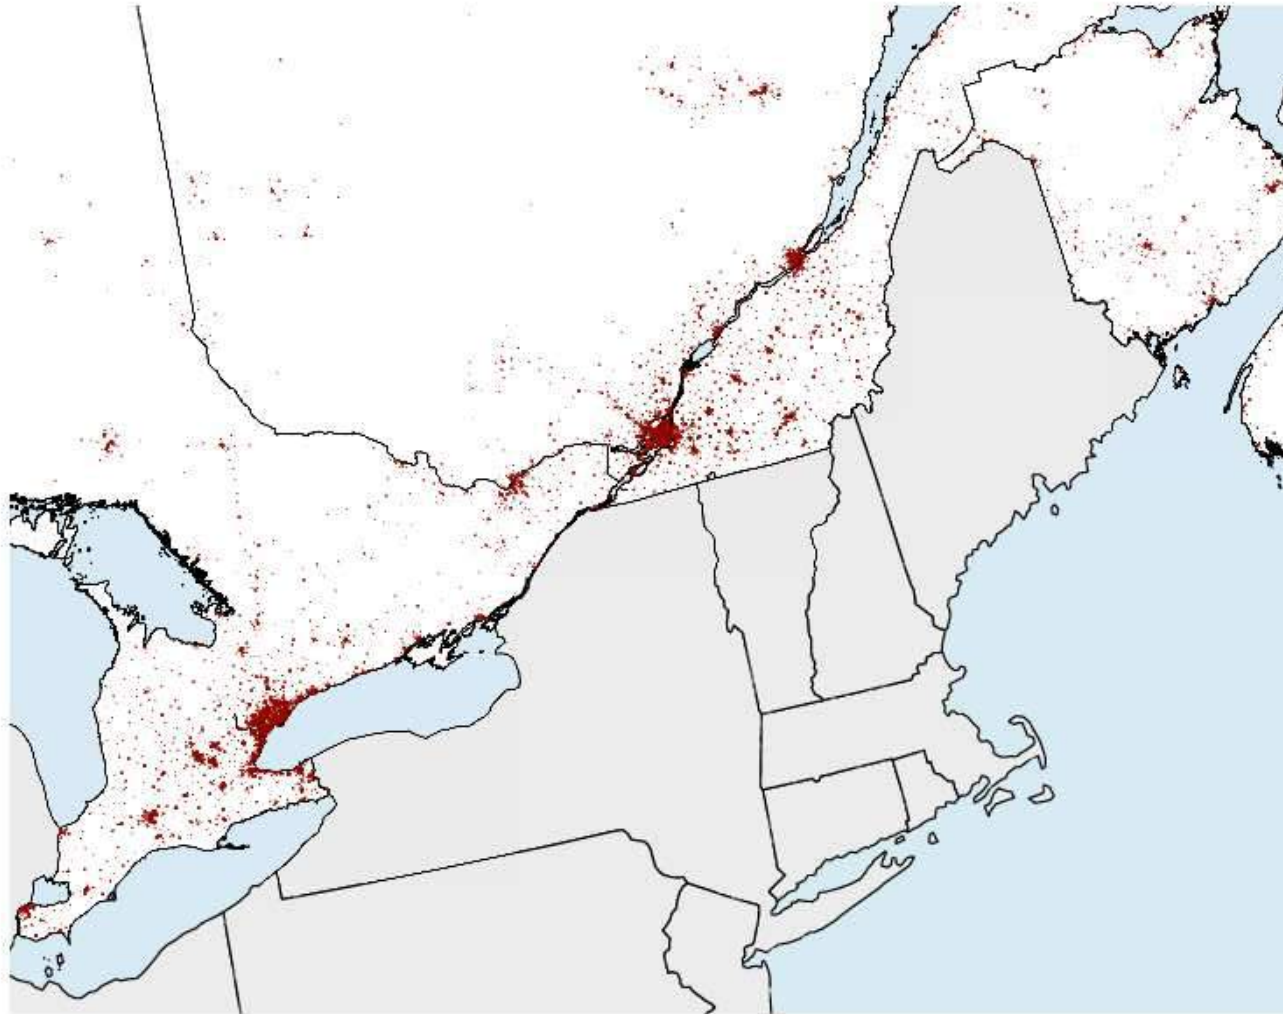
Data. Figure 6 illustrates the geographic nature of our data, with each red dot corresponding to a manufacturing establishment in 2001. Note that the data we depict are based on the manufacturing portion of the Scott's Directories National All database, and not on the ASM Longitudinal Microdata file that we use. Both datasets draw on the business register and are fairly comparable in terms of industry and geographic coverage (see [Behrens and Bougna, 2015](#)). We use the Scott's data to depict the spatial structure of manufacturing in Figure 6 since confidentiality reasons preclude us from using the ASM Longitudinal Microdata file for that purpose.

Changes over time. Table 6 summarizes our results for all 6-digit industries year-by-year from 1990 to 2009 and assess the significance of the geographic concentration patterns.³³ As can be seen from panel (A) of Table 6, the absolute geographic concentration of manufacturing has substantially decreased in Canada between 1990 and 2009 (see also [Behrens and Bougna, 2015](#)). While on average 7.6% of the bilateral distances between plants in manufacturing industries were less than 50 kilometers in 1990, only 5.6% of those distances still remained within that range in 2009, a 27.1% decrease.³⁴ The results using either employment or sales weights are very similar, with a little bit less deconcentration. Geographic concentration is stronger in terms of employment than in terms of plants, and even stronger in terms of sales than in terms of employment. Figure 7 further illustrates that point.

Panel (A) of Table 6 further shows that concentration has decreased more at shorter distances, i.e., the incentives for plants to locate in very close spatial proximity to each other has weakened over time. This likely reflects the fact that manufacturing has been bid out of cities because of higher land and labor costs there (see, e.g., [Henderson, 1997](#)). That trend also affects the employment-weighted and the sales-weighted measures, but to a lesser extent.

Panel (B) of Table 6 shows that there is also some evidence that relative geographic concentration has decreased over our study period, albeit by less than absolute geo-

Figure 6: Geographic distribution of manufacturing establishments in the south-eastern part of Canada, 2001.



Notes: Spatial distribution of manufacturing establishments in Canada in 2001, based on the Scott's Directories National All database (manufacturing portion). Those data draw on the business register, as does the ASM Longitudinal Microdata file, and are fairly comparable to the latter. We use them to depict the spatial structure of manufacturing since confidentiality reasons preclude us from using the ASM Longitudinal Microdata file.

graphic concentration. While 34.63% of industries were significantly concentrated compared to manufacturing in general in 1990, that share slightly decreased to 33.07% in 2009. The effects were a bit stronger, albeit qualitatively similar, when using either sales- or employment-weighted measures.

Comparing the unweighted to either the employment- or the sales-weighted K -densities reveals some interesting patterns. As can be seen from Figure 7, industries are on average always more concentrated in terms of employment than in terms of plant counts, and even more concentrated in terms of sales than in terms of employment. This is a manifestation of agglomeration economies, and it is consistent with the findings that more localized plants tend to be larger and more productive than less localized plants (e.g., [Holmes and Stevens, 2014](#)). Note that the ratios are increasing until about 2004, and slightly decreasing afterwards. In 2009, within 50 kilometers distance, the concentration of employment exceeds that of plant counts by about 13%, whereas the concentration of sales exceeds that of plant counts by about 20%. Note further that the de-concentration trend that we documented for the plant-count based measures also affects the employment-weighted and the sales-weighted measures of geographic concentration (see Table 7). Yet, as can be seen from Figure 7, although industries have in general become more geographically dispersed according to all three measures, the size of plant pairs in close proximity has tended to increase in relative terms regardless of whether size is measured by employment or by sales. Put differently, the process of dispersion is less pronounced when geographic concentration is measured by either employment or sales, thus suggesting that smaller plants drive a substantial part of the dispersion process, either through entry and exit or through relocation.

Figure 7: Ratios of mean employment- and sales-based CDFs to count-based CDF by distance.

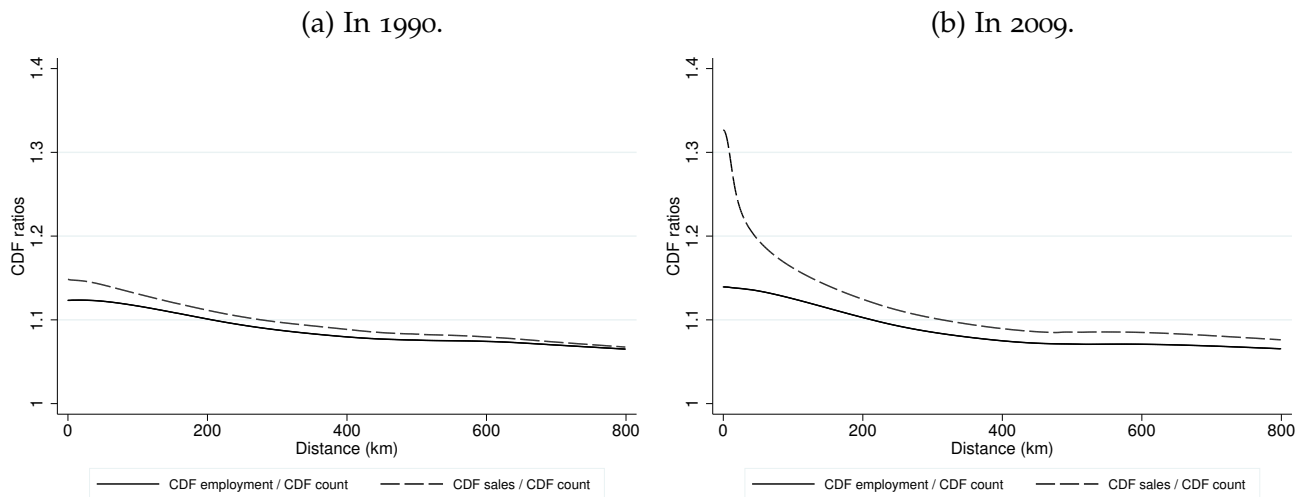


Table 8 below provides the (unweighted) K -density CDFs in 1990, 1999, and 2009 for

Table 6: Summary of geographic concentration patterns of Canadian manufacturing industries, 1990–2009.

	(1) Unweighted				(2) Employment weighted				(3) Sales weighted			
	(A) Mean of the cumulative K -densities across industries, 1990 and 2009: CDF at a distance of											
Year	10 km	50 km	100 km	500 km	10 km	50 km	100 km	500 km	10 km	50 km	100 km	500 km
1990	0.020	0.076	0.139	0.420	0.021	0.083	0.151	0.449	0.022	0.086	0.156	0.453
2009	0.013	0.056	0.107	0.373	0.015	0.063	0.121	0.397	0.017	0.068	0.126	0.403
Mean 1990–2009	0.015	0.064	0.121	0.394	0.017	0.073	0.136	0.422	0.019	0.077	0.141	0.428
% change	-36.0%	-27.1%	-22.6%	-11.3%	-28.7%	-23.3%	-20.3%	-11.4%	-21.5%	-21.2%	-19.3%	-11.0%
	(B) Share of industries with random, localized, and dispersed point patterns, 1990 and 2009.											
Year	Random	Localized	Dispersed		Random	Localized	Dispersed		Random	Localized	Dispersed	
1990	52.53	34.63	12.84		52.53	36.96	10.51		54.86	37.35	7.78	
2009	59.53	33.07	7.39		63.04	31.52	5.45		63.04	31.13	5.84	

Notes: Authors' computations based on the Annual Survey of Manufacturers Longitudinal Microdata file, 1990–2009. **Panel (A):** We report the values for the starting and the end years only since the series in between change rather smoothly (see Table 7 in the supplemental appendix for the full set of results). The means of the cumulative K -densities are based on 257 industries and are not weighted (but the CDFs for each industry are weighted by either employment in the middle columns (2), or by sales in the right columns (3)). 'Mean 1990–2009' refers to the temporal mean of the K -densities over the whole 1990–2009 period. '% change' is the percentage change between 1990 and 2009. **Panel (B):** The statistical significance of the location patterns is computed using Monte Carlo simulations with 1,000 replications following the procedure developed by [Duranton and Overman \(2005\)](#). We report the values for the starting and the end years only since the series in between decrease rather smoothly (see Table 9 in the supplemental appendix for the full set of results).

the geographically most strongly concentrated industries in Canada. As can be seen, textile and clothing related industries rank very high in that table, which thus explains why we run robustness checks later to exclude them in Table 15. Table 9 summarizes the year-on-year location patterns of industries based on the formal significance test of [Duranton and Overman \(2005\)](#). As can be seen from the table, the number of significantly localized industries has fallen over time, whereas the number of industries that display location patterns that are not significantly different from random ones has increased. These patterns hold for all geographic concentration measures (unweighted and employment or sales weighted).

Table 7: Mean of the Duranton-Overman CDFs across industries, 1990 to 2009.

Year	(1) Unweighted				(2) Employment weighted				(3) Sales weighted			
	10 km	50 km	100 km	500 km	10 km	50 km	100 km	500 km	10 km	50 km	100 km	500 km
1990	0.020	0.076	0.139	0.420	0.021	0.083	0.151	0.449	0.022	0.086	0.156	0.453
1991	0.019	0.076	0.139	0.423	0.022	0.083	0.152	0.447	0.023	0.087	0.156	0.453
1992	0.020	0.074	0.135	0.418	0.020	0.079	0.147	0.442	0.022	0.084	0.151	0.448
1993	0.019	0.072	0.132	0.416	0.020	0.079	0.145	0.440	0.021	0.082	0.148	0.446
1994	0.017	0.071	0.131	0.413	0.020	0.077	0.143	0.438	0.021	0.081	0.147	0.443
1995	0.017	0.068	0.126	0.402	0.019	0.076	0.141	0.432	0.020	0.080	0.145	0.438
1996	0.016	0.065	0.122	0.402	0.019	0.073	0.136	0.428	0.020	0.076	0.140	0.435
1997	0.016	0.066	0.123	0.401	0.017	0.072	0.135	0.427	0.019	0.077	0.140	0.433
1998	0.016	0.064	0.120	0.396	0.019	0.074	0.135	0.425	0.019	0.078	0.141	0.433
1999	0.015	0.062	0.118	0.398	0.017	0.072	0.134	0.426	0.018	0.076	0.139	0.434
2000	0.014	0.063	0.120	0.383	0.016	0.073	0.135	0.411	0.016	0.075	0.140	0.421
2001	0.013	0.061	0.118	0.383	0.015	0.072	0.136	0.412	0.016	0.076	0.142	0.421
2002	0.013	0.062	0.119	0.383	0.016	0.073	0.137	0.413	0.017	0.078	0.143	0.422
2003	0.013	0.060	0.117	0.384	0.015	0.072	0.137	0.416	0.016	0.075	0.141	0.422
2004	0.013	0.060	0.115	0.379	0.015	0.070	0.132	0.412	0.017	0.074	0.137	0.418
2005	0.012	0.059	0.113	0.379	0.014	0.068	0.130	0.409	0.016	0.072	0.134	0.415
2006	0.013	0.061	0.116	0.378	0.015	0.069	0.131	0.406	0.015	0.072	0.135	0.412
2007	0.012	0.057	0.110	0.374	0.015	0.064	0.122	0.399	0.017	0.069	0.127	0.406
2008	0.012	0.057	0.110	0.376	0.017	0.067	0.125	0.400	0.017	0.069	0.128	0.405
2009	0.013	0.056	0.107	0.373	0.015	0.063	0.121	0.397	0.017	0.068	0.126	0.403
Mean	0.015	0.064	0.121	0.394	0.017	0.073	0.136	0.422	0.019	0.077	0.141	0.428
Change	-36.0%	-27.1%	-22.6%	-11.3%	-28.7%	-23.3%	-20.3%	-11.4%	-21.5%	-21.2%	-19.3%	-11.0%

Notes: Authors' computations based on the Annual Survey of Manufacturers Longitudinal Microdata file, 1990–2009. The means of the CDF are based on 257 industries and are not weighted. Column (1) reports the count-based measure, whereas columns (2) and (3) report measures weighted by employment and by sales, respectively. See Section 3.2 for details. 'Mean' refers to the mean of the K -densities over the 1990–2009 period. 'Change' is the percentage change between 1990 and 2009.

D.2. Changes in Canadian manufacturing

Table 10 summarizes broad trends in the Canadian manufacturing sector. As shown, manufacturing fell substantially as a share of overall employment, although the absolute decline over the 1992–2016 period was much more limited. Thus, our main results should be viewed in the context of a shrinking manufacturing sector in relative terms, but not in

Table 8: Ten most localized NAICS 6-digit industries (based on plant counts).

NAICS	Industry description	CDF
(A) 1990		
315231	Women's and Girls' Cut and Sew Lingerie, Loungewear and Nightwear Manufacturing	0.62
315233	Women's and Girls' Cut and Sew Dress Manufacturing	0.55
313240	Knit Fabric Mills	0.53
315292	Fur and Leather Clothing Manufacturing	0.42
315291	Infants' Cut and Sew Clothing Manufacturing	0.32
315210	Cut and Sew Clothing Contracting	0.30
337214	Office Furniture (except Wood) Manufacturing	0.21
332720	Turned Product and Screw, Nut and Bolt Manufacturing	0.21
313110	Fibre, Yarn and Thread Mills	0.19
333511	Industrial Mould Manufacturing	0.18
(B) 1999		
315231	Women's and Girls' Cut and Sew Lingerie, Loungewear and Nightwear Manufacturing	0.63
313240	Knit Fabric Mills	0.47
315210	Cut and Sew Clothing Contracting	0.22
333220	Rubber and Plastics Industry Machinery Manufacturing	0.20
336370	Motor Vehicle Metal Stamping	0.18
332720	Turned Product and Screw, Nut and Bolt Manufacturing	0.18
336330	Motor Vehicle Steering and Suspension Components (except Spring) Manufacturing	0.17
333519	Other Metalworking Machinery Manufacturing	0.16
337214	Office Furniture (except Wood) Manufacturing	0.15
315291	Infants' Cut and Sew Clothing Manufacturing	0.14
(C) 2009		
315231	Women's and Girls' Cut and Sew Lingerie, Loungewear and Nightwear Manufacturing	0.61
322299	All Other Converted Paper Product Manufacturing	0.29
337214	Office Furniture (except Wood) Manufacturing	0.17
336370	Motor Vehicle Metal Stamping	0.17
332720	Turned Product and Screw, Nut and Bolt Manufacturing	0.16
337215	Showcase, Partition, Shelving and Locker Manufacturing	0.15
321112	Shingle and Shake Mills	0.14
331420	Copper Rolling, Drawing, Extruding and Alloying	0.13
336360	Motor Vehicle Seating and Interior Trim Manufacturing	0.13
315110	Hosiery and Sock Mills	0.13

Notes: The CDF at distance d is the cumulative sum of the K -densities up to distance d . Results in this table are reported for a distance $d = 50$ kilometers. To understand how to read that table, take 'Women's and Girls' Cut and Sew Lingerie, Loungewear and Nightwear Manufacturing' (NAICS 315231) as an example. In 1990, 62 percent of the distances between plants in that industry are less than 50 kilometers. Put differently, if we draw two plants in that industry at random, the probability that these plants are less than 50 kilometers apart is 0.62. If we, however, draw two plants at random among *all* manufacturing plants, that same probability would only be about 0.08. Clearly, this large difference suggests that the location patterns of plants in the 'Women's and Girls' Cut and Sew Lingerie, Loungewear and Nightwear Manufacturing' industry are very different from those of manufacturing in general. Plants in that industry are much closer than they 'should be' if they were distributed like overall manufacturing.

Table 9: Percentages of industries with random, localized, and dispersed patterns, 1990–2009.

Year	(1) Unweighted			(2) Employment weighted			(3) Sales weighted		
	Random	Localized	Dispersed	Random	Localized	Dispersed	Random	Localized	Dispersed
1990	52.53	34.63	12.84	52.53	36.96	10.51	54.86	37.35	7.78
1991	51.36	36.19	12.45	52.92	38.52	8.56	55.25	36.19	8.56
1992	53.70	36.19	10.12	56.42	35.02	8.56	58.37	33.46	8.17
1993	53.70	34.24	12.06	58.37	33.46	8.17	59.53	31.52	8.95
1994	49.81	36.96	13.23	57.20	33.07	9.73	60.70	30.74	8.56
1995	55.25	33.46	11.28	58.37	33.07	8.56	59.53	32.30	8.17
1996	54.09	35.41	10.51	56.03	35.41	8.56	59.53	33.46	7.00
1997	55.25	35.41	9.34	60.70	32.30	7.00	61.09	32.68	6.23
1998	55.64	34.24	10.12	58.37	35.02	6.61	61.87	32.68	5.45
1999	55.25	34.63	10.12	58.75	35.41	5.84	61.48	32.30	6.23
2000	47.86	37.74	14.40	51.75	40.47	7.78	53.31	40.47	6.23
2001	43.58	41.25	15.18	52.92	40.86	6.23	50.58	42.41	7.00
2002	45.91	39.69	14.40	50.97	41.63	7.39	54.86	37.35	7.78
2003	47.47	36.58	15.95	50.58	40.86	8.56	55.64	35.41	8.95
2004	60.31	30.35	9.34	60.31	33.07	6.61	60.70	32.30	7.00
2005	58.75	33.46	7.78	62.65	31.13	6.23	64.20	31.52	4.28
2006	60.31	30.35	9.34	60.31	33.46	6.23	62.26	33.85	3.89
2007	57.59	33.46	8.95	60.70	33.85	5.45	62.65	32.30	5.06
2008	56.03	34.24	9.73	61.48	31.91	6.61	64.59	29.96	5.45
2009	59.53	33.07	7.39	63.04	31.52	5.45	63.04	31.13	5.84

Notes: Authors' computations using Statistics Canada's Annual Survey of Manufacturers Longitudinal Microdata file. The statistical significance of the location patterns is computed using Monte Carlo simulations with 1,000 replications following the procedure developed by [Duranton and Overman \(2005\)](#). See Section 3.2 for details.

absolute terms. The estimations in Section 5.3 deal with the issue of the relative decline of manufacturing.

Table 10: Changes in the overall composition of employment in Canada, 1992–2016.

Year	(1) Manufacturing sector		(2) Goods-producing sector		(3) Total, all industries
	Employment	Share of total	Employment	Share of total	Employment
1992	1,814.5	14.25%	3,390.6	26.6%	12,730.9
2009	1,745.1	10.43%	3,720.1	22.2%	16,727.6
2016	1,694.8	9.37%	3,833.0	21.2%	18,079.9
% change (1992 – 2016)	-7.06%		+13.05%		+42.02%

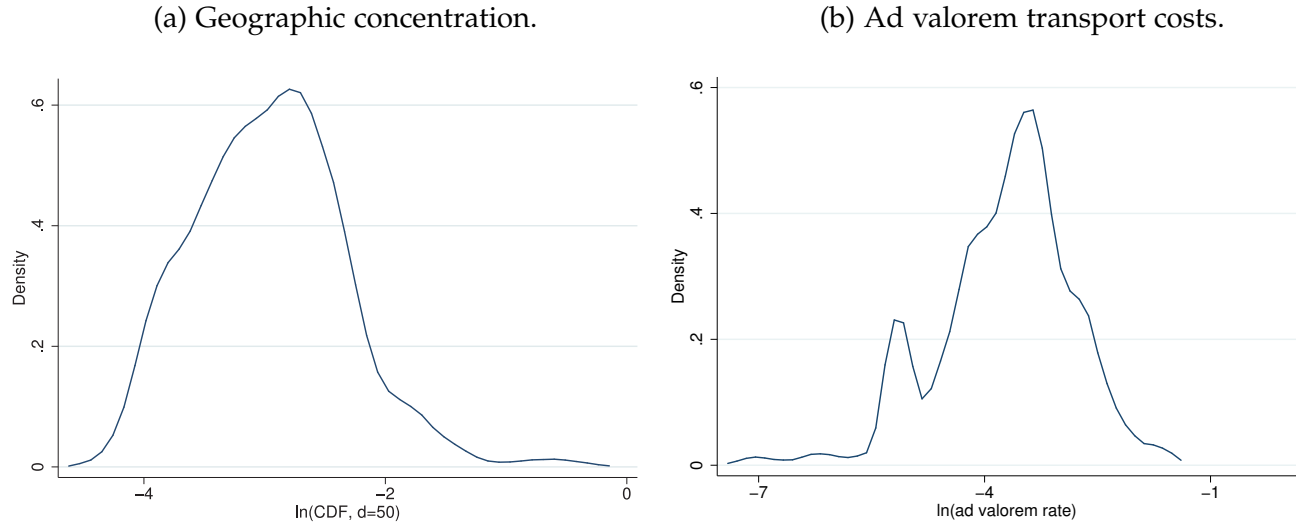
Notes: CANSIM Table 282-0008, 'Labour force survey estimates (LFS), by North American Industry Classification System (NAICS), sex and age group' for 'both sexes, 15 years and over' (figures given in 1,000s of workers). (1) 'Total, all industries' corresponds to NAICS codes 11 to 91; (2) 'Goods-producing sector' corresponds to NAICS codes 11 to 33; and (3) 'Manufacturing sector' corresponds to NAICS codes 31 to 33.

D.3. Log-transformation of our key variables

Figure 8 depicts the kernel densities of our log-transformed measures of geographic concentration (left panel) and ad valorem rates (right panel), pooled across all years (individual cross sections do not look very different). As shown, the log-transformed

distributions look fairly normal, which stems from the fact that the untransformed distributions are strongly right-skewed: there are a few strongly concentrated and a few high transport cost industries. This feature of our data suggests that a log transformation is appropriate.

Figure 8: Log-transformed key variables.



D.4. Ad valorem transport costs

Table 11 shows summary statistics for the ad valorem transport costs of the top-10 and the bottom-10 6-digit NAICS industries. As shown, the ten most expensive industries to ship have on average an ad valorem transport cost of about 12–14%, whereas the ten least expensive industries to ship have on average an ad valorem transport cost of about 0.3–0.4%. Clearly, there are enormous differences in the relative importance of shipping costs across industries. As shown by Figure 1, there is also substantial change across time as the prices of industry output and the general cost of the transport sector change.

Table 12 provides detailed results of ad valorem trucking costs for 2008 at the 4-digit level. For confidentiality reasons, we cannot report the 6-digit estimates that we use in this paper, but the 4-digit and the 6-digit estimates are very much in line with each other. As Table 12 shows, there are a number of industries with high ad valorem transport costs: ‘Cement and Concrete Product Mfg’, with 9.9%; ‘Lime and Gypsum Product Mfg’, with 15.6%; and ‘Other Non-Metallic Mineral Product Mfg’, with 9.3%. These are two orders of magnitude larger than the ad valorem transport costs for ‘Communications Equipment Mfg’ with 0.1% or ‘Computer and Peripheral Equipment Mfg’ with 0.3%. The simple average across 4-digit industries is 2.9%, with standard deviation of 2.39.

Table 11: Highest and lowest ad valorem trucking costs.

Year	Average top-10	Average bottom-10
1992	14.3%	0.40%
1993	13.8%	0.39%
1994	12.8%	0.37%
1995	12.5%	0.37%
1996	12.3%	0.36%
1997	12.7%	0.36%
1998	12.6%	0.35%
1999	12.2%	0.34%
2000	12.6%	0.36%
2001	12.8%	0.36%
2002	12.5%	0.34%
2003	12.9%	0.35%
2004	13.2%	0.37%
2005	13.2%	0.37%
2006	13.5%	0.38%
2007	13.3%	0.37%
2008	14.0%	0.39%

Notes: Statistics Canada, author's calculations, based on NAICS 6-digit industries.

Appendix E. Additional estimation results and robustness checks

E.1. Cross section in 2008

Table 13 provides detailed OLS and IV results for the cross section of industries in 2008. As can be seen from that table, all of our results reported using the pooled cross section are robust in the simple cross section.

E.2. Results for employment-weighted, sales-weighted, and five-year averages panel

Table 14 provides evidence for the robustness of our results to the choice of the dependent variable, to the distance cutoff used for the CDF, and to the year-on-year volatility of some variables included in the model **Model (P6)**. It shows that the effect of transport costs on geographic concentration is weaker—and the explanatory power of the model slightly lower—when the latter is measured using either employment- or sales-weighted CDFs. Although the key qualitative flavor of the results and the sign and significance of our key coefficient remain unchanged, the estimates using employment- or sales-weighted K -densities are slightly less sharp. Furthermore, the effect of import competition tends to be more limited to imports from Asia, and the coefficient tends to be smaller too. This suggests that much of the adaptation to import competition, partic-

Table 12: Detailed industry-level ad valorem trucking costs for 2008.

NAICS	Industry name	% AVTC
3111	Animal Food Mfg	4.8
3112	Grain and Oilseed Milling	3.3
3113	Sugar and Confectionery Product Mfg	2.7
3114	Fruit and Vegetable Preserving and Specialty Food Mfg	6.2
3115	Dairy Product Mfg	2.1
3116	Meat Product Mfg	3.8
3117	Seafood Product Preparation and Packaging	2.3
3118	Bakeries and Tortilla Mfg	3.0
3119	Other Food Mfg	2.9
3121	Beverage Mfg	4.4
3122	Tobacco Mfg	0.6
3131	Fibre, Yarn and Thread Mills	4.1
3132	Fabric Mills	1.5
3133	Textile and Fabric Finishing and Fabric Coating	5.4
3141	Textile Furnishings Mills	3.8
3149	Other Textile Product Mills	3.2
3151	Clothing Knitting Mills	0.6
3152	Cut and Sew Clothing Mfg	0.6
3159	Clothing Accessories and Other Clothing Mfg	1.0
3161	Leather and Hide Tanning and Finishing	2.3
3162	Footwear Mfg	1.9
3169	Other Leather and Allied Product Mfg	0.9
3211	Sawmills and Wood Preservation	8.9
3212	Veneer, Plywood and Engineered Wood Product Mfg	5.0
3219	Other Wood Product Mfg	4.8
3221	Pulp, Paper and Paperboard Mills	6.3
3222	Converted Paper Product Mfg	5.8
3231	Printing and Related Support Activities	4.7
3241	Petroleum and Coal Products Mfg	2.6
3251	Basic Chemical Mfg	3.4
3252	Resin, Synthetic Rubber, and Artificial and Synthetic Fibres and Filaments Mfg	4.0
3253	Pesticide, Fertilizer and Other Agricultural Chemical Mfg	4.1
3254	Pharmaceutical and Medicine Mfg	1.7
3255	Paint, Coating and Adhesive Mfg	2.8
3256	Soap, Cleaning Compound and Toilet Preparation Mfg	4.2
3259	Other Chemical Product Mfg	2.2
3261	Plastic Product Mfg	2.4
3262	Rubber Product Mfg	2.9
3271	Clay Product and Refractory Mfg	2.7
3272	Glass and Glass Product Mfg	5.4
3273	Cement and Concrete Product Mfg	9.9
3274	Lime and Gypsum Product Mfg	15.6
3279	Other Non-Metallic Mineral Product Mfg	9.3
3311	Iron and Steel Mills and Ferro-Alloy Mfg	2.2
3312	Steel Product Mfg from Purchased Steel	4.0
3313	Alumina and Aluminum Production and Processing	2.4
3314	Non-Ferrous Metal (except Aluminum) Production and Processing	0.9
3315	Foundries	1.8
3321	Forging and Stamping	1.9
3322	Cutlery and Hand Tool Mfg	1.4
3323	Architectural and Structural Metals Mfg	3.8
3324	Boiler, Tank and Shipping Container Mfg	2.1
3325	Hardware Mfg	2.9

Table 12 (continued)

3326	Spring and Wire Product Mfg	3.6
3327	Machine Shops, Turned Product, and Screw, Nut and Bolt Mfg	3.2
3329	Other Fabricated Metal Product Mfg	1.5
3331	Agricultural, Construction and Mining Machinery Mfg	1.2
3332	Industrial Machinery Mfg	1.5
3333	Commercial and Service Industry Machinery Mfg	0.6
3334	Ventilation, Heating, Air-Conditioning and Commercial Refrigeration Equipment Mfg	1.5
3335	Metalworking Machinery Mfg	1.0
3336	Engine, Turbine and Power Transmission Equipment Mfg	1.5
3339	Other General-Purpose Machinery Mfg	1.6
3341	Computer and Peripheral Equipment Mfg	0.3
3342	Communications Equipment Mfg	0.1
3343	Audio and Video Equipment Mfg	1.2
3344	Semiconductor and Other Electronic Component Mfg	0.7
3345	Navigational, Measuring, Medical and Control Instruments Mfg	0.4
3346	Mfg and Reproducing Magnetic and Optical Media	1.1
3351	Electric Lighting Equipment Mfg	3.5
3352	Household Appliance Mfg	3.3
3353	Electrical Equipment Mfg	1.7
3359	Other Electrical Equipment and Component Mfg	1.7
3361	Motor Vehicle Mfg	0.6
3362	Motor Vehicle Body and Trailer Mfg	1.1
3363	Motor Vehicle Parts Mfg	1.9
3364	Aerospace Product and Parts Mfg	0.7
3365	Railroad Rolling Stock Mfg	2.2
3366	Ship and Boat Building	2.7
3369	Other Transportation Equipment Mfg	1.7
3371	Household and Institutional Furniture and Kitchen Cabinet Mfg	1.9
3372	Office Furniture (including Fixtures) Mfg	4.1
3379	Other Furniture-Related Product Mfg	1.8
3391	Medical Equipment and Supplies Mfg	0.5
3399	Other Miscellaneous Mfg	2.7

Notes: Statistics Canada, author's calculations.

ularly from low-wage countries which are responsible for the bulk of exit in Canadian manufacturing, occurs for smaller plants and firms. The residual transport cost variable remains significantly negative in all specifications that we estimate, irrespective of how we construct the dependent variable and irrespective of the distance at which we evaluate geographic concentration. Note that this differs from the results using count-based measures, where the coefficients become insignificant after about 150 kilometer in the panel specification. In a nutshell, changes in transport costs have a significant effect on the geographic concentration of economic activity, no matter whether we consider plants, employment, or sales to measure that concentration. Last, we also re-estimate the model by averaging all variables over five year periods. Doing so reduces the year-on-year volatility of some variables (e.g., the trade variables), and allows for slowly moving variables like R&D expenditures to be potentially better identified in the regressions. It also deals with business cycle aspects that may drive the changes in the geographic concentration of industries. The last three columns of Table 14 show that our basic qualitative

Table 13: Estimation results for the cross section in 2008, OLS and IV.

Variables	(X1')	(X2')	(X3')	(X4')	(X5')	(X6')	(X7')	(X8')	(X9')
	Base	Trade	IO-links	Trade&IO	Controls	Full	Purged	iv-2SLS Q5	iv-2SLS Q3
AVTC	-0.219 ^a (0.041)	-0.244 ^a (0.048)	-0.041 (0.028)	-0.086 ^a (0.032)	-0.128 ^b (0.052)	-0.074 ^b (0.036)			
AVTC residual							-0.074 ^b (0.036)	-0.119 ^a (0.041)	-0.103 ^b (0.050)
Asian share of imports		0.183 (0.451)		-0.372 (0.298)		-0.545 ^c (0.316)	-0.547 ^c (0.317)	-0.595 ^c (0.306)	-0.577 ^c (0.313)
OECD share of imports		0.359 (0.498)		0.007 (0.320)		-0.206 (0.300)	-0.207 (0.263)	-0.220 (0.289)	-0.215 (0.291)
NAFTA share of imports		-0.161 ^a (0.407)		-0.022 (0.273)		-0.243 (0.262)	-0.245 (0.263)	-0.240 (0.251)	-0.242 (0.251)
Asian share of exports		-1.016 ^b (0.497)		-0.238 (0.401)		0.232 (0.343)	0.234 (0.344)	0.202 (0.327)	0.214 (0.326)
OECD share of exports		-0.863 ^c (0.485)		-0.116 (0.403)		0.438 (0.369)	0.440 (0.370)	0.424 (0.354)	0.430 (0.354)
NAFTA share of exports		0.224 (0.381)		0.141 (0.274)		0.478 ^c (0.262)	0.480 ^c (0.262)	0.498 ^c (0.252)	0.491 ^c (0.254)
Input distance			-0.184 ^c (0.103)	-0.343 ^a (0.102)		-0.247 ^b (0.112)	-0.247 ^b (0.112)	-0.235 ^b (0.107)	-0.239 ^b (0.108)
Output distance			-0.511 ^a (0.105)	-0.290 ^a (0.100)		-0.501 ^a (0.111)	-0.501 ^a (0.111)	-0.515 ^a (0.106)	-0.510 ^a (0.107)
Average minimum distance			-0.266 ^a (0.064)	-0.300 ^a (0.066)		-0.317 ^a (0.064)	-0.318 ^a (0.064)	-0.305 ^a (0.062)	-0.309 ^a (0.063)
Industry controls included	No	No	No	No	Yes	Yes	Yes	Yes	Yes
Observations (number of NAICS ind.)	257	257	257	257	257	257	257	257	257
R^2	0.102	0.166	0.613	0.632	0.198	0.702	0.702	—	—
First-stage R^2								1008.95	550.41
First-stage F test of excluded instruments								0.883	0.826

Notes: ^a, ^b and ^c denote coefficients significant at the 1%, 5%, and 10% levels, respectively. Standard errors in parentheses are clustered at the industry level. The dependent variable is the count-based Duranton-Overman K -density CDF at 50 kilometers distance. We have 257 industries. Our measures of input and output distances, as well as average minimum distance, are computed using $N = 5$ (see Appendices A.1 and A.2 for details). A constant term is included but not reported. (X5')–(X9') include the following industry controls: Total industry employment; Firm Herfindahl index (employment based); Mean plant size; Share of plants affiliated with multiplant firms; Share of plants controlled by foreign firms; Natural resource share of inputs; Energy share of inputs; Share of hours worked by all workers with post-secondary education; In-house R&D share of sales. 'AVTC residual' denotes the residual of the regression of ad valorem trucking costs on industry multi-factor productivity. (X5') only includes our industry-level controls. (X8') and (X9') use the rank-bin of the cross-sectional rates as an instrument, with quintile bins (Q5) in the former and tertile bins (Q3) in the latter.

findings are unchanged when replacing year-on-year variations with five-year averages. Yet, the results are a bit sensitive to the volatility of some of the variables when using year-on-year changes, as shown by the substantial decrease in the coefficient on transport costs.

E.3. Excluding specific industries

Our results may be partly driven by a small number of sectors that were subject to major changes over our study period. For example, the Canadian textile and clothing industry experienced a remarkable downward trend in the number of plants and in its geographic concentration in the wake of the end of the Multi-Fibre Arrangement in 2005 (see [Behrens, Boualam, and Martin, 2017](#)). Given that the textile and clothing industry contains some of the initially most strongly agglomerated sectors in Canada (see Table 8 for details), the large changes in those sectors may drive some of our key results. That this is not the case, and that all of our main findings in the panel regressions are robust to the exclusion of those sectors, is shown in Table 15. We also run our panel regressions by excluding the ‘high-tech’ sectors, and the results are qualitatively unchanged.

E.4. Incremental distance results

We run a variety of regressions to compute the incremental change in the ad valorem transport cost coefficients for various distance bands. The left half of Table 16 summarizes our results for the cross section, whereas the right half reports the same results for the panel. To save space, we only report results for the full specifications with the residual transport cost measure (**Models (X6) and (P6)**). We define the incremental distance of the CDF between distances $d_1 < d_2$ as $\Delta\gamma_i(d_1, d_2) = \text{CDF}_i(d_2) - \text{CDF}_i(d_1)$, with all variables taken in logs. We estimate the marginal effects by ‘distance bands’. The results in Table 16 complement those depicted by Figure 2 in the text.

E.5. Dependence on the transport mode

One may be worried that our results pick up the peculiar geography of Canada, where seaborne trade occurs in a highly localized way through the major ports at the two coasts, whereas land borne trade with the U.S. (and Mexico) occurs in a much more diffuse way along the long southern border. This may, in turn, affect the transport cost coefficient. To address this potential concern, we first include as a control the average minimum distance of plants in each industry from one of the five major Canadian container ports. As can be seen from **Models (X12) and (P12)** in Table 17, the ‘distance to major container

Table 14: Estimation of (P6) using employment-weighted CDFs, sales-weighted CDFs, and five year averages.

Variables	(1) Employment weighted			(2) Sales weighted			(3) Unweighted, five year averages		
	CDF 10km	CDF 100km	CDF 500km	CDF 10km	CDF 100km	CDF 500km	CDF 10km	CDF 100km	CDF 500km
AVTC residual	-0.158 ^b (0.077)	-0.150 ^b (0.072)	-0.148 ^a (0.053)	-0.134 ^c (0.076)	-0.127 ^c (0.070)	-0.137 ^a (0.045)	-0.377 ^a (0.085)	-0.361 ^a (0.076)	-0.315 ^a (0.060)
Asian share of imports	-0.684 ^b (0.312)	-0.531 ^b (0.252)	-0.241 ^c (0.145)	-0.713 ^b (0.349)	-0.604 ^b (0.276)	-0.285 ^c (0.162)	-1.463 ^b (0.579)	-1.012 ^a (0.357)	-0.383 ^c (0.202)
OECD share of imports	-0.377 (0.264)	-0.232 (0.217)	0.008 (0.164)	-0.305 (0.286)	-0.186 (0.236)	0.043 (0.176)	-0.770 (0.566)	-0.351 (0.336)	-0.006 (0.236)
NAFTA share of imports	-0.312 (0.244)	-0.208 (0.198)	-0.018 (0.141)	-0.262 (0.276)	-0.195 (0.226)	0.003 (0.159)	-0.821 (0.518)	-0.477 (0.317)	-0.104 (0.201)
Asian share of exports	0.264 (0.483)	0.368 (0.389)	0.065 (0.130)	0.217 (0.507)	0.299 (0.398)	0.082 (0.106)	0.322 (0.539)	0.366 (0.439)	0.051 (0.211)
OECD share of exports	0.212 (0.295)	0.330 (0.210)	0.181 ^c (0.094)	0.349 (0.288)	0.424 ^c (0.216)	0.280 ^a (0.096)	0.360 (0.386)	0.450 (0.314)	0.266 (0.191)
NAFTA share of exports	0.111 (0.310)	0.276 (0.206)	0.098 (0.075)	0.190 (0.303)	0.318 (0.213)	0.169 ^b (0.076)	0.265 (0.383)	0.442 (0.296)	0.180 (0.149)
Input distance	-0.256 ^a (0.063)	-0.238 ^a (0.054)	-0.186 ^a (0.032)	-0.256 ^a (0.064)	-0.239 ^a (0.056)	-0.180 ^a (0.033)	-0.258 ^a (0.073)	-0.246 ^a (0.059)	-0.221 ^a (0.043)
Output distance	-0.234 ^a (0.053)	-0.222 ^a (0.048)	-0.127 ^a (0.030)	-0.200 ^a (0.056)	-0.193 ^a (0.048)	-0.113 ^a (0.029)	-0.374 ^a (0.069)	-0.383 ^a (0.062)	-0.239 ^a (0.044)
Minimum distance	-0.312 ^a (0.050)	-0.246 ^a (0.039)	-0.119 ^a (0.026)	-0.327 ^a (0.054)	-0.249 ^a (0.039)	-0.131 ^a (0.026)	-0.400 ^a (0.067)	-0.297 ^a (0.043)	-0.141 ^a (0.032)
Number of NAICS industries	257	257	257	257	257	257	257	257	257
Number of years	17	17	17	17	17	17	4	4	4
Observations (NAICS × years)	4,369	4,369	4,369	4,369	4,369	4,369	1,028	1,028	1,028
R ²	0.318	0.371	0.381	0.294	0.359	0.376	0.517	0.599	0.598

Notes: ^a, ^b and ^c denote coefficients significant at the 1%, 5%, and 10% levels, respectively. Standard errors in parentheses are clustered at the industry level. Our measures of input and output distances, as well as average minimum distance, are computed using $N = 5$ (see Appendices A.1 and A.2 for details). A constant term is included in all regressions but not reported. All regressions include industry and year dummies and the following industry controls: Total industry employment; Firm Herfindahl index (employment based); Mean plant size; Share of plants affiliated with multiplant firms; Share of plants controlled by foreign firms; Natural resource share of inputs; Energy share of inputs; Share of hours worked by all workers with post-secondary education; In-house R&D share of sales. 'AVTC residual' denotes the residual of the regression of AVTC on industry multi-factor productivity.

Table 15: Estimation of (P6), excluding textile and high-tech industries.

Variables	(1) Excluding textiles industries			(2) Excluding high-tech industries		
	CDF 10km	CDF 100km	CDF 500km	CDF 10km	CDF 100km	CDF 500km
AVTC residual	-0.213 ^a (0.077)	-0.210 ^a (0.072)	-0.193 ^a (0.049)	-0.396 ^a (0.145)	-0.324 ^b (0.128)	-0.205 ^a (0.068)
Asian share of imports	-0.568 ^c (0.322)	-0.508 ^c (0.282)	-0.211 (0.174)	-1.517 ^a (0.554)	-1.035 ^a (0.350)	-0.380 ^b (0.155)
OECD share of imports	-0.035 (0.275)	0.007 (0.241)	0.137 (0.181)	-0.860 (0.530)	-0.474 (0.333)	-0.084 (0.177)
NAFTA share of imports	-0.097 (0.251)	-0.062 (0.221)	0.076 (0.156)	-0.878 ^c (0.499)	-0.531 ^c (0.317)	-0.133 (0.157)
Asian share of exports	0.627 (0.440)	0.505 (0.358)	0.096 (0.130)	0.468 (0.490)	0.469 (0.378)	0.111 (0.121)
OECD share of exports	0.471 ^b (0.186)	0.413 ^b (0.161)	0.249 ^b (0.097)	0.346 (0.236)	0.424 ^b (0.170)	0.271 ^a (0.098)
NAFTA share of exports	0.400 ^b (0.196)	0.348 ^b (0.170)	0.128 (0.080)	0.149 (0.246)	0.275 (0.179)	0.124 (0.085)
Input distance	-0.458 ^a (0.051)	-0.439 ^a (0.049)	-0.315 ^a (0.036)	-0.387 ^a (0.075)	-0.346 ^a (0.057)	-0.245 ^a (0.038)
Output distance	-0.265 ^a (0.043)	-0.245 ^a (0.040)	-0.155 ^a (0.029)	-0.333 ^a (0.051)	-0.336 ^a (0.044)	-0.216 ^a (0.030)
Average minimum distance	-0.289 ^a (0.041)	-0.265 ^a (0.038)	-0.142 ^a (0.026)	-0.321 ^a (0.053)	-0.257 ^a (0.038)	-0.128 ^a (0.026)
Number of NAICS industries	229	229	229	198	198	198
Number of years	17	17	17	17	17	17
Observations (NAICS × years)	3,893	3,893	3,893	3,366	3,366	3,366
R ²	0.516	0.532	0.539	0.481	0.556	0.553

Notes: ^a, ^b and ^c denote coefficients significant at the 1%, 5%, and 10% levels, respectively. Standard errors in parentheses are clustered at the industry level. Our measures of input and output distances, as well as average minimum distance, are computed using $N = 5$ (see Appendices A.1 and A.2 for details). A constant term is included in all regressions but not reported. All regressions include industry and year dummies and the following industry controls: Total industry employment; Firm Herfindahl index (employment based); Mean plant size; Share of plants affiliated with multiplant firms; Share of plants controlled by foreign firms; Natural resource share of inputs; Energy share of inputs; Share of hours worked by all workers with post-secondary education; In-house R&D share of sales. 'AVTC residual' denotes the residual of the regression of AVTC on industry multi-factor productivity. Our definition of high-tech sectors is based on the US Bureau of Labor Statistics classification by Hecker (2005). This definition of high-tech industries is 'input based'. An industry is 'high-tech' if it employs a high proportion of scientists, engineers or technicians. As shown by Hecker (2005), these industries are also usually associated with a high R&D-to-sales ratio, and they also largely – but not always – produce goods that are classified as 'high-tech' by the Bureau of Economic Analysis.

ports' variable is insignificant in the cross section, and significantly negative in the panel. This suggests that industries that do move closer to import entry points tend to concentrate more geographically, because import entry points are few and themselves located in a few specific places. As an additional and arguably cleaner test, we also estimate heterogeneous effects for transport costs based on how much of an industry's trade is likely to be seaborne versus land borne. To this end, we compute the relative share of the industry's exports to (or imports from) NAFTA compared to those to Asia. Because

Table 16: Effects of transport costs on geographic concentration at different spatial scales.

Variables	(1) Cross section (X6) by incremental change in the CDF				(2) Panel (P6) by incremental change in the CDF			
	$\Delta\gamma_i(10, 25)$	$\Delta\gamma_i(25, 50)$	$\Delta\gamma_i(50, 100)$	$\Delta\gamma_i(100, 500)$	$\Delta\gamma_i(10, 25)$	$\Delta\gamma_i(25, 50)$	$\Delta\gamma_i(50, 100)$	$\Delta\gamma_i(100, 500)$
AVTC residual	-0.072 ^b (0.034)	-0.058 ^c (0.033)	-0.046 (0.031)	-0.022 (0.019)	-0.254 ^a (0.078)	-0.239 ^a (0.079)	-0.230 ^a (0.069)	-0.106 (0.089)
Asian share of imports	-0.296 (0.289)	-0.233 (0.276)	-0.122 (0.229)	0.265 ^c (0.152)	-1.029 ^b (0.434)	-0.724 ^b (0.338)	-0.352 (0.235)	0.583 (0.429)
OECD share of imports	-0.235 (0.281)	-0.157 (0.272)	0.007 (0.228)	0.417 ^b (0.165)	-0.450 (0.374)	-0.174 (0.285)	0.102 (0.211)	0.721 (0.455)
NAFTA share of imports	-0.184 (0.244)	-0.125 (0.235)	0.013 (0.196)	0.347 ^a (0.132)	-0.526 (0.359)	-0.284 (0.268)	0.007 (0.190)	0.587 (0.372)
Asian share of exports	0.592 (0.414)	0.502 (0.397)	0.250 (0.314)	-0.649 (0.397)	0.630 (0.427)	0.658 (0.405)	0.421 (0.264)	-0.782 (0.714)
OECD share of exports	0.380 (0.305)	0.324 (0.297)	0.137 (0.267)	-0.307 ^c (0.185)	0.545 ^a (0.198)	0.662 ^a (0.224)	0.470 ^a (0.156)	-0.112 (0.304)
NAFTA share of exports	0.347 (0.250)	0.331 (0.234)	0.197 (0.191)	-0.259 ^c (0.139)	0.440 ^b (0.211)	0.541 ^b (0.215)	0.431 ^a (0.162)	-0.191 (0.274)
Input distance	-0.302 ^a (0.073)	-0.258 ^a (0.071)	-0.231 ^a (0.071)	-0.100 ^c (0.056)	-0.332 ^a (0.061)	-0.322 ^a (0.056)	-0.315 ^a (0.054)	-0.193 ^a (0.041)
Output distance	-0.407 ^a (0.070)	-0.415 ^a (0.068)	-0.378 ^a (0.065)	-0.198 ^a (0.052)	-0.341 ^a (0.046)	-0.340 ^a (0.045)	-0.302 ^a (0.045)	-0.123 ^a (0.039)
Minimum distance	-0.361 ^a (0.047)	-0.324 ^a (0.046)	-0.267 ^a (0.048)	-0.087 ^b (0.036)	-0.298 ^a (0.041)	-0.243 ^a (0.043)	-0.204 ^a (0.038)	-0.038 (0.036)
Industry dummies	No	No	No	No	Yes	Yes	Yes	Yes
Observations	4,369	4,369	4,369	4,369	4,369	4,369	4,369	4,369
R-squared	0.712	0.684	0.665	0.417	0.481	0.417	0.436	0.168

Notes: ^a, ^b and ^c denote coefficients significant at the 1%, 5%, and 10% levels, respectively. The dependent variable is the change $\Delta\gamma_i(d_1, d_2) = \gamma_i(d_1) - \gamma_i(d_2)$ in the unweighted (count based) Duranton-Overman K -density CDF between distance d_1 and d_2 . We have 17 years and 257 industries. Standard errors in parentheses are clustered at the industry level. Our measures of input and output distances, as well as average minimum distance, are computed using $N = 5$ (see Appendices A.1 and A.2 for details). A constant term is included in all regressions but not reported. All regressions include industry and year dummies and the following industry controls: Total industry employment; Firm Herfindahl index (employment based); Mean plant size; Share of plants affiliated with multiplant firms; Share of plants controlled by foreign firms; Natural resource share of inputs; Energy share of inputs; Share of hours worked by all workers with post-secondary education; In-house R&D share of sales. 'AVTC residual' denotes the residual of the regression of AVTC on industry multi-factor productivity.

trade with Asia is mostly seaborne, whereas trade with NAFTA is mostly land borne, a higher value for that variable implies that a larger share of that industry's international trade goes by truck. We create a dummy variable for the top (Q1) and the bottom (Q4) quartile of industries according to that measure, and interact it with our transport costs variable. We also experimented with including directly all trade shares with the interaction terms, but the resulting variables are too colinear to allow for meaningful estimation of the effects.

Table 17: Port distance and heterogeneous effects by prevalence of mode.

Variables	(X12) Port distance	(P12)	(X13) Relative NAFTA exports	(P13)	(X14) Relative NAFTA imports	(P14)
AVTC residual	-0.068 ^b (0.031)	-0.245 ^a (0.082)	-0.092 ^b (0.039)	-0.307 ^a (0.104)	-0.068 ^b (0.032)	-0.284 ^a (0.108)
Distance to major container ports	-0.160 (0.125)	-0.235 ^b (0.098)				
Relative NAFTA exports (Q1)			-0.069 (0.042)	-0.052 ^b (0.020)		
AVTC × relative NAFTA exports (Q1)			-0.095 ^c (0.052)	0.007 (0.022)		
Relative NAFTA imports (Q1)					-0.029 (0.041)	-0.030 ^c (0.017)
AVTC × relative NAFTA imports (Q1)					-0.112 ^b (0.047)	-0.022 (0.156)
Trade shares included	Yes	Yes	No	No	No	No
Industry dummies	No	Yes	No	Yes	No	Yes
Observations (NAICS × years)	4,369	4,369	4,369	4,369	4,369	4,369
R ²	0.728	0.525	0.728	0.499	0.728	0.498

Notes: ^a, ^b and ^c denote coefficients significant at the 1%, 5% and 10% levels, respectively. Standard errors are clustered at the industry level and given in parentheses. The dependent variable is the count-based Duranton-Overman *K*-density CDF at 50 kilometers distance. We have 17 years and 257 industries. Our measures of input and output distances, as well as average minimum distance, are computed using $N = 5$ (see Appendices A.1 and A.2 for details). A constant term is included in all regressions but not reported. All models include year dummies, the input-output distances, the minimum distance, and the following industry controls: Total industry employment; Firm Herfindahl index (employment based); Mean plant size; Share of plants affiliated with multiplant firms; Share of plants controlled by foreign firms; Natural resource share of inputs; Energy share of inputs; Share of hours worked by all workers with post-secondary education; In-house R&D share of sales. **Models (X13), (P13), (X14), and (P14)** include 'Relative NAFTA imports (Q4)' and 'AVTC × relative NAFTA imports (Q4)', but they turn out insignificant.

Models (X13) and (X14) in Table 17 show that there are cross-sectional differences for high- and low-seaborne-trade industries in terms of how trucking costs affect geographic concentration. The interaction term for the top quartile is negative and significant for both exports and imports, thus showing that high transport cost industries that trade a lot with Asia are slightly more geographically concentrated than high transport cost industries that trade a lot with NAFTA, all else equal. This effect may be due to the difference in the geography of access to markets (sea vs land). The terms involving the bottom quartile are never significant. A formal test rejects the null hypothesis of

equality of the interaction terms Q_1 and Q_4 (not reported in the table) for both exports and imports in the cross section at the 5% level.

As shown by (P13) and (P14), the same effect shows up in the panel in terms of levels (the coefficient on the dummy variables being negative and significant for the top quartile industries). The important point to note in all regressions is that our estimates for transport costs remain very stable.

E.6. Other robustness checks

We ran a large number of additional robustness checks in our panel regressions, which we do not reported in detail here. However, we provide a brief summary and description of those checks.

Information and communication technologies. We investigate whether changes in information and communication technologies (ICT) may lead to more geographic dispersion. To this end, we use the ICT investment variables from the KLEMS database, interacted with the other variables of the model, to check whether changes in communication costs have the same effect than changes in transport costs. We did not get any significant coefficients—neither for the direct effects, nor for the interaction terms.

Heterogeneous transport cost coefficients. We deal with one aspect of this in Section 5.3, where we estimate different coefficients by mode. We also estimated additional models with heterogeneous coefficients since transport costs differ across industries. To this end, we split our sample into high versus low transport cost industries, using a ‘below median’–‘above median’ criterion to define high and low. The two coefficients were statistically identical. We also treated decreasing/increasing transport costs in an asymmetric way as they may have asymmetric impacts. Again, the two coefficients were fairly close.

Input-output links. We replaced our measures of input and output linkages with the industry ‘material share to sales’ ratio, a proxy for reliance on intermediate inputs. That variable turns out to be insignificant in our regressions, whereas the other coefficients are largely unaffected. We also compute measures of ‘upstreamness’ of industries following [Antràs, Chor, Fally, and Hillberry \(2012\)](#). Using those measures, we split industries into the top quintile Q_5 (most upstream industries) and the bottom quintile Q_1 (most downstream industries). We then reestimate the model by interacting the transport and trade variables with those upstream-downstream dummies to capture potentially different impacts on different industries in the vertical production chain. When including our

input-output measures, splitting by upstreamness has virtually no effects on our main coefficients, which suggests that our input-output measures capture quite well vertical industry links. When excluding those measures, we find that more downstream industries are more sensitive to both transport costs and import competition, although the differential effects are quite imprecisely estimated.

Labor intensity. We also run regressions where we control for the ‘labor intensity’ of an industry (not just high-skilled workers vs low-skilled workers). We constructed different measures using the quantity index of labor and the quantity index of capital from the KLEMS data, but these variables turned out again to be insignificant in our regressions.

Non-linear transport costs. We further experimented with different non-linear transport cost specifications. More precisely, we estimated the effect of transportation costs with a spline, allowing the coefficients to vary between ad valorem rates of 0 to 0.05% (low), 0.05 to 15% (moderate), and 15% or greater (high). These are admittedly arbitrary categories, but ones that we believe make intuitive sense. The results are, by and large, consistent with the simpler specification that we use. Yet, we find that at low levels, the effect of transportation costs is positive or insignificant. At moderate levels, the coefficient is negative and always significant, and at high levels the coefficient is negative and insignificant. Transport costs thus seem to matter most strongly in the intermediate range.

Standardized non-log-transformed variables. Finally, we also ran the model using non-log-transformed variables that we standardized to obtain beta coefficients. We did this to check our key panel results when the 2008 cross-section estimates of transport costs are not soaked up by the industry fixed effects. Our main results are robust to this, however as expected, the estimates are less precise because of the right-skew in the distribution of transport costs and geographic concentration measures (see Appendix C.3 and Figure 8). It is worth noting that the instrumentation works equally well in the standardized variables case than in the log-transformed case. In a nutshell, although the estimates are less precise, they remain qualitatively unchanged, especially in the iv regressions where the coefficients remain highly significant.

Notes

²¹There are approximately 34,000 trucking firms in Canada, as measured by the mean over 2001 to 2009 from Statistics Canada’s Business Register. Let s_m be the revenue share of firm m . The entropy

$E = \sum_m s_m \log(1/s_m)$ that measures the skewness in the distribution of revenues across firms has mean $\bar{E} = 3.61$ over the period. Using the ‘numbers equivalent’—which can be interpreted as the number of firms that would be present if revenues were evenly spread across them—this is the equivalent to a market served by $10^{\bar{E}} = 4,070$ firms. This large number of competitors suggests that the assumption of a perfectly competitive industry is reasonable.

²²For simplicity, we assume that carriers must ship the same commodity between the two regions in both directions. We also abstract from ‘deadheading’ for cargo and from multistop pickups.

²³Alternatively, we could assume that consumers pay the transport costs, i.e., the manufacturer maximizes $(p_{rs}^i - m_r^i)Q_{rs}^i(p_{rs}^i + t_{rs})$ instead of $(p_{rs}^i - m_r^i - t_{rs})Q_{rs}^i(p_{rs}^i)$. With CES demands, this yields $p_{rs}^i = \frac{\sigma_i}{\sigma_i - 1} (m + t/\sigma_i)$ which is the free-on-board (FOB) price that increases with shipping distance (Martin, 2012). In both cases, the cost-insurance-freight (CIF) price is the same. In our case, consumers pay transport costs $T = [\sigma/(\sigma - 1)]t$, whereas in the other case they pay $T = t$. By definition of the ad valorem rate $\tau = p^{\text{cif}}/p^{\text{fob}}$, we have $\tau_{rs}^i = 1 + t_{rs}^i/m_r^i$ in our case and $\tau_{rs}^i = (m_r^i + t_{rs}^i)/(m_r^i + \sigma_i^{-1}t_{rs}^i)$ in the alternative case. As the qualitative behavior of τ_{rs}^i is the same in both cases, we stick to the simpler specification where the manufacturer pays the transport costs to the carrier.

²⁴The shifter A could differ between regions. This amounts to replacing $M \equiv M_r/M_s$ with $\tilde{M} \equiv (A_r M_r)/(A_s M_s)$ and it does not change our analysis if $\tilde{M} \geq 1$.

²⁵See Rosenthal and Strange (2004) and Combes and Gobillon (2015) for the empirics of agglomeration economies; and Duranton and Puga (2004) and Behrens and Robert-Nicoud (2015) for the theory.

²⁶One can also imagine that the carriers’ cost c depends on the volume of trade between regions via, e.g., density economies. See Mori and Nishikimi (2002) and Behrens, Gaigné, and Thisse (2009) for theoretical models.

²⁷Estimates using a Herfindahl index of plant-level concentration are qualitatively similar.

²⁸We also constructed proxies for labor market conditions using the non-production to production workers ratio and other educational characteristics of the workforce. The latter are available at a more aggregated industry level (L -level) from Statistics Canada’s KLEMS database (e.g., the share of hours worked by all workers with a university degree, and the labor productivity index). These measures, however, turned out to be statistically insignificant.

²⁹An alternative would be to use shift-share instruments as in Autor, Dorn, and Hanson (2014), but since trade shares are controls and not our main variable of interest we do not further pursue this direction.

³⁰Because of confidentiality reasons, we do not use the finer W -level matrices since this would make disclosure of results more problematic. However, the tests we ran using those matrices yield very similar results to the ones we report in this paper. Using the L -level matrix also provides smoother series of input-output linkages than those obtained using the confidential W -level national input-output tables.

³¹We have no micro-geographic information on final demand and thus cannot include it in our output linkage measures. Using a population-weighted market potential measure as a proxy is infeasible because of the very strong persistence through time. However, our industry fixed effects are likely to control for slow-changing final demand due to changes in the population distribution.

³²Our dataset has the standard problem of reporting only a plant’s primary sector of activity. Hence, it is possible that a plant operates in multiple sectors, so that our measures still partly pick up own-industry location patterns. There is not much we can do about this. We experimented with measures where we exclude all plants within the same 4-digit industry, and the results do not change qualitatively.

³³In computing the K -densities, we follow Behrens and Bougna (2015) and choose a cutoff distance of 800 kilometers. This means that, although we use all plants in each industry to compute the K -densities, we do not evaluate them numerically beyond 800 kilometers.

³⁴Table 8 in the supplemental appendix provides a summary of the geographically most concentrated industries in selected years. Various textile, machinery, and automotive related industries top that list.

References

- [1] Antràs, Pol, Davin Chor, Thibault Fally, and Russell Hillberry. 2012. "Measuring the upstreamness of production and trade flows." *American Economic Review Papers and Proceedings* 102(3): 412–416.
- [2] Autor, David H., David Dorn, and Gordon H. Hanson. 2013. "The China syndrome: Local labor market effects of import competition in the United States." *American Economic Review* 103(6): 2121–2168.
- [3] Behrens, Kristian, Brahim Boualam, and Julien Martin. 2017. "Are clusters resilient? Evidence from Canadian textile industries." CEPR Discussion Paper #12184, Centre for Economic Policy Research, London, UK.
- [4] Behrens, Kristian, and Frédéric Robert-Nicoud. 2015. "Agglomeration theory with heterogeneous agents." In: Duranton, Gilles, J. Vernon Henderson, and William C. Strange (eds.) *Handbook of Regional and Urban Economics*, vol. 5. North-Holland: Elsevier, pp. 171–245.
- [5] Behrens, Kristian, Carl Gaigné, and Jacques-François Thisse. 2009. "Industry location and welfare when transport costs are endogenous." *Journal of Urban Economics* 65(2): 195–208.
- [6] Combes, Pierre-Philippe, and Laurent Gobillon. 2015. "The empirics of agglomeration economies." In: G. Duranton, J.V. Henderson, and W.C. Strange (eds.), *Handbook of Regional and Urban Economics*, vol. 5. North-Holland: Elsevier, pp. 247–348.
- [7] Duranton, Gilles, and Henry G. Overman. 2005. "Testing for localisation using micro-geographic data." *Review of Economic Studies* 72(4): 1077–1106.
- [8] Duranton, Gilles, and Diego Puga. 2004. "Micro-foundations of urban agglomeration economies." In: J. Vernon Henderson, and Jacques-François Thisse (eds.), *Handbook of Regional and Urban Economics*, vol. 4. North-Holland: Elsevier B.V., pp. 2063–2117.
- [9] Henderson, J. Vernon. 1997. "Medium-sized cities." *Regional Science and Urban Economics* 27(6): 583–612.

- [10] Holmes, Thomas J., and John J. Stevens. 2014. "An alternative theory of the plant size distribution, with geography and intra- and international trade." *Journal of Political Economy* 122(2): 369–421.
- [11] Lewbel, Arthur. 2012. "Using heteroscedasticity to identify and estimate mismeasured and endogenous regressor models." *Journal of Business and Economic Statistics* 30(1): 67–80.
- [12] Martin, Julien. 2012. "Markups, quality, and transport costs." *European Economic Review* 56(4): 777–791.
- [13] Mori, Tomoya, and Koji Nishikimi. 2002. "Economies of transport density and industrial agglomeration." *Regional Science and Urban Economics* 32(2): 167–200.
- [14] Redding, Stephen J., and Anthony J. Venables. 2004. "Economic geography and international inequality." *Journal of International Economics* 62(1), 63–82.
- [15] Rosenthal, Stuart S., and William C. Strange. 2001. "The determinants of agglomeration." *Journal of Urban Economics* 50(2): 191–229.