

A Service of



Leibniz-Informationszentrum Wirtschaft Leibniz Information Centre

Behrens, Kristian; Brown, W. Mark

Working Paper

Transport costs, trade, and geographic concentration: Evidence from Canada

Document de travail, No. 2017-09

Provided in Cooperation with:

Department of Economics, School of Management Sciences (ESG UQAM), University of Quebec in Montreal

Suggested Citation: Behrens, Kristian; Brown, W. Mark (2017): Transport costs, trade, and geographic concentration: Evidence from Canada, Document de travail, No. 2017-09, Université du Québec à Montréal, École des sciences de la gestion (ESG UQAM), Département des sciences économiques, Montréal

This Version is available at: https://hdl.handle.net/10419/234754

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.





DOCUMENT DE TRAVAIL / WORKING PAPER

No. 2017-09

Transport costs, trade, and geographic concentration: Evidence from Canada

Kristian Behrens et W. Mark Brown

Décembre 2017



Département des sciences économiques École des sciences de la gestion Université du Québec à Montréal This is a draft chapter / article. The final version is available in 'Handbook of International Trade and Transportation' edited by Bruce A. Blonigen and Wesley W. Wilson, published in 2018, Edward Elgar Publishing Ltd http://dx.doi.org/10.4337/9781785366154.

The material cannot be used for any other purpose without further permission of the publisher, and is for private use only.

Transport costs, trade, and geographic concentration: Evidence from Canada

 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

Document de travail No. 2017-09

Décembre 2017

Département des Sciences Économiques
Université du Québec à Montréal
Case postale 8888,
Succ. Centre-Ville
Montréal, (Québec), H3C 3P8, Canada
Courriel: brisson.lorraine@ugam.ca

Site web: http://economie.esg.uqam.ca

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

Copyright (2017): Kristian Behrens et W. Mark Brown. De courts extraits de texte peuvent être cités et reproduits sans permission explicite des auteurs à condition de référer au document de travail de manière appropriée.

Transport costs, trade, and geographic concentration: Evidence from Canada

Kristian Behrens* W. Mark Brown[†]

December 18, 2017

Abstract

Our objective is threefold. First, we explain how to estimate transport costs and the geographic concentration of industries using trucking microdata and geocoded plant-level data. Second, we document that transport costs explain between 25% to 57% of the observed relationship between trade and distance across Canada's economic regions. Last, we show that changes in transport costs have a substantial impact on geographic concentration patterns for vertically linked industries, depending on the strength of the links. A one standard deviation increase in transport costs leads to a 0.02 standard deviation decrease in geographic concentration for industry pairs at the bottom decile of the input-output coefficient distribution, whereas the corresponding effect at the top decile is a 0.02 standard deviation increase. This gap between weakly and strongly linked industries stands up to a wide range of specifications and is robust to instrumental variables estimations.

Keywords: Transport costs; trade; geographic concentration; Canada.

JEL Classification: R12; C23; L60.

^{*}Université du Québec à Montréal, Canada; National Research University Higher School of Economics, Russian Federation; 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

"He had been taught, of course, that history, along with geography, was dead." (William Gibson. 1999. All Tomorrow's Parties. Ace Books: New York, p. 165)

1 Introduction

As transport costs for goods have substantially fallen through time (see, e.g., Glaeser and Kohlhase, 2004), the point of view has grown that they "don't matter anymore" (see, e.g., Cairncross, 1991; Friedman, 2005). In particular, so goes the argument, vertical supply chains and firms' location choices are no longer strongly conditioned by the costs of shipping goods across establishments and regions. This view conflicts with a large body of academic research in geography and trade that has substantiated that transport costs still matter in shaping trade flows (e.g., Head and Mayer, 2013, 2014). Most of that literature, however, relies on either distance or infrastructure as a proxy for transport costs (see, e.g., Chandra and Thompson, 2000; Michaels, 2008; Duranton et al., 2014). While infrastructure, distance, and transport costs are linked, it is unclear what the strength of that link is and, moreover, how infrastructure or distance map into (changes in) transport costs. It is further unclear how changes in transport costs, ultimately, reshape the geographic concentration of industries. We believe it necessary, therefore, to measure transport costs directly in order to better understand how they affect trade flows and geographic patterns of economic activity. However, utilizing direct measures of transport costs to look at trade and geography is difficult because of a lack of data.

Most work to date that utilizes transport costs (e.g., Hummels, 2007) relies on the difference between cost-insurance-freight (CIF) and free-on-board (FOB) prices reported on international trade data to obtain estimates. While useful, because they are primarily based on water borne trade, these data are not necessarily appropriate to measure domestic transport costs, which are typically incurred overland. We take advantage of new data developed by Statistics Canada to measure domestic transport costs and trade. Built from transaction-level trucking records, these data provide a direct measure of transport costs and, because we are also able to estimate the value of shipments, allow us to measure transport costs on an ad valorem basis. When combined with measures of regional trade flows and information on the precise location of manufacturing plants, we can begin to answer some seemingly basic but understudied questions: (i) How can we measure transport costs using microdata?; (ii) how do transport costs affect trade?; (iii) how can we measure geographic concentration using microgeographic data?; and (iv) how do changes in transport costs affect the geographic concentration of industries and of vertically linked industry pairs? While addressed using data on domestic transport costs and trade, we believe the answers to these questions are also relevant to future research in international trade.

We focus on the effect of trucking-related transport costs on trade and the location of industries because of the importance of that mode. By value, trucking is the most important mode for moving goods in Canada and between Canada and the United States. For goods moved by truck and rail domestically, about 90% by value are moved by truck. Including additional modes (e.g., pipelines) lowers this share, but there is little doubt trucking is by far the most important domestic mode. Trucking is also the most important mode for Canada-U.S. trade, accounting for 72% of exports and 50% of imports in 2005. Trucking has the added advantage of being a highly competitive sector with about 34,000 firms in Canada that, after taking into account the distribution of revenues across firms, is the equivalent to a market served by 4,070 firms. This large number of competitors makes the simplifying assumption of a perfectly competitive sector tenable.

Previewing our key results, we find that transport costs are main drivers of both interregional trade flows and the geographic concentration of industries. Our estimates suggest that transport costs explain between 25% to 57% of the observed relationship between trade flows and distance in Canada, which is above previous estimates that gave transport costs a much smaller contribution that — at the most generous — ranged from 10% (Allen, 2014) to 28% (Head and Mayer, 2013). These figures, which are obtained from international trade data, are at or below our most conservative estimates, as expected. That is, because domestic trade relies more on expensive surface transportation modes, a stronger relationship between distance and transport costs results. A higher estimated elasticity of trade with respect to transport costs and a lower estimated elasticity of trade with respect to distance thus accounts for the stronger estimated contribution of transport costs to interregional trade than to international trade. Furthermore, our findings suggest transport costs have a material effect on the location of industries. Increases in transport costs generally tend to disperse manufacturing plants, be it within the same industry or within paired industries. As suggested by theory the effect is,

¹These figures are based on a tabulation of the Surface Transportation File from Statistics Canada for the period 2004 to 2012. The file includes shipments of truck and rail carriers and so excludes goods moved by pipeline, marine, and air. Hence, the overall modal share for trucking will be less. While transport costs are measured using the trucking mode, it is broadly representative of the transport costs incurred by manufactures when sourcing inputs within the sector. Across the set of commodities produced and used by manufacturing industries, trucking accounts for more than three-quarters of revenues to carriers across 2-digit Standard Classification of Transported Goods (SCTG) commodities; only 'Basic chemicals' and 'Pulp, newsprint, paper and paper board' are exceptions.

²Estimates are derived from the Bureau of Transportation Statistics' North American Transborder Freight Data database (available online at https://transborder.bts.gov/).

³The number of firms is the mean over 2001 to 2009 as reported by Statistics Canada's Business Register. Industry concentration is measured using the entropy-based numbers equivalent. Entropy is $E = \sum_m s_m \log(1/s_m)$, where s_m is the revenue share of firm m. Its mean \overline{E} over the period is 3.61 and the numbers equivalent is $10^{\overline{E}} = 4,070$. The latter can be interpreted as the number of firms that would be present if revenues were evenly spread across them.

however, mediated by input-output links. A one standard deviation increase in transport costs leads to a (statistically significant) 0.02 standard deviation decrease in geographic concentration for industry pairs at the bottom decile of the input-output coefficient distribution, whereas the corresponding effect at the top decile is a (statistically insignificant) 0.02 standard deviation increase. This gap between weakly and strongly linked industries stands up to a wide range of specifications and is robust to instrumental variables estimations that deal with key endogeneity concerns. At the top decile of the distribution of input-output coefficients, the effect of changes in transport costs on the geographic concentration of these input-output linked industries is essentially zero, whereas it is significantly negative and quantitatively large at the bottom decile. Taken as a whole, the statistically robust and economically meaningful effect of transport costs suggests they are a key determinant of the spatial economy and still matter substantially for interregional trade and the geographic organization of economic activity by holding vertically linked industries together.

The remainder of the chapter is organized as follows. The next section (Section 2: Modeling and estimating transport costs) develops a conceptual framework to model transport costs, and based on this, we illustrate how we measure them. This is followed by an analysis of the relationship between transport costs and the patterns of trade (Section 3: Transport costs and trade), which serves to motivate the remainder of the chapter. It focusses first on measuring the geographic concentration of industries (Section 4: Geographic concentration in Canada) and then moves on to relate changes in transport costs to the geographic concentration of individual industries and of vertically linked industry pairs (Section: 5: Transport costs and geographic concentration). This is followed by some concluding remarks (Section 6: Conclusions). We relegate details on the data that we use and additional results to a set of appendices.

2 Modeling and estimating transport costs

We first develop a simple conceptual framework, based on Behrens and Picard (2011), within which to model transport costs. That framework highlights the key aspects we need to take into account when estimating those costs. It also points to major endogeneity concerns that need to be dealt with when estimating the impact of transport costs on trade and geographic concentration.

Building on this framework, we establish conceptually how transport costs should be measured and then apply this to transaction-level estimates of transport costs for goods moved by truck. This serves as the foundation for the subsequent analysis that estimates how much transport costs influence the pattern of trade and, in particular, the geographic concentration of industry.

2.1 Conceptual framework

Consider two regions, f (fronthaul) and b (backhaul). There are M_f manufacturers (shippers) in region f, and M_b manufacturers in region b. Without loss of generality, we assume that f is the larger region. We henceforth refer to $M \equiv M_f/M_b \geq 1$ as the relative size of region f. Shipping goods between regions requires the services of freight carriers. Since our focus is on trucking, and since that sector is highly competitive in North America, we henceforth assume that there is perfect competition between carriers who operate at constant returns to scale.

Let m_f denote a manufacturer's marginal cost of production in region f and industry i. He faces demand $Q^i_{fb}(p^i_{fb}(m_f))$ for his commodity in the distant market b when quoting a delivered price $p^i_{fb}(m_f) \equiv p(m_f(\mathbf{X}^i_f), t^c_{fb})$, where \mathbf{X}^i_f is a vector of industry-region specific covariates, e.g., factor costs. To alleviate notation, we mostly suppress the industry index i in what follows.

Shipping requires the services of freight carriers who charge a per unit freight rate t_{fb}^c to ship commodity c from market f to market b.⁴ The index f is a mnemonic for the fronthaul part of the trip. There is also a backhaul part: the carrier needs to return from market b to the initial market f, irrespective of whether the truck is loaded or not. Hence, the carrier will transport goods from manufacturers in market b to market f, where demand for those goods is $Q_{bf}(p_{bf}(m_b))$ at the price $p_{bf} \equiv p(m_b(\mathbf{X}_b^i), t_{bf}^c)$.

Total demand for transport services from f to b and from b to f, conditional on the pair of freight rates t_{fb}^c and t_{bf}^c , is given by

$$D_{fb}(t_{fb}^c) = M_f Q_{fb}(p_{fb}(m_f(\mathbf{X}_f^i)), t_{fb}^c)$$
 and $D_{bf}(t_{bf}^c) = M_b Q_{bf}(p(m_b(\mathbf{X}_b^i)), t_{bf}^c)$.

Carriers provide transport services in both directions and face a logistics problem: they must commit to the capacity required by the largest demand on a return trip, i.e., the capacity required for the return trip is that in the direction of the largest demand for transport services. Taking into account this backhaul problem, the carriers' profits are thus given by:

$$\pi(t_{fb}^c, t_{bf}^c) = S_{fb}t_{fb}^c + S_{bf}t_{bf}^c - 2c(\mathbf{Y}, d) \max\{S_{fb}, S_{bf}\},\tag{1}$$

where S_{fb} denotes the supply of transport services from f to b, and where $2c(\mathbf{Y}, d)$ is the carriers' physical cost of a return trip that they must commit to. This cost depends on the distance d of a one-way trip, and on a vector \mathbf{Y} of carrier-specific factors like fuel prices and wages (it also includes the productivity of carriers in general).

⁴From the shippers' perspective, in addition to the payment to carriers, the cost of transportation services also includes inventory costs among other considerations (Baumol and Vinod, 1970; also, see Train and Wilson, 2008, for an application).

A competitive equilibrium is given by non-negative freight rates, t_{fb}^c and t_{bf}^c , and supplies, S_{fb} and S_{bf} , 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., $S_{fb} = D_{fb}$ and $S_{bf} = D_{bf}$; and (iii) carriers' profits (1) are maximized and equal to zero because of free entry. Using expression (1), profit maximization implies that $t_{fb}^c + t_{bf}^c = 2c(\mathbf{Y}, d)$. If $S_{fb} > S_{bf}$, then $t_{fb} = 2c(\mathbf{Y}, d)$ and $t_{bf} = 0$. The reverse holds if $S_{fb} < S_{bf}$. Hence, $t_{fb}^c > 0$ and $t_{bf}^c > 0$ requires that $S_{fb} = S_{bf}$ and that $t_{fb}^c + t_{bf}^c = 2c(\mathbf{Y}, d)$. Put differently, transport costs in both directions are positive if and only if freight rates adjust to 'balance' trade in the two directions: $D_{fb}(t_{bf}^c) = D_{bf}(t_{bf}^c)$.

To obtain 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 — conditional on freight rates that are taken as given — are equal to:⁶

$$p^c(m_f(\mathbf{X}_f^i);t_{fb}^c) = \frac{\sigma_i}{\sigma_i - 1}(m_f(\mathbf{X}_f^i) + t_{fb}^c), \quad \text{and} \quad p^c(m_b(\mathbf{X}_b^i);t_b^c) = \frac{\sigma_i}{\sigma_i - 1}(m_b(\mathbf{X}_b^i) + t_{bf}^c),$$

on the fronthaul and the backhaul parts of the trip, respectively. Here, σ_i denotes the industry-specific (constant) price elasticity of demand the manufacturers' face.

With CES demands, $Q_{fb} = A \cdot (p_{fb})^{-\sigma_i}$ and $Q_{bf} = A \cdot (p_{bf})^{-\sigma_i}$, where A is a shifter.⁷ Hence, to equalize quantities demanded in both directions requires that:

$$M_f \left[\frac{\sigma_i}{\sigma_i - 1} (m_f(\mathbf{X}_f^i) + t_{fb}^c) \right]^{-\sigma_i} = M_b \left[\frac{\sigma_i}{\sigma_i - 1} (m_b(\mathbf{X}_b^i) + t_{bf}^c) \right]^{-\sigma_i}.$$

 5 We can also consider the case of corner solutions, where the difference in freight rates is large enough so that freight rates on the backhaul part of the trip effectively fall to zero. This requires that $D_{fb}(2c) > D_{bf}(0)$. See Behrens and Picard (2011) for details. Zero freight rates in one direction are an extreme case that captures the idea that carriers are willing to transport at steep discounts in the direction of excess capacity. Consider for example "[...]the growing imbalance in manufacture trade between China and the U.S., which has become an issue for the transport sector as it creates important logistics problems associated with the 'empties'. About 60% of the containers shipped from Asia to North America in 2005 came back empty, and those that did come back full were often transported at a steep discount for lack of demand [...] shipping companies charge an average of \$1,400 to transport a 20-foot container from China to the United States. From the United States to China, companies charge much less: \$400 or \$500." (Behrens and Picard, 2011, p.280).

⁶Our specification abstracts from the fact that FOB prices do tend to increase with distance, at least for international shipments (see, e.g., Martin, 2012). This could be due to the fact that firms charge higher prices to more distant consumers if they are less price sensitive. The elasticity of demand for distant shipments may be lower if, for example, only richer consumers can buy goods as the CIF price increases with distance.

⁷We could assume that A differs between regions (A_f and A_b). This amounts to replacing $M \equiv M_f/M_b$ with $\widetilde{M} \equiv (A_f M_f)/(A_b M_b)$ in what follows, and it does not change our analysis.

We thus have

$$M^{-1/\sigma_i} \left[(m_f + t_{fb}^c) \right] = m_b + 2c(\mathbf{Y}, d) - t_{fb}^c, \tag{2}$$

where we have used $t_{fb}^c + t_{bf}^c = 2c(\mathbf{Y}, d)$, and where $m_f \equiv m_f(\mathbf{X}_f^i)$ and $m_b \equiv m_b(\mathbf{X}_b^i)$ to alleviate notation. Solving equation (2), we obtain the fronthaul freight rate as follows:

$$t_{fb}^{c} = \frac{1}{1 + M^{-1/\sigma_i}} \left[m_b - M^{-1/\sigma_i} m_f + 2c(\mathbf{Y}^c, d) \right].$$
 (3)

The freight rate for the backhaul trip can be recovered from $t^c_{fb} + t^c_{bf} = 2c(\mathbf{Y}, d)$, and the ad valorem rate is retrieved from $\tau^c_{fb} = 1 + t^c_{fb}/m_f$ and given by:

$$\tau_{fb}^{c} = \frac{1}{1 + M^{-1/\sigma_i}} \left[1 + \frac{m_b}{m_f} + \frac{2c(\mathbf{Y}^c, d)}{m_f} \right]. \tag{4}$$

In what follows, (3) and (4) will be key objects in our empirical analysis. Before proceeding, three important comments are in order.

First, as can be seen from (3) and (4), *freight rates are heterogeneous* along many dimensions. They depend, among others things, on the type of commodity shipped (e.g., dry bulk, liquid bulk, container), the industry in which shippers operate (which determines demand conditions via the price elasticities), the distance shipped, shippers' productivity (and characteristics that correlate with that productivity), and carriers' productivity.⁸ Carefully accounting for those dimensions is key in the econometric analysis.

Second, freight rates are fundamentally endogenous. They are prices that are set to clear markets and as such they do reflect supply and demand conditions. Note that even if freight rates are fully determined by carriers' costs — given competition in the trucking industry — costs themselves are endogenous to the spatial structure of the economy. Higher productivity in manufacturing in region f (a lower m_f), because of agglomeration economies (e.g., Combes et al., 2012; Combes and Gobillon, 2015), maps into lower prices for manufactured goods and affects freight rates by changing demand patterns. Imbalances in the endogenously determined geographic distribution of economic activity create imbalances in shipping patterns and influence freight rates via backhaul problems. Freight rates also depend on factor costs and on the distance shipped, both of which are again endogenous to the equilibrium distribution of economic activity. The key message is that freight rates and the geographic distribution of

⁸As seen from (4), the backhaul problem is more severe for carriers with high costs. If firm size has a substantial effect on the decision to be "loaded" on the backhaul, this needs to be controlled for. We do so using carrier fixed effects — which control for carrier-specific costs — in all our estimations.

⁹Note also that manufacturers and carriers may sort in non-random ways across locations (e.g., Forslid and Okubo, 2014; Gaubert, 2015). Because freight rates depend on the productivity of shippers and carriers, this creates additional endogenous variation in rates.

economic activity are jointly determined in equilibrium, and dealing with that simultaneity and the resulting endogeneity issues is crucial to assess the causal effect of transport costs on geographic concentration. We return to that point later in greater detail.

Last, most of the economic geography and international trade literature has subsumed τ_{fb}^c by an exogenous iceberg trade cost and disregarded how that trade cost depends on the different dimensions of heterogeneity and how it reacts to changes in the spatial structure of the economy. While the assumption of an exogenous trade cost may be useful in some contexts, it definitively is untenable when the key objective is to investigate the complex relationship between transportation and economic geography. Much of the literature has thus used geographic distance as an exogenous proxy for trade costs. While reasonable, the absence of time-series variation in the distance variables precludes strong identification since many omitted variables positively correlated with distance (the 'dark trade' costs; see Head and Mayer, 2013) cannot be taken into account.¹⁰

2.2 Estimating transport costs

Transport costs can be measured on a multiplicative (i.e., iceberg) or an additive (i.e., ad valorem) basis (see, e.g., Irarrazabal et al., 2015). As we have already argued, exogenous iceberg transport costs ignore how these costs depend on many dimensions of heterogeneity and spatial structure. Additionally, multiplicative transport costs imply higher priced goods are more expensive to ship, because they are a constant share of the value of the good. By construction, additive transport costs are not related to the value and, moreover, allow the effect of transport costs on delivered prices to vary between high and low value goods, that, in turn, have been shown to affect consumption patterns across markets (Alchian and Allen, 1964). As a result, in the empirical international trade literature the now standard approach is to estimate transport costs on an ad valorem basis ('ad valorem transport costs'; AVTC).

Internationally, transport costs are typically estimated using customs documents that measure trade on a fob and CIF basis (Hummels, 2007; see also Globerman and Storer, 2010). Domestically, due to a lack of information about the value of goods being shipped, statistical agencies measure transport costs on a per tonne kilometre (and sometimes a per kilometre)

¹⁰Even more sophisticated measures such as Generalized Transport Costs (e.g., Combes and Lafourcade, 2005) provide little help with the time-series variation and are, therefore, highly colinear with distance. Storeygard (2016) interacts road distances with oil prices to construct a time-varying proxy for transport costs in Africa. While this approach is an improvement and plausible at an aggregate level (e.g., diesel prices are a fairly good predictor for overall trucking costs in Canada), it provides no between-industry variation that is useful for looking at the geographic concentration of specific industries.

¹¹Of course, trade costs such as insurance may be multiplicative (i.e., proportional to value), but we set this aside here to simplify the discussion.

basis.¹² Hence, transport costs on international trade are measured on an ad valorem basis, but often with considerable error (see Head and Mayer, 2013), while domestically transport costs are measured more directly and promise more accurate estimates, but ones that do not provide a means to observe their effect on relative prices across space.

Regarding transport costs within countries, Combes and Lafourcade (2005) estimate generalized transport costs (GTC) for trucking in France, using detailed data on factor costs and road networks. Since their estimates do not vary by industry, they cannot be used to learn anything about the geographic concentration of individual industries. Also, the very high correlation of their measure with distance limits it usefulness in the time-series dimension. To obtain industry-specific estimates that vary across time, we have to zoom in on transport costs as a share of goods prices, i.e., ad valorem based measures. As of now, we are unaware of work that estimates transport costs on domestic trade on an ad valorem basis and uses it in a panel.

2.2.1 Ad valorem rates

Measuring transport costs on an ad valorem basis allows us to estimate their effect on relative prices. To make this concrete, consider the origin f (factory gate) price (p_f^c) and destination b (delivered) price (p_{fb}^c) of commodity c. The delivered unit price is simply the origin unit price plus transport costs incurred to move that unit: $p_{fb}^c = p_f^c + r_{fb}^c$, where r_{fb}^c is the revenue to carriers (the cost to shippers) of moving a unit of commodity c from f to b and is a measure of t_{fb}^c .¹³ Therefore, the effect of ad valorem transport costs on relative prices is $p_{fb}^c/p_f^c = 1 + r_{fb}^c/p_f^c$, where $\tau_{fb}^c \equiv 1 + r_{fb}^c/p_f^c$ is the ad valorem rate. This is an additive measure of transport costs, where τ_{fb}^c varies with the price of the commodity being shipped. We express ad valorem rates as either τ_{fb}^c or as a $\tau_{fb}^c - 1$, where the latter is transport costs as a proportion of the origin price.

Of course, one of our primary challenges is developing an estimate of r_{fb}^c . To do so, we use information derived from Statistics Canada's Trucking Commodity Origin-Destination Survey (TCOD). Using shipping documents (waybills), it collects information on the origin and destination and (network) distance of each shipment, as well as their tonnage and commodity. Crucially, carrier revenues are also reported and, therefore, we have a means to measure r_{fb}^c directly. Missing, however, is the value of goods shipped, p_f^c . The latter is derived by using value per tonne estimates, which can be thought of as a unit price, estimated using a special 2008 transaction level international trade file that reports both the value and tonnage of exports and imports by mode of transport.¹⁴ These 6-digit Hs commodity estimates (indexed h) are first

¹²See, for example, Cansim Table 403-0004 (http://www5.statcan.gc.ca/cansim/home-accueil?lang=eng).

¹³For the sake of simplicity we ignore other margins like wholesale that may also affect the delivered price.

¹⁴See Appendix 6 for more detail on the source data and the construction of value per tonne estimates.

concorded with the Standard Classification of Transported Goods (sctg; indexed c) used to classify shipments on the tool ($p_{2008}^h \to p_{2008}^c$). Export price indices, \mathbb{P}_t^c , where t denotes time, are then used to project the value per tonne estimates backwards and forwards: $p_t^c = \mathbb{P}_t^c \times p_{2008}^c$. Finally, multiplying the value per-tonne estimates by the tonnage q of each shipment k provides an estimate of their value: $x_{tk} = p_t \times q_k$, with the commodity class suppressed to simplify notation. Note that x_{tk} is weighted to ensure trade within and across provinces by commodity and year add to known trade totals from the input-output accounts (see Appendix 6 for more detail). The estimated ad valorem rate for goods shipped from f to b is:

$$\tau_{fbt} - 1 = \frac{\sum_{k} r_{fbtk}}{\sum_{k} x_{fbtk}} = \frac{r_{fbt}}{x_{fbt}}.$$
 (5)

The ad valorem rate (5) can be calculated for aggregate flows or by commodity and, when concorded with NAICS codes, can be expressed on an industry basis. We return to that point later in Section 5.1.

2.2.2 Estimating revenue to carriers

As we have shown in the conceptual framework, freight rates depend on a number of factors. They are usually commodity specific, because the carrier's cost depends on the commodity shipped. They also depend on the distance shipped (e.g., economies of long haul, density economies, etc.) and, obviously, on the carrier's productivity, as per $c(\cdot)$. This suggests that freight rates should be estimated controlling for carrier fixed effects, commodity fixed effects, and distance shipped. We control for this when estimating the ad valorem rates. Namely, ad valorem rates are estimated for a representative shipment over a 500 kilometres fixed distance, i.e., we set $d=\overline{d}$ in the ad valorem rates. Hence, our rates should be independent of the distance shipped. Furthermore, in the estimation of τ^c_{fb} in (4), we include carrier and commodity fixed effects. This allows us to control for the $c(\mathbf{Y}^c,d)$ part of (4), as well as for other factors such as carriers' differential probability of being loaded on the return trip.

In concrete terms, ad valorem rates are estimated first by using a model to predicted trucking firm (carrier) revenues for a 500 kilometres trip by commodity for the average tonnage using shipment (waybill) data from Statistics Canada's Trucking Commodity Origin-Destination Survey TCOD (see Brown, 2015, for details, and Behrens et al., 2015, for an application). We estimate the 'prices' charged by trucking firms as a function of distance shipped, tonnage, and a set of commodity and firm fixed effects. We assume firms set prices such that both fixed and variable (linehaul) costs are just covered, as should be the case in a competitive market. Firms are assumed to set prices based on a fixed component and kilometres shipped: $r_{m,k}^c = \alpha + \beta d_k$, where $r_{m,k}^c$ is the revenue earned by carrier m for shipment k composed of commodity k0, k2 is

the fixed-price component, β is rate per kilometre, and d_k is the distance shipped. As firms may also price on a per tonne-km basis, and assuming firms set prices based on an unknown average tonnage q^* shipped, this implies that the rate per tonne-km is $r_{m,k}^c = \alpha + (\beta/q^*)(d_kq^*)$. For loads less (greater) than q^* the implicit price per tonne-km will be scaled upward (downward) to ensure that the price on a per-km basis is maintained, which is captured by the following function:

$$r_{m,k}^c = \alpha + \left[\frac{\beta}{q^*} + \phi(q^* - q_k)\right] d_k q_k = \alpha + \left(\frac{\beta}{q^*} + \phi q^*\right) d_k q_k - \phi d_k q_k^2, \tag{6}$$

where q_k is the actual tonnage shipped and $\phi(q^* - q_k)$ is the scaling factor. Factoring out the known tonnage results in a flexible function that allows firms to price using either rule (or some hybrid of the two). Equation (6) can be estimated using a simple quadratic form

$$r_{m,k}^c = \alpha + \delta d_k q_k + \omega d_k q_k^2 + \lambda_{mc} + \mathbf{X}\gamma + \epsilon_{m,kc}$$
, with $\delta = \beta/q^* + \phi q^*$ and $\omega = -\phi$, (7)

augmented with carrier-commodity fixed effects, λ_{mc} , and a vector **X** of controls to account for the quarter the shipment was made, the effect of empty backhauls on prices (distance shipped), and fuel prices; and an error term, $\epsilon_{m,kc}$. The model is estimated across three types of carriers — truck-load, less-than-truck load, and specialized — with the estimated rate being the weighted average by value of the three by commodity (see Brown, 2015, for a detailed discussion of the data and model). The entire period, 1994 to 2009, is used to predict prices in order to bring as much information as possible to bear on the cross-sectional estimates by commodity. To simplify the presentation of the model, we suppress notation for time and its interaction with fixed and linehaul costs. Predictions from the model closely match observed annual prices on aggregate. Prices charged to shippers are predicted by commodity using their average tonnage for a 500 kilometre trip. The final predicted price, \hat{r}_t^c , is the weighted average by value of the

¹⁵Each carrier type differs based on the business model or technology used. Truck-load specialized carriers typically transport full loads from point to point, while less-than-truck load carriers move partial loads from depot to depot where shipments are consolidated and broken down for local distribution. Specialized carriers differ from truck load and less-than-truck load by the type of capital used (e.g., tank trailers versus box).

¹⁶Because prices are measured on a quarterly basis, fixed and variable costs are permitted to vary through time. Time enters as trends through a spline with knots set to reflect changing trends in a trucking price index generated from the same file.

¹⁷While we do not directly control for the time costs of transportation they will be, at least partially, embedded in the transport prices (which would capture quality of service for time-dependent trips). We also do not include origin-destination (dyadic) fixed effects in the model to account for the effect of backhauls on prices. This is because the addition of origin-destination fixed effects to a model that already includes carrier-commodity fixed effects would leave little variation for most origin-destination pairs. Furthermore, as a practical matter, prior to 2004 there is no detailed geography (i.e., postal codes) to build a constant geography through time, which could introduce considerable error. The absence of dyadic fixed effects should not be a problem since we control in what follows for geography using instrumental variables.

three carrier types by commodity and year. Ad valorem trucking rates $\hat{\tau}_t^c$ are measured using the value of shipments for the average tonnage shipped across commodities.¹⁸

Finally, using an industry-commodity concordance, the ad valorem transport rates in 2008 for commodities, τ^c_{2008} , are aggregated to an industry basis, τ^i_{2008} . We need to transform the estimates to an industry basis since we observe geographic patterns of industries and not commodities. Moreover, by moving to an industry basis we can use industry-based price indices to project the ad valorem rates through time. These are less subject to sampling variation than commodity-based estimates that are ultimately drawn from the TCOD. To generate a time series, yearly trucking 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 ad valorem transport rate time series:

$$\tau_t^i = \frac{\mathbb{P}_t^{\text{trans}}}{\mathbb{P}_t^i} \tau_{2008}^i. \tag{8}$$

While (8) already takes into account a wide variety of factors expected to influence freight rates, they may still be influenced by industry-region specific components, via the manufacturers' marginal cost $m_f(\mathbf{X}_i^f)$, and an industry-specific component, notably through the industry's price elasticity of demand. Again, this needs to be controlled for in the analysis, which is accomplished through the inclusion of industry fixed effects. Furthermore, freight rates depend on shippers' productivity (and characteristics that correlate with that). Since we do not have shipper information, systematic differences in shipper characteristics may still correlate with our ad valorem rates below, and this will require some careful instrumenting and controls to get rid of those effects.

3 Transport costs and trade

Declining tariff and non-tariff barriers to trade realized through successive rounds of GATT negotiations and the implementation of a series of regional trade agreements (e.g., NAFTA) has increased the relative importance of transportation as a source of trade costs (Hummels, 2007). As Hummel's (2007, p.136) notes, "[...] exporters paid \$9 in transport costs for every \$1 they paid in tariff duties" and as a result transport costs have a significant effect on relative prices across exporters. This has increased interest in measuring transport costs and identifying their effect on trade patterns (see Anderson and van Wincoop, 2004; Head and Mayer, 2013, 2014).

¹⁸Since the value of shipments are not reported, they have to be estimated by multiplying the average tonnage shipped for each commodity by their respective value per tonne derived from an 'experiment export trade file' produced only in 2008, as explained in Section 2.2.1.

Building off of this literature, we develop estimates of the effects of transport costs on trade flows, the relationship between transport costs and distance, and decomposition of the effect of distance itself on trade into transport cost-related and other, to use Head and Mayer's (2013) language, 'dark' trade costs.

Our interest in the influence of transport costs on trade is not made in isolation. We are interested in determining how transport costs influence the location and co-location of industry. There has been a tendency to assume that because transport costs have fallen to such an extent through the last century (see Glaeser and Kohlhase, 2004) that they no longer influence trade patterns and, by implication, the location choices of firms. While there is ample evidence that distance influences the volume of trade, with an elasticity averaging about -1 (Head and Mayer, 2014) the direct role of transport costs, as opposed to other distance related trade costs, is far from clear. For instance, Head and Mayer (2013), estimate transport costs contribute somewhere between 4% and, at most, 28% of the overall effect of distance on trade, with the remainder attributed to 'dark' trade costs. So to help motivate the discussion of plant location and co-location, we develop the necessary components to decompose the effect of distance on trade into portions attributable to transport costs and a residual, i.e., 'dark' costs.

3.1 Estimating the effects of transport costs on trade

At issue is the degree to which transport costs influence the volume of trade. We start with aggregate trade in order to keep comparability with (much of) the gravity literature. The standard approach is to estimate the influence of transport costs on trade through cross-sectional variation using the simple model suggested by Head and Mayer (2014):

$$\ln x_{fb} = \alpha + \delta_f + \phi_b + \beta \ln \tau_{fb} (d_{fb}, dark_{fb}) + \ln \epsilon_{fb}. \tag{9}$$

where δ_f and ϕ_b are origin and destination fixed effects, respectively, and τ_{fb} , as noted above, captures the effect of transport costs on the delivered price relative to the origin price. Using mean trade between Economic Regions (ERS) from 2004 to 2012, results in an OLS estimate of $\beta = -6.40$ (with robust standard error of 0.308) from (9), which is in line with Head and Mayer's (2014) meta analysis. While equation (9) is suggestive that transport costs have a strong influence on trade, because other 'dark' trade costs (e.g., marketing costs) vary with distance it is difficult to isolate the effect of transport costs. However, since distance is invariant over time, the time series variation in transport costs provides a means to isolate its effect from (time-invariant) 'dark' trade costs correlated with distance.

To overcome this, our approach is to estimate trade across a panel of ERS using the following estimation equation:

$$\ln x_{fbt} = \alpha + \gamma_{fb} + \delta_{ft} + \phi_{bt} + \sigma_t + \beta \ln \tau_{fbt} + \ln \epsilon_{fbt}. \tag{10}$$

where γ_{fb} is an origin-destination (dyadic) fixed effect, δ_{ft} and ϕ_{bt} are respective time varying origin and destination fixed effects, and σ_t is a time fixed effect. Equation (10) is estimated using OLS. The dyadic fixed effect controls for structural difference across region pairs (e.g., flows driven by strong upstream-downstream buyer-supplier links), while allowing for time series variation in bilateral flows.

There are two primary challenges associated with estimating (10). First, in practice, we measure transport costs for a shipment of units q between f and b. Hence, the value of the shipment is $x_{fbt} = p_t q_{fbt}$. This means for any trade equation with ad valorem transport costs on the right-hand-side we have to pay attention to the mechanical relationship induced by also having p_t on the left-hand-side. Depending on the elasticity of demand, a rise in p_t will result in a fall in τ_{fbt} and a rise in x_{fbt} . Therefore, we need to control for any generalized rise in prices, while also allowing for ad valorem transport costs to continue to vary. This is important because, a rise in prices affects the delivered price relative to the origin price, which, in turn, will affect quantity demanded at the destination. A change in ad valorem transport costs resulting from a change in prices is still relevant. To purge x_{fbt} of the effect of a general rise prices, we first recognize that $lnx_{fbt} = lnp_t + lnq_{fbt}$. Hence by including time dummies (σ_t) , the price effect is taken into account. This approach, however, does not address shifts in the commodity composition of trade through time and so the model also needs to be estimated using dissaggregated trade flows as well. Hence, we also use commodity flows to control for differences in prices (at the commodity level) and compositional effects.

Second τ_{fbt} is potentially endogenous, because the volume of trade can influence the prices charged by carriers. The direction of this relationship is, however, unclear. First, rising trade between an origin and destination will result in a rightward shift in the demand curve for transport services, pushing prices upwards, all else being held equal. This will occur because less efficient, high cost carriers enter the market. It may also occur because rising trade between an origin and destination results in an imbalance of trips, putting upward pressure on prices as the relative demand for the outbound portion of the journey increases relative to the inbound trips. ¹⁹ Of course, rising trade may also result in lower prices under some circumstances. If an origin-destination pair has a low volume of trips, rising trade may reduce carrier costs if the number of trips increases the likelihood of a return load for that particular trip — that is, at the end of a leg the truck would not have to wait for an extended period of time at the destination for a return load, or have to travel empty (deadhead), to find a return load elsewhere.

If higher trade leads to higher transport rates, the estimated effect of transport costs on trade will be underestimated. Of course, the opposite holds true if rising trade leads to lower

¹⁹See Behrens and Picard (2011) for a theoretical treatment and Jonkeren et al. (2011) and Tanaka and Tsubota (2016) for empirical applications.

transport rates. In the cross-section, we find a positive association between the number of trips between regions and balance of trips on rates per tonne-km (see Appendix B). Hence, to the extent it influences the results, endogeneity likely leads to an underestimation of the effect of transport costs on trade levels.

Table 1: Fixed effects estimates of economic region trade as function of transport costs.

	(1)	(2)	(3)	(4)	(5)	(6)
					Deflated	Commodity
Transport costs	-1.048 ^a	-2.662 ^a	-3.538 ^a	-5.364 ^a	-5.331 ^a	-7.363 ^a
	(0.118)	(0.151)	(0.200)	(0.209)	(0.208)	(0.093)
Year fixed effects	Yes	Yes	Yes	Yes	No	No
Commodity-year fixed effects	No	No	No	No	No	Yes
Observations	25,830	25,830	25,830	25,830	25,830	169,200
R-squared	0.104	0.143	0.146	0.159	0.158	0.150

Notes: Transport costs are measures as $\ln(\tau_{fb})$. Trade is also log transformed. Robust standard errors are in parentheses. All models include time varying origin and destination fixed effects and dyadic (origin-destination pair) fixed effects and are estimated using a balanced panel of 2,870 region pairs over 9 years. Within region trade is excluded as well as trade less than 1 km and greater than 5,000km. Estimates reported in column (1) include all observations to calculate the ad valorem rate, while estimates reported in columns (2) to (4) are based on progressively more restrictive samples. Column (2) excludes observations that may be errors or are idiosyncratic (e.g., those where revenues are set to \$1). Column (3) further excludes observations with ad valorem rates above the 95 percentile for a given 2-digit commodity-distance class, while column (4) additionally excludes observations with an ad valorem rate above 2. Estimates for columns (5) and (6) use the same sample restrictions as column (4). Column (5) reports the estimates using trade deflated at the 5-digit commodity level. Column (6) reports estimates based on trade by 2-digit commodity and includes 2-digit commodity-dyadic fixed effects and commodity-year fixed effects. Huber-White robust standard errors are in parentheses. Coefficients significant at: a 1%; b 5%; and c 10%.

The elasticity of trade with respect to transport costs is presented in Table 1. Trade is measured between provincial Economic Regions (ERS), where the sample of flows is restricted to a balanced panel of region pairs that traded continuously over the entire 9-year period from 2004 to 2012. Within region trade is excluded and flows are constructed from shipments that move more than 1km and less that 5,000km, which tend to be idiosyncratic. Transport costs are measured as $\ln(\tau_{fb})$ in order to reflect their influence on relative prices. One of the challenges of shifting from a cross-sectional dataset to a panel is the considerable amount of noise that can be introduced by errors in the micro-data shipment file. To address this, we progressively eliminated suspect observations from the sample used to estimate ad valorem rates.

Including all of the observations to estimate the ad valorem rate (see Table 1, column (1)) results in a statistically significant elasticity of about -1. However, the elasticity on transport costs becomes much more negative, while remaining significant, as we progressively use

more restrictive samples. Column (2) excludes poorer quality observations.²⁰ Column (3) further excludes observations with ad valorem rates above the 95 percentile for a given 2-digit commodity-distance class, while column (4) additionally excludes observations with an ad valorem rate above 2. Using these more restrictive samples, the transport cost elasticity falls monotonically to -5.4, a pattern consistent with the progressive elimination of errors on the file. Estimates for columns (5) and (6) use the same sample restrictions as column (4). Column (5) reports the estimates using trade deflated at the 5-digit commodity level and provides qualitatively similar estimates. Column (6) reports estimates based on trade by 2-digit commodity and includes 2-digit commodity-dyadic fixed effects and commodity-year fixed effects. Its estimated elasticity is even lower at -7.4. These are estimates that are close to the median (-5.03) and average (-6.74) identified in Head and Mayer's (2014) meta analysis, but are ones that only take into account transport costs. Hence, we find strong evidence that transport costs, through their effect on relative prices, influence the patterns of trade. Still left open to question is the degree to which transport costs matter. We can provide a perspective on this, by using these estimates to decompose the standard gravity model distance coefficient into transportation and 'dark' costs.

3.2 Estimating the effects of distance on transport costs

Before assessing the effect of distance on transport costs, we first develop a description of their relationship. Table 2 reports the mean ad valorem rate, as well as the revenue per tonne-km and average value per tonne by distance class over the 2004 to 2012 period. We expect ad valorem rates to increase with distance, but at a decreasing rate as the effect of fixed costs on transport rates fall relative to variable (linehaul) costs and as the commodity composition of trade shifts towards higher value goods for whom higher transport costs have smaller relative effect on delivered prices.²¹ This is indeed the case. Ad valorem rates increase (monotonically)

²⁰It is important to keep in mind that these data are built from shipment level files intended to estimate national aggregates. We, therefore, had to take care to mitigate the effects of possible errors stemming from data capture/entry and other sources and idiosyncratic observations on our estimates. As a first step, observations were excluded with extremely low tonnages or extremely low/high revenues (e.g., revenues of \$1), which can strongly influence ad valorem rates. Also shipments are excluded if they have tonnages that are above those generally permitted by regulations and, therefore, do not represent common shipments. Finally, observations are excluded if the weight used to benchmark the value of the shipment to provincial trade totals is above a threshold (i.e., above 500), because these weights can strongly influence estimates of ad valorem rates if applied to a small number of error driven or idosyncratic observations.

²¹The notion that rates increase at a decreasing rate is a long-standing tenet of transport known as the 'tapering principle' (see, e.g., Locklin, 1972). Rates 'taper', i.e., decrease, with distance. One reason is that fixed costs of loading and unloading can be spread over more miles. In a multimodal setting, modal choice generates a concave

with distance, from a low of 1.5% for shipments between 1 and 50km to 8.1% for shipments above 2,500km. Reflecting fixed costs associated with shipping goods (e.g., time at terminals and overhead costs for the carrier), revenues per tonne-km are high over short distances, \$0.93 from 1 to 50km, but fall rapidly with distance as prices reflect the averaging down of fixed costs per shipment such that above 500km rates are about \$0.10 per tonne-km. Also meeting expectations, trade flows shift towards higher value commodities over longer distances, with the value per tonne of goods shipped over 2,500km being thrice greater than those shipped 1 to 50km. Finally, demonstrating the additive nature of transport costs, ad valorem rates differ considerably across commodities (see Figure 1). For instance, for goods shipped between 1,000 and 2,500km the lower and upper quartile commodity rates are 0.03 and 0.16, a five fold difference. Furthermore, confirming the mean ad valorem rate pattern, the full distribution of rates rises with distance.

Table 2: Trucking ad valorem rates, revenue per tonne-km, and value per tonne by distance class.

	Ad valorem rate	Revenue per tonne-km	Average value per tonne
	$(au_{fb}-1)$	(dollars)	(dollars)
1 to 50km	0.015	0.93	1,274
51 to 100km	0.019	0.31	1,243
101 to 500km	0.029	0.18	1,521
501 to 1,000km	0.041	0.11	1,973
1,001 to 2,500km	0.057	0.10	2,698
2,501 to 5,000km	0.081	0.09	4,026

Notes: On a shipment basis, observations are excluded where ad valorem rates are above the 95th percentile by 2-digit commodity-distance class, if ad valorem rate are above 2, and a series of other restrictions to exclude poorer quality observations. Furthermore, shipments of less that 1km and greater than 5,000km are excluded. Revenue per tonne-km are not comparable to published totals, because of these restrictions, and because observations are weighted to ensure trade adds to provincial control totals from the input-output accounts. The results we report are average annual rate (2004 to 2012).

Estimates of the effect of distance on transport costs (measured by $\ln(\tau_{fb})$) are presented in Table 3. Their association is estimated using a pooled set of ER flows from 2004 to 2012. As with the trade flow regressions, we progressively improve the quality of the estimates of transport costs, with the most restricted sample corresponding to the descriptive statistics presented in Table 2. As Head and Mayer (2013) observe, improving the quality of the measurement of transport costs (e.g., eliminating shipments with an ad valorem rate above 2) reduces the estimated elasticity. However, we find a low end estimate of 0.041, while they estimate it to be

distance-cost relationship since the lower envelope of a family of concave costs functions — one for each mode — is itself concave.

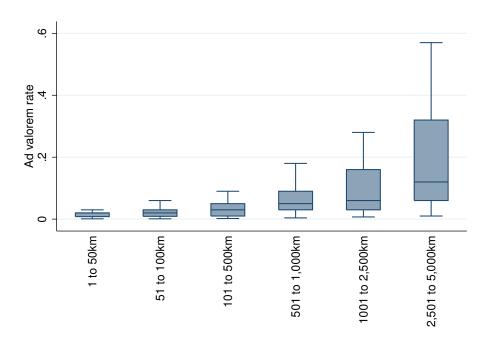


Figure 1: Commodity-based ad valorem rates by distance class.

Notes: Ad valorem rates $\tau_{fb}-1$ are calculated across commodities by distance class over the 2004 to 2012 period. Commodities are based on the 2-digit sctg classification. The same restrictions to the data are applied as in Table 2.

o.o26. All in all, our estimated elasticities are higher than most estimates from the literature, which tend to rely on international trade flows that are dominated by the lower cost marine transportation mode.

3.3 How important are 'dark' trade costs?

Finally, we turn to how large is the effect of transport costs on trade as compared to distance? The effect of distance on trade is, of course, undeniable. As Figure 2 illustrates, the density of shipments fall sharply with distance.²² That is, while up to 100km the density of shipments rises as market size increases, beyond 100km the density of shipments drops off rapidly.²³ Hence, even in this simple exposition, the influence of distance on trade is clearly evident. An elasticity of -1 for distance and trade comes as no surprise. Head and Mayer (2013) note that the coefficient on distance (δ) in the standard trade model can be decomposed into a two parts,

²²See Hillberry and Hummels (2008), for comparable U.S. evidence using Commodity Flow Survey Microdata. Behrens, Mion, Murata, and Suedekum (2017) show how this pattern emerges from a heterogeneous firms model featuring choke prices and costly interregional trade (see their Figure 8).

²³While we restrict the analysis to trucking freight rates, the kernel density is based on shipments by truck and rail. Hence the drop in shipments after 100km is not due to rail modal share of shipments rising with distance.

Table 3: Effect of distance on ad valorem transport costs.

	(1)	(2)	(3)	(4)
Distance	0.072^{a}	0.059 ^a	0.049 ^a	0.041^{a}
	(0.002)	(0.001)	(0.001)	(0.001)
Observations	34,529	34,529	34,529	34,529
R-squared	0.059	0.072	0.085	0.095

Notes: Transport costs are measured as $ln(\tau_{fb})$, which is the dependent variable. All models are based on a set of economic region flows pooled over the 2004 to 2012 period and include year fixed effects. Robust standard errors clustered on year are in parentheses. Distance and transport costs are log transformed. Trade flows less than 1km and greater than 5,000km are excluded. Estimates reported in column (1) include all observations to calculate the ad valorem rate, while estimates reported in columns (2) to (4) are based on progressively more restrictive samples. Column (2) excludes potentially poorer quality observations. Column (3) further excludes observations with ad valorem rates above the 95 percentile for a given 2-digit commodity-distance class, while column (4) additionally excludes observations with an ad valorem rate above 2. Huber-White robust standard errors are in parentheses. Coefficients significant at: a 1%; b 5%; and c 10%.

the effect of trade costs on trade and distance on trade costs:

$$\delta \equiv \frac{\partial \ln \text{Trade}}{\partial \ln \text{Trade Costs}} \frac{\partial \ln \text{Trade Costs}}{\partial \ln \text{Distance}} = -\epsilon \rho \approx -1, \tag{11}$$

where ϵ is the elasticity of trade with respect to trade costs and ρ is the elasticity of distance on trade costs. Ostensibly, we have developed estimates of both above, albeit using transport costs instead of trade costs. As they note, reasonable estimates of ϵ and ρ do not add to -1, with the remainder attributed to unobserved 'dark' trade costs. We use our estimates of these parameters to decompose the elasticity of distance with respect to trade δ into an observed transport cost and an unobserved 'dark' trade cost component.

Using columns (4)–(6) in Table 1 and columns (2)–(4) in Table 3, and assuming a distance elasticity of -1, we find that 'dark' trade costs account for 57–78% of trade impediments, whereas transport costs account for 22–43% (see column (1) of Table 4). These figures are lower than the ones reported for international trade by Head and Mayer (2013) of 72–96% or Allen (2014) of 90%. Hence, we estimate the share of transport costs within Canada to be about double international trade estimates. As expected, there are less 'dark' trade costs within countries than across countries.

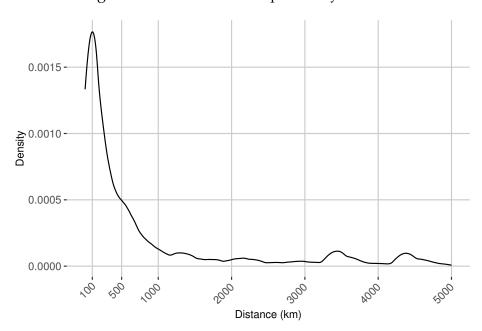


Figure 2: Distribution of shipments by distance.

It is also likely the case that ours is still a low estimate, because we are assuming the elasticity of distance on trade is -1. Bemrose, Brown, and Tweedle (2016) find the distance elasticity δ to be between -0.86 and -0.76.²⁴ Using these measured elasticities on distance results in a low- and high-end estimated contribution of transport costs of 25% and 57%, respectively (see columns (2) and (3) of Table 4). The broad conclusion we draw is that while 'dark' trade costs matter, transport costs themselves cannot be relegated to being of secondary importance when considering trade costs. As we will show, this is apparent in the influence of transport costs on the location and co-location of industry.

4 Geographic concentration in Canada

We have discussed in detail the estimation of transport costs in the foregoing section. The second key ingredient for our analysis is a measure of the geographic concentration of industries. Figure 3 displays two examples of geographic patterns in south-western Ontario: (i) the agglomeration of a single industry ('Motor Vehicle Manufacturing') in panel (a); and (ii) the coagglomeration of two industries ('Motor Vehicle Manufacturing' and 'Motor Vehicle Parts Manufacturing') in panel (b). While Figure 3 visually suggests that these industries are agglomerated and coagglomerated, we need to measure those patterns precisely to use them in

²⁴These are estimated for distances of 500km or more, where trade is measured between hexagons (75km per side) and Forward Sortation Areas (FSAS), respectively.

Table 4: Share of transport costs vs 'dark' trade costs.

		$\widehat{\delta} \equiv \partial \ln x_{fb} / \partial \ln d_{fb}$				
		Trade literature	Canada, economic			
		meta estimate	region estimates			
		(1)	(2) (3)			
$\widehat{\epsilon}$	$\widehat{ ho}$	-1.00	-o.86	-0.76		
-5.331	0.041	0.219	0.254	0.288		
	0.049	0.261	0.304	0.344		
	0.059	0.315	0.366	0.414		
-7.363	0.041	0.302	0.351	0.397		
	0.049	0.361	0.420	0.475		
	0.059	0.434	0.505	0.572		

Notes: Estimates of $\hat{\epsilon}$ from columns (5) and (6) of Table 1. Estimates of $\hat{\rho}$ from columns (2)–(4) of Table 3. Estimates of $\hat{\delta}$ from Head and Mayer (2014) in column (1); and from Bemrose, Brown, and Tweedle (2016) in columns (2) and (3).

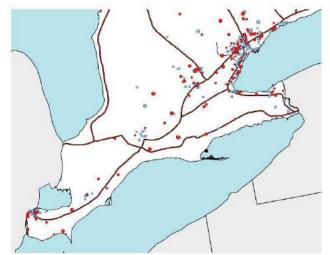
our subsequent analysis.

Figure 3: Motorvehicle manufacturing industries, agglomeration and coagglomeration patterns.

(a) Agglomeration (NAICS 3363).

(b) Cogglomeration (NAICS 3361-3363).





4.1 Measuring geographic concentration

There are various ways to measure geographic concentration, and an in-depth discussion is beyond the scope of this chapter (see Combes, Mayer, and Thisse, 2008). The two most well-known measures are the Ellison and Glaeser (1997) index, and the Duranton and Overman (2005) *K*-densities. In what follows, we use the latter to exploit the microgeographic nature

of our data. Roughly speaking, the Duranton-Overman (henceforth, DO) measure looks at how close establishments are relative to each other by considering the distribution of bilateral distances between them. The K-density is a kernel smoothed version of the distribution of bilateral distances between pairs of plants, either unweighted or weighted by the plants' employment. It gives — for each distance d — the share of bilateral distances between pairs of plants in the industry. Since the K-density is a distribution, we can also compute the associated cumulative distribution (CDF), which captures the share of plant-pairs in the industry that are located closer than distance d from each other. The DO measure can also be used to test the statistical significance of the observed geographic patterns. The idea is to apply sampling and bootstrap techniques to compare the observed distribution of bilateral distances to a set of distances obtained from samples of randomly drawn plants among all locations where we observe manufacturing plants.

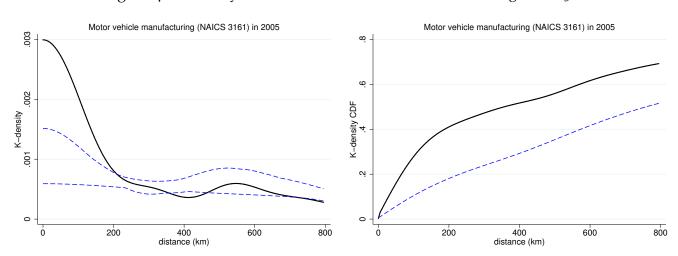


Figure 4: K-density and CDF for 'Motor Vehicle Manufacturing' in 2005.

Figure 4 shows an example of the DO K-density (left panel) and its associated CDF (right panel). The blue dashed lines depict the 90% confidence band in the left panel, and the cumulative midpoint of that confidence band in the right panel. As one can see, 'Motor Vehicle Manufacturing' (NAICS 3361) is geographically concentrated at short distances (less than 200 kilometres). That concentration is statistically significant. The CDF in the right panel provides a

²⁵The Ellison-Glaeser index has two shortcomings. First, it requires pre-defined spatial units (e.g., Census divisions in Canada or counties in the U.S.) and it is hence sensitive to the well-known 'Modifiable Areal Unit Problem': depending on the spatial units chosen, the index will take different values for the same spatial distribution of economic activity (see, e.g., Behrens and Bougna, 2015, who compute the EG index for different spatial units in the Canadian case). Second, the index is largely 'aspatial' in the sense that the relative position of the spatial units does not matter. Exchanging the geographic positions of the census divisions of British Columbia and Quebec will not affect the value of the index. The Duranton-Overman index fixes those two important problems.

natural measure of the *strength* of that concentration. As one can see, about 40% of plant-pairs in that industry are located less than 200 kilometres from each other, whereas under a random allocation that same share would be slightly less than 20%.

In what follows, we measure the geographic concentration of industries using their DO CDFs at different distances d. For most of our analysis we will set d to 100 kilometres, but we will show that our key results are not very sensitive to that choice. While many studies look at the geographic concentration across industries we will also look at the geographic concentration across industry pairs. To do so, we extend the DO measure to look at the *coagglomeration* of two different industries (see Duranton and Overman, 2005, 2008). In that case, the distribution applies to the bilateral distances between all pairs of plants in the two industries (i.e., pairs ij where i is in one of the industries and j in the other). We will use both measures in what follows. However, for reasons made clear below, using the coagglomeration measures provides much stronger results as we have many more observations and as we can interact our transport cost measures with input-output patterns across industries.

4.2 Descriptives

How widespread is the geographic concentration of industries in Canada? Table 5 — which summarizes results of K-density estimations for 85 4-digit manufacturing industries and 3,570 unique 4-digit industry pairs — shows that between 50–70% of industries and industry pairs are not randomly located. In other words, the majority of industries have location patterns that depart significantly from 'randomness'. The trend over time is, however, clearly towards less localization and more spatial randomness in location patterns.

Table 5: Shares of		

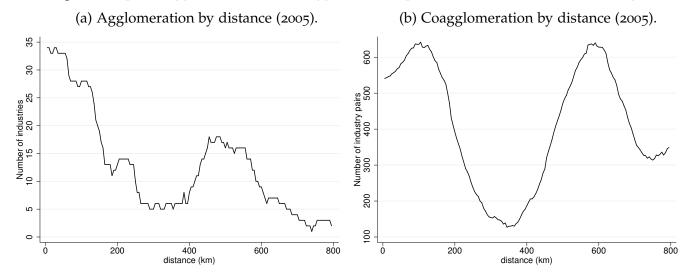
	Agglomeration: 4-digit mfg industries			Coagglomeration: 4-digit mfg industry pairs		
Year	% agglomerated	% dispersed	% random	% agglomerated	% dispersed	% random
2001	56.47	9.41	34.12	60.67	9.55	29.78
2003	47.06	12.94	40.00	56.67	10.81	32.52
2005	52.94	9.41	37.65	51.26	12.35	36.39
2007	50.59	12.94	36.47	46.72	13.53	39.75
2009	49.41	12.94	37.65	43.08	13.47	43.45
2011	48.24	15.29	36.47	39.89	14.76	45.35
2013	37.65	14.12	48.24	31.46	14.93	53.61

Notes: Results for 85 4-digit NAICs industries, and $(85 \times 84)/2 = 3,570$ unique 4-digit industry pairs. The significance tests are based on counterfactual distributions using 200 random permutations and global confidence bands (see <u>Duranton and Overman, 2005</u>).

Figure 5 depicts the number of significantly agglomerated industries (left panel) or coagglomerated industry pairs (right panel) by distance. The left panel shows that the bulk of industries

are agglomerated at fairly short distances, below 100 kilometres. The right panel shows that a substantial share of industry pairs are coagglomerated at longer distances, with a first peak at about 100 kilometres and a second peak at about 600 kilometres. The second peak is consistent with agglomeration of industries in *different* metro areas, which are specialized in different industry mixes. The first peak is consistent with coagglomeration at intermediate distances, possibly to exploit input-output links that operate at longer distances than knowledge spillovers or labor market related aspects (see the left panel).²⁶ Note that, as Figure 2 shows, the distribution of shipments in Canada peaks at about 100 kilometres, which corresponds fairly precisely to the distance at which a large share of industry pairs are coagglomerated. While there is a slight inflection in the density of shipments at around 500-600 kilometres, Figure 2 shows that shipments at that distance are far less frequent that at 100 kilometres. This suggests that the coagglomeration at around 600 kilometres may not mainly arise due to input-output links, because it is already corresponds to a long distance to ship.

Figure 5: Spatial agglomeration and coagglomeration profiles in Canadian manufacturing.



We can gain some first insights on the importance of input-output links for the coagglomeration of industries by running regressions in the spirit of Ellison et al. (2010), Faggio et al. (2015), and Behrens and Guillain (2017). To this end, we regress the coagglomeration measures for each industry pair ij on measures of the strength of input-output links, as well as on other covariates that are traditionally associated with the geographic concentration of industries (the 'Marshallian' covariates that proxy for labor market pooling and knowledge spillovers). We

²⁶The second peak is unlikely to contain information that is useful to identify the causal mechanisms driving agglomeration and coagglomeration. Those mechanisms operate at shorter distances, so that our focus will be on short-distance patterns (e.g., 100 kilometres for input-output links) in what follows.

provide a more detailed description of our data, as well as the construction of the different variables, in Appendix A.

Table 6: Coagglomeration of industries and 'Marshallian covariates'.

	(1)	(2)	(3)	(4)
Input-output share (all industries)	0.039 ^a			
	(0.007)			
Input-output share (all manuf.)		0.045^{a}		
		(0.006)		
Input-output share (manuf., excl. self)			0.038^{a}	0.035^{a}
			(0.006)	(0.006)
Occupational labor similarity	0.049^{a}	0.044^{a}	0.048^a	0.054^{a}
	(0.005)	(0.005)	(0.005)	(0.005)
Patent citation patterns	0.008^{b}	0.007^{c}	0.006	0.006
	(0.004)	(0.004)	(0.004)	(0.004)
Year fixed effects	Yes	Yes	Yes	No
Industry fixed effects	Yes	Yes	Yes	No
Industry-year fixed effects	No	No	No	Yes
Observations	17,705	17,705	17,705	17,705
R-squared	0.842	0.843	0.842	0.875

Notes: Estimation for 2001–2009, two year steps. The dependent variable is the K-density CDF at 100 kilometres. We have 3,570 unique industry pairs per year. The fixed effects are constructed as in Ellison et al. (2010) and Behrens and Guillain (2017). 'Input-output share (all industries)' is the input-output share computed using the full input-output tables. 'Input-output share (all manuf.)' uses the manufacturing input-output sub-table only, rescaled to sum to one. Finally, 'Input-output share (manuf., excl. self)' rescales those latter shares by excluding the own-industry consumption. The construction of 'Occupational labor similarity' and 'Patent citations' are detailed in the appendix. Huber-White robust standard errors are in parentheses. Coefficients significant at: a 1%; b 5%; and c 10%.

Table 6 summarizes our estimation results. As one can see, stronger input-output links are statistically associated with more coagglomeration: industries with stronger vertical links tend to locate closer together. While this finding is suggestive of the role of transport costs, it is not a proper test. Indeed, as recently noted by Combes and Gobillon (2015, p.336, our emphasis):

"[the] strand of literature [following Ellison et al. (2010)] is an interesting effort to identify the mechanisms underlying agglomeration economies. Ultimately though, it is very difficult to give a clear interpretation of the results, and the conclusions are mostly descriptive [...] For instance, according to theory, two industries sharing inputs have more incentive to colocate when trade costs for these inputs are large. In that perspective, variables capturing input-output linkages should be caused to interact with a measure of trade costs, but this is not done in the literature."

We return to this important point below where we run a proper test that: (i) interacts transport costs with input-output coefficients; and (ii) addresses the potential endogeneity of transport

costs, input-output coefficients, and their interaction.

The estimations in Table 6, which follow the extant literature, have a number of shortcomings that we will address later. In particular, as mentioned before, we will incorporate transport costs into the analysis. We will also use a wider range of controls, exploit the within variation of our panel data, and carefully instrument various covariates that we suspect of endogeneity. Previewing our key results, transport costs and input-output links are important for explaining the coagglomeration patterns of industries. How industries relate to each other in space depends, among others, on how strongly they are linked by vertical relationships and on how costly it is to ship their outputs and source their inputs. We think that this result is novel to the literature and shows that, despite their historically low levels, transport costs — and changes therein — still matter greatly for the spatial structure of the economy.

5 Transport costs and geographic concentration

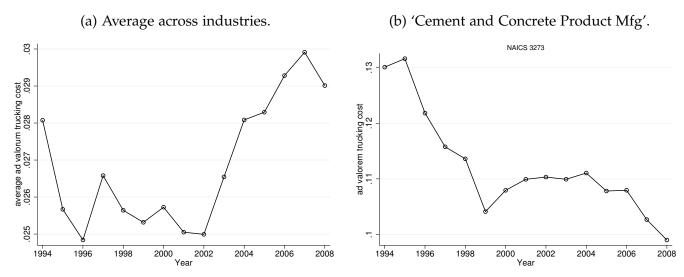
5.1 Industry-level transport costs

Our measures of geographic concentration are estimated at the industry level, whereas the ad valorem transport costs are estimated at the commodity level. We hence first need to aggregate our commodity-level transport costs to the industry level. This is accomplished by using an industry-commodity concordance to transform commodity-level transport costs into industry-level estimates, where domestic trade by commodity is used to weight transport costs in instances where multiple commodities match to an industry.²⁷

After aggregating the transport costs to the industry level, we find that over the 1994 to 2008 period the (simple) average 4-digit industry ad valorem transport cost in Canada was 2.7%. Panel (a) of Figure 6 shows that ad valorem transport costs change substantially over time. They first fell from 1994 to 2002 — due to decreasing labor costs and constant fuel prices — and then sharply rose until 2008 due to essentially increasing fuel and commodity prices. Since our study period for the geographic concentration is from 2001 to 2009, we focus in what follows on an episode of increasing transport costs. The maximum (across all years pooled) is 13.5% for 'Lime and Gypsum Product Manufacturing' (NAICS 3274) and 11.2% for 'Cement and Concrete Product Manufacturing' (NAICS 3273). The temporal evolution of transport costs in the latter sector is depicted in panel (b) of Figure 6. The minimum (across all years pooled) is 0.14% for 'Computer and Peripheral Equipment Manufacturing' (NAICS 3341) and 0.08% for 'Communications Equipment Manufacturing' (NAICS 3342). The largest time variation between

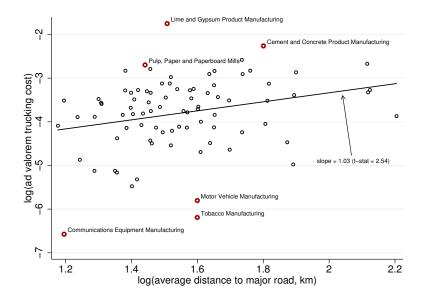
²⁷Namely, we us нs codes to bridge 4-digit NAICs industries and 5-digit Standard Classification of Transported Goods codes used to classify trucking shipments.

Figure 6: Changes in industry-level ad valorem transport costs over time.



1994 to 2008 is for 'Petroleum and Coal Product Manufacturing' (NAICS 3241), which has seen a steep fall in its ad valorem trucking rates from about 20% down to 3% during the ramp up to the 2008 spike in oil prices. Overall, these examples reveal that there is substantial between-industry variation and time-series variation. This suggests that working with the within-dimension of the panel is both important and feasible to achieve identification.

Figure 7: Industry average distance to major roads and AVTC in 2001.



Another reason for using the within dimension can be seen from Figure 7, which plots the average distance of plants in the different industries from major roads against our industry-level ad valorem transport costs. As the figure shows, there is a positive correlation between

average distance from major roads and industry-level trucking costs. In words, industries where plants are on average further away from major roads are also industries that seem to pay more to ship their goods.²⁸ As explained in Section 2.2, our commodity-level ad valorem rates are purged from carrier, commodity, and distance-shipped effects. However, they are not purged from industry effects, as embodied, e.g., in the spatial structure of industries. The correlation in Figure 7 can be due to the fact that plants with different characteristics locate in different areas and face different transport costs (recall from Section 2.1 that transport costs embody plant productivity, especially at the ad valorem level). If, for example, more productive and larger plants locate further away from major roads, and if those plants have lower ad valorem transport costs, our industry-level measures will pick this up. Furthermore, it is unclear how aggregating from the commodity to the industry level introduces additional 'industry effects' into the ad valorem rates. For all these reasons, we will include systematically industry — or industry-pair — fixed effects in our regressions to control for these effects.

Table 7: Purging ad valorem transport costs.

	(1)	(2)	(3)	(4)
Average distance to major roads	0.848^{a}	0.861^a	-0.118	-0.115
	(0.199)	(0.206)	(0.081)	(0.080)
Average plant size (employment)	0.152		-0.110	
	(0.102)		(0.128)	
Industry share of exporting plants	-0.804^{a}	-0.792^{a}	0.448^{b}	0.354
	(0.111)	(0.124)	(0.211)	(0.235)
Industry employment		0.094		-0.108
		(0.114)		(0.153)
Industry number of plants		-0.637		-0.113
		(0.627)		(1.216)
Year fixed effects	Yes	Yes	Yes	Yes
Industry fixed effects	No	No	Yes	Yes
Observations	420	420	420	420
R-squared	0.169	0.167	0.963	0.963

Notes: The dependent variable is $\ln(\tau_{fb}-1)$. All other variables in logs too. The different explanatory variables are constructed as industry averages or industry totals using our plant-level dataset. We compute for each manufacturing plant the distance to the nearest major road using GIs and the map of major North American roads from the U.S. Geological Survey. We then average those plant-level distances by 4-digit industry. Huber-White robust standard errors are in parentheses. Coefficients significant at: a 1%; b 5%; and c 10%.

Table 7 shows regressions of our industry-level ad valorem transport costs on different industry characteristics. While some characteristics are statistically significantly associated with

 $^{^{28}}$ The R^2 of the univariate regression is fairly low, just short of 0.07. The pooled regression for the years 2001, 2003, 2005, 2007, and 2009, with year fixed effects, yields a coefficient of 0.95 (T-stat of 4.71) and exactly the same R^2 as the 2001 regression. None of the year dummies are significant. The result holds with clustering by year.

transport costs in the cross-sectional dimension, the effects tend to go away or become substantially weaker once industry fixed effects are included. Hence, with those fixed effects, our ad valorem transport costs no longer correlate strongly with the different industry characteristics. This again vindicates the inclusion of fixed effects and the use of the within dimension of the data for estimation.

5.2 Estimating the effects of transport costs on geographic concentration

There is a large body of theoretical literature that focuses on the impacts of transport costs on the spatial organization of the economy (see Fujita, Krugman, and Venables, 1999; Combes, Mayer, and Thisse, 2008; and Fujita and Thisse, 2013 for reviews). The leitmotif of that literature is that transport costs matter for the spatial structure of economic activity. However, the precise impacts of changing transport costs — whether it is agglomerative or dispersive — on geographic concentration is theoretically unclear and depends on the specificities of the models used. The ambiguity in results stems from the fact that transport costs are both an agglomeration and a dispersion force, depending on the spatial structure of the economy (e.g., Krugman, 1991). Serving a geographically dispersed demand makes higher transport costs a priori a dispersion force. However, when that 'dispersed' demand becomes less dispersed, higher transport costs become an agglomeration force since firms want to be close to the large markets. The relative strength of transport costs as an agglomeration or as a dispersion force further depends on the additional agglomeration or dispersion forces that are included in the models. For example, when urban congestion impedes geographic concentration, lower transport costs are dispersive since firms want to relax congestion when shipping becomes cheap enough (e.g., Helpman, 1998; Behrens et al., 2017).

The impact of transport costs on geographic concentration is hence an empirical question. Yet, surprisingly little is known empirically. Many studies use either changes in infrastructure or changes in 'market access' to investigate the spatial impact of changes in 'transport costs' (see Redding and Turner, 2015, for a recent survey). Redding and Sturm (2008) and Brülhart, Carrère, and Trionfetti (2012) use the collapse of the ex Soviet block to look at how border regions were differentially affected by it. While these papers have a clear exogenous source of variation and thus provide valuable insights as to the causal effects of changes in market access, none of them answers the question of whether transport costs are agglomerative or dispersive. The evidence they provide is at best indirect, and it does not relate to transport costs or the geographic concentration of individual industries per se.

A full model of the geographic concentration of manufacturing should have the following ingredients: (i) a geographically dispersed final and intermediate demand that tends to disperse industries in the presence of transport costs; (ii) urban costs that limit the concentration

of industries into a few locations, especially when transport costs are low so that markets can be served from anywhere; and (iii) distance-sensitive agglomeration economies (e.g., due to knowledge spillovers or local labor pools) which tend to pull industries together, especially when transport costs are low so that the dispersed demand can be served at low costs from any location. With those three ingredients, lower transport costs will tend to geographically concentrate industries if (i) and (iii) dominate (ii), whereas the reverse holds otherwise. As we said before, this is in the end an empirical question and the answer is likely to differ across industries depending on the relative strength of the three ingredients.

5.2.1 Changes in agglomeration patterns

We first look at the effects of transport costs on the geographic concentration of single industries. We follow Behrens, Bougna, and Brown (2015) and construct a panel of Duranton-Overman K-density CDFs. There are two differences between their approach and ours. First, while Behrens et al. (2015) work at the 6-digit level, we work here at a more aggregate 4-digit level. Second, we work with private Business Register data from 2001 to 2009, in two year steps, instead of with the confidential Annual Survey of Manufacturers Longitudinal Microdata File from 1993 to 2008. This implies that we replicate the results of Behrens et al. (2015) using a different dataset, a different time period, and a different level of industrial aggregation. While this is a plus and allows us to check the robustness of the results, the downside is that we have about ten times fewer observations, which makes the estimations necessarily less precise. However, the qualitative results hold true and — as we will see later — also hold robustly in the coagglomeration analysis.

We run regressions of the following form:

$$\gamma_{i,t}(d) = (\tau_{i,t} - 1)\beta_T + \mathbf{X}_{i,t}\beta_X + \alpha_t + \mu_i + \varepsilon_{i,t}, \tag{12}$$

where $\gamma_{i,t}(d)$ is the K-density CDF for industry i in year t at distance d; $\tau_{i,t}-1$ 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 (measures of international trade exposure and input-output distances of industries); α_t and μ_i are year and industry fixed effects, respectively; and $\varepsilon_{i,t}$ is the error term. The latter is assumed to be independently and identically distributed. Our main coefficient of interest is β_T , which captures the impact of changes in transport costs on changes in the geographic concentration of industries. In what follows, we present both OLS and IV estimates of equation (12), with robust standard errors.

Endogeneity concerns and IV strategies. As mentioned before, transport costs both influence the geographic structure of economic activity and are partly determined by that same structure.

Hence, to estimate the causal effect of transport costs on geographic concentration requires us to deal with several potential endogeneity problems. Those problems are related to M and m_b/m_f in our conceptual framework of Section 2.

Consider first the backhaul problem. If markets are symmetric — i.e., of equal size M=1 and productivity $m_f=m_b$ — then $t_{fb}^c=t_{bf}^c=c(\mathbf{Y},d)$. In that case, freight rates are symmetric and equal to the marginal cost of the carriers on each leg of the trip. Assume next that M>1, so that $M^{-1/\sigma_i}<1$: market f is larger than b, so there is more demand for transport services from f to b than the converse. Assume further that $m_f=m_b$. Then it is easily seen from (4) that $t_{fb}^c>c(\mathbf{Y},d)>t_{bf}^c$. Freight rates are larger on the fronthaul trip than on the backhaul trip, because of imbalances in shipping patterns due to market size. These imbalances — due to the spatial structure of economic activity — 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 model of agglomeration). Hence, the geographic concentration of industries — by decreasing M — will tend to directly affect transport costs via the carriers' backhaul problem. This reverse causality channel needs to be controlled for appropriately.

Several empirical contributions have substantiated the potential importance of the backhaul problem (see also Appendix B for estimates using our data).³¹ Jonkeren et al. (2011) show that there are trade imbalances across the northwestern European inland waterways, and they find that these imbalances have a substantial causal effect on transport prices for inland water shipping, raising it in the direction of excess demand. Closer to our study, Tanaka and Tsubota (2016, abstract) show for the Japanese trucking industry that "a 10% increase in the front-haul transport flow relative to back-haul transport flow leads to a 1.3% decrease in the front-haul freight rate relative to back-haul freight rate." This is due to the fact that, for Japan, the negative effect of backhaul problems seems to be dominated by the productivity gains stemming from density economies in transportation. Whatever the direction of the bias, we need to control for the fact that transport costs directly depend on the spatial structure of economic activity.

²⁹The same holds true when $m_f < m_b$ — firms in f are more productive than firms in b — and if M=1. In that case, we have again $t_{fb}^c > c(\mathbf{Y},d) > t_{bf}^c$: freight rates are larger on the fronthaul trip than on the backhaul trip, because of imbalances in shipping patterns due to productivity differences.

 $^{^{30}}$ One 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.

³¹Wicksell and Taussig (1918, p.407) already pointed out that "the increased number of ships going from England to America with full load, and bound to go back in ballast or with insufficient cargo, must increase the transport charges on goods going one way and diminish the cost of sending goods the other way." See Clark, Dollar, and Micco (2004) and Blonigen and Wilson (2008) for the problem of imbalances in an international trade context. They show that directional trade imbalances between the U.S. and their trading partners have significant effects on freight rates.

The second major endogeneity concern comes from the fact that the geographic distribution of economic activity affects both the productivity of the shippers and of the carriers. As stated above, if $m_f < m_b$ in our conceptual framework — firms in f are more productive than firms in b — and assuming that M=1, we have $t_{fb}^c>c(\mathbf{Y},d)>t_{bf}^c$: freight rates are larger on the fronthaul trip than on the backhaul trip, because of imbalances in shipping patterns due to productivity differences. Regional productivity differences are pervasive and may stem from three major sources: (i) more productive firms sort into specific locations (e.g., large cities); (ii) some regions make firms more productive by offering specific locational advantages (e.g., natural resources or infrastructure); and (iii) some regions make firms more productive via agglomeration economies or selection effects. Whatever the reason, the literature on agglomeration economies has substantiated a significant and causal effect between regional size, M_f , and productivity, $1/m_f$.³² The geographic concentration of industries — by increasing productivity due to all sorts of agglomeration economies — will tend to directly affect transport costs. It affects transport costs by potentially decreasing carriers' costs through densitiv economies, which are agglomeration economies for carriers, and it affects transport costs by increasing demand due to productivity gains. When working with ad valorem rates (4), there is also an additional effect: increased productivity of manufacturers due to agglomeration leads to lower prices (because of productivity gains), which mechanically tends to increase ad valorem rates.

Since transport costs are endogenous to the spatial structure of economic activity — through backhaul problems, density economies in transportation, and agglomeration economies in manufacturing — we need to find instruments to deal with the source of endogeneity. In what follows, we will use the following empirical strategies which build on Behrens, Bougna, and Brown (2015) and Ellison et al. (2010).

First, we regress our measure of ad valorem transport costs on industrial multi-factor productivity (MFP) indices to purge them from productivity effects. We then use the residual from those regressions as our explanatory variable. We refer to these as the ad valorem transport cost residuals (AVTCR), and they are by construction orthogonal to MFP.

Second, we instrument the industry price indices in (8) using the corresponding U.S. versions. The idea underlying that strategy is the following. As explained before, the spatial distribution of an industry has direct effects on the ad valorem rates by affecting the industry's productivity (through agglomeration economies) and by affecting the carriers' productivity (through backhaul problems and density economies). Both of these effects influence the price indices in (8). To get rid of these effects, we construct ad valorem transport cost instruments based on (8), where the U.S. price indices are used. This provides valid instruments if the

³²See Rosenthal and Strange (2004) and Combes and Gobillon (2015) for the empirics; and Duranton and Puga (2004) and Behrens and Robert-Nicoud (2015) for the theory.

geographic structure of the Canadian industry has no impact on the U.S. price index of that industry. Since Canada is ten times smaller than the U.S., it is unlikely that changes in the geographic distribution of Canadian industries drive changes in the U.S. industry price indices.

Third, the input-output coefficients in Table 6 are also likely to be endogenous. Industries may end up colocating for reasons unrelated to agglomeration economies — e.g., because they use the same natural resources — and they thus end up buying from each other because they are close and this allows to save costs via input substitution. To deal with this problem, we follow Ellison et al. (2010) and instrument this variable using its U.S. counterpart from the input-output benchmark tables. The idea is the same as for the instrumentation of the price indices. For the instrument to be valid, we need that industries that spuriously colocate in the U.S. and start buying from each other do so for different reasons than in Canada. We further need cross-border trade in those industries to be not so large as to make the geographic structure of the Canadian industry have a direct effect on the U.S. input-output structure. The former seems fairly unlikely — though we cannot exclude it — while we control for the latter by excluding industries that are strongly linked by cross-border trade.

Results. Table 8 summarizes our results.³³ Observe first that the transport cost variable is negative and highly significant in the pooled estimates of columns (1) and (2). The specifications where we use the MFP purged transport cost residual give slightly more negative estimates, but the difference is not significant. Overall, the cross-sectional results suggest that high transport cost industries are more dispersed, or increases in transport costs tend to disperse industries. We cannot dissociate these two explanations from the pooled regressions. In columns (3) and (4), we use U.S. instruments for the ad valorem rates to replicate columns (1) and (2). Both when instrumenting the ad valorem transport cost, or the MFP purge residual ad valorem rate, the results are basically unchanged, with slightly more negative coefficients.³⁴ Note that the instrument is strong in the cross section.

When turning to the within estimates in columns (5)–(10), one can see that we lose precision. The ad valorem rates are not significant, no matter whether or not we purge them from productivity effects, or we instrument them, or we add additional controls.³⁵ As mentioned

³³We do not include industry controls, contrary to Behrens et al. (2015). The reason is that these controls are largely insignificant and have only little explanatory power.

³⁴We cannot formally test the equality of IV and OLS estimates since the models cannot be nested. However, there is substantial overlap of the confidence intervals and the IV estimates are contained in the OLS confidence intervals, thus suggesting that endogeneity is not a major concern in columns (1) and (2) of Table 8.

³⁵Table 7 shows that industry fixed effects explain the largest chunk of the variation in transport costs. We thus do not have enough variation in Table 8 to precisely estimate the effect of transport costs in a small panel of 5 years once we include industry fixed effects. Note also that theory does not tell us what coefficient estimate we should expect to find.

33

Table 8: Transport costs and the agglomeration of industries.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
			U.S. IV	U.S. IV			U.S. IV	U.S. IV		
Ad valorem trucking rate	-0.318 ^a		-0.331 ^a		0.032		0.005			
	(0.042)		(0.042)		(0.028)		(0.026)			
Ad valorem trucking rate residual		-0.325^a		-0.330 ^a		-0.014		0.005	0.001	0.002
		(0.043)		(0.042)		(0.026)		(0.026)	(0.016)	(0.020)
Average input-output distance									-0.888^a	-0.882^{a}
									(0.125)	(0.116)
NAFTA share of imports										-0.045
										(0.143)
OECD share of imports										-0.029
										(0.089)
Asian share of imports										-0.087
										(0.130)
NAFTA share of exports										-0.051
										(0.120)
OECD share of exports										-0.090
										(0.089)
Asian share of exports										-0.005
										(0.044)
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry fixed effects	No	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Observations	420	420	420	420	420	420	420	420	420	420
R-squared	0.103	0.108	0.103	0.108	0.961	0.961	0.961	0.961	0.978	0.978

Notes: The dependent variable is the Duranton-Overman K-density CDF at 100 kilometres distance. All variables are standardized so that the coefficients measure effect sizes. See the appendix for a detailed description of the variables. 'Average input-output distance' captures the industry-average distance of plants to their suppliers and clients. See Behrens, Bougna, and Brown (2015) for details. Huber-White robust standard errors are in parentheses. Coefficients significant at: a 1%; b 5%; and c 10%.

before, compared to Behrens et al. (2015), we have about ten times fewer observations, so that we just do not have enough data to estimate these relationships precisely. Note, however, that the input-output distance variable is negative and precisely estimated, as in Behrens et al. (2015), and that the import shares are negative (though insignificant). Both suggest that industries where suppliers or clients tend to disperse are experiencing dispersion too, and that increased import competition also has a dispersive effect. We will show in the next section that all these results can be precisely estimated using the coagglomeration patterns of industries.

5.2.2 Changes in coagglomeration patterns

To more precisely estimate the effects of transport costs on the geographic concentration of industries, we now turn to the coagglomeration patterns between industry pairs. Using these patterns has two significant advantages. First, we can make use of the input-output coefficients between industries to obtain a precise measure of how interdependent the industries are in terms of buyer-supplier relationships. Although the inclusion of the 'input-output distance measure' in Table 8 intuitively captures how far an industry is from its potential suppliers or clients, that measure is indirect at best. Second, we have many more observations to work with. By looking at the coagglomeration patterns of our 85 4-digit industries, we increase the number of cross-sectional observations from 85 to 3,570. As we will see, this allows us to precisely measure the relationships that we are interested in. Let us emphasize that, as far as we know, this material on changes in coagglomeration patterns taking into account (i) transport costs, (ii) input-output links, and (iii) the interaction between the two, is new to the literature. It provides a theory-based test for the importance of transport costs for industry location, as suggested recently by Combes and Gobillon (2015).

We now run regressions of the following form:

$$\gamma_{ij,t}(d) = (\tau_{ij,t} - 1)\beta_T + \mathbf{X}_{ij,t}\beta_X + \mu_{(i,j,t)} + \varepsilon_{ij,t}, \tag{13}$$

where $\gamma_{ij,t}(d)$ is the coagglomeration K-density CDF for industries i and j in year t at distance d; $\tau_{ij,t}-1$ is the average of our measures of ad valorem transport costs (8) of industries i and j in year t, i.e., $\tau_{ij,t}-1=(\tau_{i,t}+\tau_{j,t}-2)/2$; $\mathbf{X}_{ij,t}$ is a vector of time-varying industry-pair controls (including measures of international trade exposure and input sourcing patterns from primary and from service industries); $\mu_{(i,j,t)}$ are different combinations of industry, year, industry-year, and industry-pair fixed effects, respectively; and $\varepsilon_{ij,t}$ is the error term. The latter is assumed to be independently and identically distributed. Our main coefficient of interest is again β_T , which captures the impact of changes in transport costs on changes in the geographic concentration of industry pairs. If $\beta_T < 0$, this means that an increase in the average transport cost faced by these two industries tends to pull those industries further apart. Such industries

hence tend to coagglomerate — maybe to exploit agglomeration economies that are unrelated to transportation — as shipping their goods becomes cheaper. Our dependent variable is the CDF of the Duranton-Overman coagglomeration measures at 100 kilometres. Our choice of distance is motivated by Figures 2 and 5, which show that shipping and coagglomeration patterns peak at that distance. In what follows, we present both old and IV estimates of equation (13), with robust standard errors.

Table 9 runs our baseline specifications with various sets of fixed effects, and without any controls. As shown, except for specifications (2) and (6), all coefficients on transport costs are negative and highly significant. Specification (2) may suffer from an endogeneity bias, and the estimations in columns (3) and (4) — which use either the MFP purged ad valorem residual or the U.S. instrument — seem to confirm this hypothesis. Once productivity effects and potential endogeneity in prices are purged, the coefficient β_T is precisely estimated and significantly negative. This suggests that backhaul problems and agglomeration economies tend to bias the estimated coefficient on ad valorem transport costs towards zero, as suggested by the theory.³⁶ In our preferred specifications (8) to (10), where we include industry-pair fixed effects to focus on the within dimension, and where we either instrument or use the MFP residual or both, we consistently find elasticities around -0.05 to -0.12 which are precisely estimated.

We next introduce the input-output coefficient, as well as its interaction with transport costs. Table 10 summarizes our key regressions, where we interact the transport cost variable with the strength of input-output links between the industries. All regressions use the within dimension of our data by including industry-pair fixed effects. These control for a wide range of time-invariant factors that may drive the coagglomeration of industries but are unrelated to transport costs.³⁷ Most importantly, including these pairwise fixed effects controls for natural advantages and infrastructure, which are both hard to control for in the between dimension and which are hard to measure exhaustively (see Ellison and Glaeser, 1999; Ellison et al., 2010). As these characteristics vary slowly over time, we can be confident that they are soaked up by our fixed effects over the short 2001–2009 period we consider.

As one can see from columns (1) and (2) in Table 10, the coefficient on transport costs barely changes, whereas the coefficient on the interaction with the input-output coefficient is positive and significant. Since we are in the within dimension, the coefficient on the input-output coefficient itself becomes less precisely estimated and loses significance (note that input-output tables are slow changing). The positive coefficient on the interaction term means that *industries*

³⁶Contrary to Tanaka and Tsubota (2016), backhaul problems seem to dominate density economies in Canada. This may be explained by the radically different spatial structure of Canada, where average shipping distances are much longer than in Japan.

³⁷Note that the patent citations data we use in Table 6 have no time dimension, and so it cannot be separately identified when using industry-pair fixed effects.

36

Table 9: Transport costs and the coagglomeration of industries.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
			мғг purged	U.S. IV	мғг purged					мғр purged
					+ U.S. IV					+ U.S. IV
Ad valorem trucking rate	-0.247 ^a	-0.012		-0.047 ^a		-0.004	-0.012 ^b		-0.125 ^a	
	(0.008)	(0.011)		(0.012)		(0.015)	(0.006)		(0.008)	
Ad valorem trucking rate residual			-0.069 ^a		-0.047 ^a			-0.069 ^a		-0.050 ^a
			(0.011)		(0.012)			(0.005)		(0.007)
Year fixed effects	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes	No	Yes
Industry fixed effects	No	Yes	Yes	Yes	Yes	No	No	No	No	No
Industry-year fixed effects	No	No	No	No	No	Yes	No	No	Yes	No
Industry-pair fixed effects	No	No	No	No	No	No	Yes	Yes	Yes	Yes
Observations	17,430	17,430	17,430	17,430	17,430	17,430	17,430	17,430	17,430	
R-squared	0.064	0.838	0.838	_	_	0.870	0.952	0.952	0.984	_

Notes: The dependent variable is the Duranton-Overman K-density CDF at 100 kilometres distance. All variables are standardized so that the coefficients measure effect sizes. See the appendix for a detailed description of the variables. Huber-White robust standard errors are in parentheses. Coefficients significant at: a 1%; b 5%; and c 10%.

Table 10: Transport costs, input-output links, and the coagglomeration of industries.

	(1)	(2)	(3) Share controls	(4) Excl. NAICS3	(5) Excl. naics3
Ad valorem trucking cost residual (AVTCR)	-0.073 ^a	-0.076 ^a	-0.083 ^a	-0.073 ^a	-0.080 ^a
()	(0.005)	(0.006)	(0.005)	(0.006)	(0.005)
Input-output share (all industries)	0.027	(0.000)	(0.00)	(0.000)	(0.003)
,	(0.014)				
Input-output share (all industries) \times AVTCR	0.033^{b}				
	(0.015)				
Input-output share (manuf., excl. self)	, ,,,	0.009	0.013	0.043 ^a	0.050^{a}
		(0.011)	(0.011)	(0.017)	(0.016)
Input-output share (manuf., excl. self) \times AVTCR		0.032^{a}	0.029^{a}	0.029^{b}	0.023^{c}
		(0.010)	(0.010)	(0.014)	(0.012)
Occupational labor similarity	-0.049 ^a	-0.048^a	-0.058 ^a	-0.054 ^a	-0.064 ^a
	(0.014)	(0.014)	(0.014)	(0.014)	(0.014)
Input-output share primary industries			-0.141 ^a		-0.144 ^a
			(0.013)		(0.013)
Input-output share service industries			0.037^{a}		0.035^{a}
			(0.008)		(0.008)
Year fixed effects	Yes	Yes	Yes	Yes	Yes
Industry fixed effects	No	No	No	No	No
Industry-year fixed effects	No	No	No	No	No
Industry-pair fixed effects	Yes	Yes	Yes	Yes	Yes
Observations	17,430	17,430	17,430	16,495	16,495
R-squared	0.952	0.952	0.953	0.951	0.952

Notes: The dependent variable is the Duranton-Overman K-density CDF at 100 kilometres distance. All variables are standardized so that the coefficients measure effect sizes. See the appendix for a detailed description of the variables. In columns (4) and (5), we exclude all 4-digit industry pairs that belong to the same 3-digit industry. Huber-White robust standard errors are in parentheses. Coefficients significant at: a 1%; b 5%; and c 10%.

that are more strongly linked by input-output relationships tend to disperse less as transport costs increase than industries that are less strongly linked. This effect is depicted in panel (i) of Figure 8 below, which graphs the total effect of transport costs on geographic concentration at the 10th percentile of the input-output coefficient distribution (left panel) and at the 90th percentile of that distribution (right panel). As can be seen, the effect is negative and significant at the 10th percentile, whereas it is insignificant at the 90th percentile. Examples of industry pairs in the former include 'Animal Food Manufacturing' (NAICS 3111) and 'Grain and Oilseed Milling' (NAICS 3112); or 'Motor Vehicle Manufacturing' (NAICS 3361) and 'Motor Vehicle Parts Manufacturing' (NAICS 3363). These strongly linked industry pairs seem to not move away from each other as transport costs change. Column (3) shows that these results are robust to the inclusion of controls for the share of business services or primary inputs sourced by the industries.³⁸ Columns (4) and (5) replicate our previous regressions by excluding all industry pairs

³⁸As expected, industry pairs that source more primary inputs are generally more dispersed as those inputs might be dispersed and differ between industries. Also, industry pairs that source more business services are more coagglomerated as business services are themselves highly spatially concentrated.

ij where both *i* and *j* are in the same 3-digit industry. The reason for running this robustness check is that some of our data (e.g., the *L*-level input-output tables) are only available at a level of aggregation between the 3- and 4-digit levels. More details are give in Appendix A. As one can see, although the coefficients drop slightly in magnitude and are a bit less precisely estimated, our main results are robust.

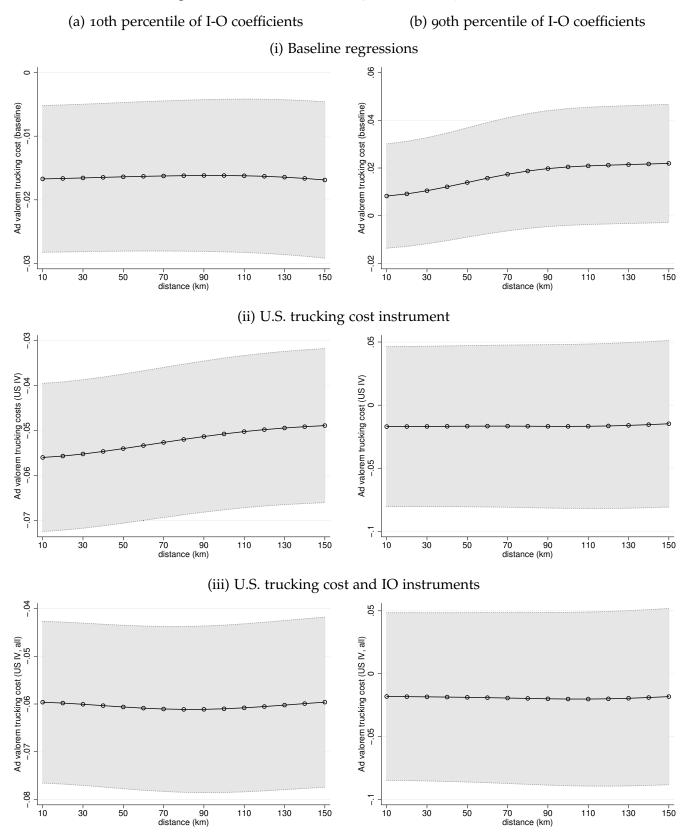
Since all of our variables are standardized, we can interpret the magnitude of their estimated coefficients as effect sizes. As one can see, the effect size of transport costs is large: it is almost as large as the sum of the (absolute value of the) effect sizes of the labor market variable and the input-output coefficient. In other words, transport costs have a first-order effect on the geographic concentration of industry pairs. This suggests that transport costs still matter, and that any full analysis of the geographic concentration of industries should carefully consider that variable.

Table 11 replicates our key regressions using different sets of instruments. As one can see from that table, the coefficient on the interaction between input-output links and transport costs is positive and significant in specifications (1), (2), (3), and (7). In these specifications, we either use the raw ad valorem transport cost, or the MFP purged residual, or the U.S. instrument for the transport cost. The estimated coefficients on the interaction terms are between 0.033 and 0.047. As one can see from specifications (4), (5), (8), and (9), the coefficient remains positive and gets larger, but is less precisely estimated, once we instrument simultaneously the ad valorem transport cost, the input-output share, and the interaction term. There are two possible explanations for this. First, since we exploit the within dimension, the instrument may be weak because the input-output tables do not change much (even in five year using the U.S. benchark tables for 1997, 2002, and 2007 from the BEA). Second, the instrument may be invalid because the geographic concentration of some Canadian industries may have a direct effect on the U.S. input-output coefficients.³⁹ While we cannot do much about the former point, we can deal with the latter by excluding industries with large cross-border trade.⁴⁰ Columns (6) and (10) show that dropping the industry pairs where one of the two industries is in the top decile of exporters or importers with the U.S. and Mexico leads the coefficients on the interaction term to be much larger and statistically significant. This is reassuring and leads us to believe that the instruments are valid and strong. Although the magnitude of the estimated interaction term varies across specifications, the estimated coefficient on the ad valorem trucking rate is always significant and negative, with very stable values between -0.05 and -0.07 across the

³⁹Recall that imports and exports do not directly enter the input-output coefficients, which are measured using national interindustry links. Hence, if the U.S. imports a lot from Canada in some industries, this might lead to its input-output coefficients to be small.

 $^{^{40}}$ The first-stage shows that the U.S. input-output shares are strong instruments in the cross-section, but less so in the within dimension. However, even in the within dimension the first stage F-stat is above 10.

Figure 8: Ad valorem trucking cost effects by distance.



40

Table 11: Transport costs, input-output links, and the coagglomeration of industries (IV estimations).

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
							Excl. NAICS3	Excl. NAICS3	Excl. NAICS3	Excl. NAICS3
						Drop trade				Drop trade
Ad valorem trucking rate (AVTC)	-0.017 ^a		-0.053 ^a	-0.061 ^a		-0.050 ^a	-0.050 ^a	-0.073 ^a		-0.068 ^a
	(0.006)		(0.007)	(0.009)		(0.012)	(0.007)	(0.024)		(0.017)
Ad valorem trucking rate residual		-0.073 ^a			-0.060^{a}				-0.066^a	
		(0.005)			(0.009)				(0.024)	
Input-output share (all industries)	0.033^{b}	0.027^{c}	0.023^{b}	-0.139	-0.113	-0.228	0.065^{a}	-0.276	-0.102	-0.367
	(0.014)	(0.014)	(0.011)	(0.111)	(0.118)	(0.160)	(0.017)	(0.494)	(0.497)	(0.255)
Input-output share (all industries) \times AVTC	0.036^a		0.047^{a}	0.040		0.084^{a}	0.044 ^b	0.115		0.258^{a}
	(0.013)		(0.012)	(0.034)		(0.030)	(0.017)	(0.070)		(0.062)
Input-output share (all industries) \times AVTCR		0.033^{b}			0.046				0.153^{c}	
		(0.015)			(0.036)				(0.088)	
Occupational labor similarity	-0.046^a	-0.049 ^a	-0.050 ^a	-0.047 ^a	-0.046^a	-0.045 ^a	-0.057 ^a	-0.047 ^a	-0.052 ^a	-0.043 ^a
	(0.014)	(0.014)	(0.010)	(0.010)	(0.010)	(0.011)	(0.011)	(0.017)	(0.019)	(0.012)
Instruments	No	No	U.S. AVTC	U.S. all	U.S. all	U.S. all	U.S. AVTC	U.S. all	U.S. all	U.S. all
Multifactor productivity purged AVTC	No	Yes	No	No	Yes	No	No	No	Yes	No
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry-pair fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	17,430	17,430	17,430	17,430	17,430	15,105	16,495	16,495	16,495	14,299
R-squared	0.952	0.952	_			<u> </u>				

Notes: The dependent variable is the Duranton-Overman K-density CDF at 100 kilometres distance. All variables are standardized so that the coefficients measure effect sizes. See the appendix for a detailed description of the variables. In columns (7)–(10), we exclude all 4-digit industry pairs that belong to the same 3-digit industry. In columns (6) and (10), we drop all industry pairs where one of the two industries is in the top decile of exporters or importers with the U.S. and Mexico. Specifications with 'U.S. AVTC' instrument the AVTC using the U.S. counterpart. Specifications with 'U.S. all' use U.S. instruments for AVTC, for the input-output coefficient, and the interaction term. Huber-White robust standard errors are in parentheses. Coefficients significant at: a 1%; b 5%; and c 10%.

different specifications.

Panels (ii) and (iii) of Figure 8 depict the total effect of trucking costs on the geographic concentration of industry pairs at the 10th percentile of the input-output coefficient distribution (left panel) and the 90th percentile of that distribution. As can be seen, the coefficient is negative and significant at the 10th percentile, and insignificant at the 90th percentile. Furthermore, the coefficients do not change much with the distance threshold we use to measure the K-density CDF. Note also that the effects of ad valorem transport costs on geographic concentration become more negative as we move from the raw ad valorem transport cost (panel (i)) to the one instrumented using U.S. price indices only (panel (ii)), and finally to the one using both U.S. price indices and U.S. input-output coefficients to instrument all variables (panel (iii)).

Finally, as a last robustness check, Table 12 replicates a number of our foregoing regressions by including as additional covariates the industries' import and export shares with various trading partner groups (as in Table 8). Consistent with the results in Behrens et al. (2015) and Behrens, Boualam, and Martin (2017) we find that more import exposure is associated with less geographic concentration. The results for exports are the opposite.

6 Conclusions

There is an extensive literature that *indirectly* looks at the impacts of transport costs on various economic outcomes (see Redding and Turner, 2015, for a recent survey). Most of that literature uses distance (or changes in infrastructure) as a proxy for levels of (or changes in) transport costs (e.g., Chandra and Thompson, 2000; Michaels, 2008; Duranton et al., 2014). The fundamental reason for this approach is that it is quite hard to obtain reliable transport cost data, except for maritime shipping where for and CIF prices can be more readily observed from shipping manifests.⁴¹ While distance, infrastructure, and transport costs are linked, it is unclear what the strength of that link is and, moreover, how changes in infrastructure (or how distance) map into changes in (or levels of) transport costs. It is also unclear how changes in transport costs, ultimately, reshape the geographic concentration of industries.

In this chapter, we have shown how we can estimate transport costs using trucking microdata, and how we can combine them with geocoded plant-level data to investigate the impacts of transport costs on the geographic concentration of industries. On the methodological side, one key message is that transport costs are endogenous to the spatial structure of the econ-

⁴¹Contrary to what one may think, the endogeneity of transport costs — which are prices that clear markets — is not a primary concern. The reason is that infrastructure is equally endogenous, so that much of the literature has been concerned with developing credible instrumentation strategies to deal with non-random assignment into treatment (e.g., Baum-Snow, 2007; Michaels, 2008; Duranton and Turner, 2011).

Table 12: The coagglomeration of industries with geography and trade.

	(1)	(2)	(3)	(4)	(5)	(6)
					Share controls	Excl. NAICS3
Ad valorem trucking rate residual (AVTCR)	-0.078^{a}	-0.063 ^a	-0.063 ^a		-0.071 ^a	-0.067 ^a
	(0.006)	(0.006)	(0.006)		(0.006)	(0.006)
Input-output share (mfg., excl. self)	0.007	0.013	0.011	0.021^{c}	0.017	0.050^{a}
	(0.011)	(0.011)	(0.011)	(0.011)	(0.011)	(0.015)
Input-output share (mfg., excl. self) × AVTCR	0.033^{a}	0.030^{a}	0.030^{a}		0.025^{a}	0.014
	(0.010)	(0.010)	(0.010)		(0.009)	(0.012)
NAFTA import share	-0.039		-0.079 ^a	-0.113^a	-0.104 ^a	-0.105 ^a
	(0.026)		(0.026)	(0.025)	(0.026)	(0.027)
OECD import share	-0.122 ^a		-0.166 ^a	-0.183 ^a	-0.186^a	-0.185 ^a
-	(0.017)		(0.017)	(0.017)	(0.017)	(0.018)
Asian import share	-0.072^{a}		-0.115 ^a	-0.143 ^a	-0.133 ^a	-0.134 ^a
•	(0.022)		(0.022)	(0.021)	(0.022)	(0.023)
NAFTA export share		0.135^{a}	0.191^{a}	0.200^{a}	0.193	0.196^{a}
•		(0.020)	(0.020)	(0.020)	(0.020)	(0.021)
OECD export share		0.108^{a}	0.156^{a}	0.162^{a}	0.164^{a}	0.165^{a}
•		(0.018)	(0.018)	(0.018)	(0.018)	(0.019)
Asian export share		-0.012	0.005	0.005	0.002	0.002
•		(0.008)	(0.008)	(0.008)	(0.008)	(0.009)
Occupational labor similarity					-0.044 ^a	-0.049 ^a
,					(0.013)	(0.014)
Input-output share primary industries					-0.199^a	-0.199 ^a
					(0.013)	(0.013)
Input-output share service industries					0.016^{b}	0.014 ^c
1 1					(0.008)	(0.008)
Ad valorem trucking rate				-0.011 ^c	` ,	, ,
8				(0.006)		
Input-output share (mfg., excl. self) \times AVTC				0.019 ^c		
, , , , , , , , , , , , , , , , , , , ,				(0.011)		
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Industry fixed effects	No	No	No	No	No	No
Industry-year fixed effects	No	No	No	No	No	No
Industry-pair fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	17,430	17,430	17,430	17,430	17,430	16,495
R-squared	0.953	0.953	0.954	0.953	0.954	0.954

Notes: The dependent variable is the Duranton-Overman K-density CDF at 100 kilometres distance. All variables are standardized so that the coefficients measure effect sizes. See the appendix for a detailed description of the variables. Huber-White robust standard errors are in parentheses. Coefficients significant at: a 1%; b 5%; and c 10%.

omy, and we have proposed a number of instrumental variables strategies to deal with that problem. On the empirical side, we have shown that transport costs are key drivers of both interregional trade patterns and geographic concentration of industries. Our estimates suggest that 25% to 57% of the relationship between trade flows and distance in Canada is explained by transport costs. Furthermore, our estimates also suggest that increases in transport costs tend to disperse manufacturing industries, but that the effect is weaker the more strongly industries are connected by input-output links. At the 90th percentile of the distribution of input-output coefficients, the effect is essentially zero, whereas it is significantly negative and large at the

10th percentile. These results suggest that transport costs are a key determinant of the spatial economy, and that their magnitude is large. The claim that we live in 'brave new frictionless world' is exaggerated and does not hold up to solid empirical scrutiny. The world is not yet flat, and transport costs still matter to a large extent.

Our results are important and potentially policy relevant. There is indeed a recent literature that emphasizes that any cost-benefit analysis in transportation has to take into account how transport costs affect the geographic concentration of economic activity (see, e.g., Venables, 2007, and Kanemoto, 2013, who discuss the issue of transport policy neglecting the endogenous response of the spatial organization). Although this point is well-taken and potentially important, little is known about the possible magnitudes of the effects, so that even back-of-the-envelope calculations seem out of reach. We believe that this chapter partly fills this gap by showing how the effects can be quantified and by providing a first set of benchmark numbers. We hope that our estimates and data will be useful for future research aiming at a better understanding of how frictions for shipping goods shape the geographic landscape.

Acknowledgements. We thank the editors, Wesley Wilson and Bruce Blonigen, as well as Julien Martin, David Evans, and Larry McKeown for helpful comments; and Théophile Bougna, Afshan Dar-Brodeur, and Jesse Tweedle for their research assistance. We are grateful to Bill Kerr for sharing the patent citation data with us. This paper has been written for inclusion in the volume *Handbook of International Trade and Transportation*, edited by Bruce Blonigen and Wesley Wilson. Behrens gratefully acknowledges financial support from the CRC Program of the Social Sciences and Humanities Research Council (SSHRC) of Canada for the funding of the *Canada Research Chair in Regional Impacts of Globalization*. This study was funded by the Russian Academic Excellence Project '5-100'. Any remaining errors are ours.

References

- [1] Alchian, Armen A., and William R. Allen. 1964. *University Economics*. Belmont, Calif.: Wadsworth.
- [2] Allen, Treb. 2014. "Information frictions in trade." Econometrica 82(6): 2041–2083.
- [3] Anderson, James E., and Eric van Wincoop. 2004. "Trade costs." *Journal of Economic Literature* 42(3): 691–751.
- [4] Baum-Snow, Nathaniel. 2007. "Did highways cause suburbanization?" *Quarterly Journal of Economics* 122(2): 775–805.

- [5] Baumol, William J., and Hrishikesh D. Vinod. 1970. "An inventory theoretic model of freight transportation demand." *Management Science* 16(7), 413–421.
- [6] Behrens, Kristian, Brahim Boualam, and Julien Martin. 2017. "Are clusters resilient? Evidence from Canadian textile industries." In progress, Université du Québec à Montréal.
- [7] 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.
- [8] Behrens, Kristian, Théophile Bougna, and W. Mark Brown. 2015. "The world is not yet flat: Transport costs matter!" CEPR Discussion Paper #10356, Centre for Economic Policy Research, London, UK.
- [9] 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.
- [10] Behrens, Kristian, Giordano Mion, Yasusada Murata, and Jens Suedekum. 2017. "Spatial frictions." *Journal of Urban Economics* 97**(A)**: 40–70.
- [11] Behrens, Kristian, and Rachel Guillain. 2017. "The determinants of coagglomeration: Evidence from functional employment patterns." In progress, Université du Québec à Montréal.
- [12] Behrens, Kristian, and Pierre M. Picard. 2011. "Transportation, freight rates, and economic geography." *Journal of International Economics* 85(2): 280–291.
- [13] 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 B.V., pp. 171–245.
- [14] Bemrose, Robby, W. Mark Brown, and Jesse Tweedle. 2016. "Going the distance: Estimating the effect of provincial borders on trade when geography (and everything else) matters." In progress, Statistics Canada.
- [15] Blonigen, Bruce A., and Wesley W. Wilson. 2008. "Port efficiency and trade flows." *Review of International Economics* 16(1), 21–36.
- [16] Brown, W. Mark. 2015. "How much thicker is the Canada–US border? The cost of crossing the border by truck in the pre-and post-9/11 eras." *Research in Transportation Business & Management* 16: 50–66.

- [17] 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.
- [18] Cairncross, Frances. 2001. The Death of Distance: How the Communications Revolution Is Changing our Lives. Cambridge, MA: Harvard Business Review Press.
- [19] 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.
- [20] Clark, Ximena, David Dollar, and Alejandro Micco. 2004. "Port efficiency, maritime transport costs, and bilateral trade." *Journal of Development Economics* 75(2): 417–450.
- [21] Combes, Pierre-Philippe, Gilles Duranton, Laurent Gobillon, Diego Puga, and Sébastien Roux. 2012. "The productivity advantages of large cities: Distinguishing agglomeration from firm selection". *Econometrica* 80**(6)**, 2543–2594.
- [22] Combes, Pierre-Philippe, and Laurent Gobillon. 2015. "The empirics of agglomeration economies." In: Duranton, Gilles, J. Vernon Henderson, and William C. Strange (eds.) *Handbook of Regional and Urban Economics, vol.* 5. North-Holland: Elsevier B.V., pp. 247–348.
- [23] 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.
- [24] Combes, Pierre-Philippe, Thierry Mayer, and Jacques-François Thisse. 2008. *Economic Geography: The Integration of Regions and Nations*. Princeton Univ. Press: Princeton, NJ.
- [25] 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.
- [26] Duranton, Gilles, and Henry G. Overman. 2008. "Exploring the detailed location patterns of U.K. manufacturing industries using microgeographic data." *Journal of Regional Science* 48(1): 213–243.
- [27] Duranton, Gilles, and Henry G. Overman. 2005. "Testing for localization using microgeographic data." *Review of Economic Studies* 72(4): 1077–1106.

- [28] 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, Elsevier: North-Holland, pp. 2063–2117.
- [29] Duranton, Gilles, and Matthew A. Turner. 2012. "Urban growth and transportation." *Review of Economic Studies* 79(4): 1407–1440.
- [30] Ellison, Glenn D., and Edward L. Glaeser. 1999. "The geographic concentration of industry: Does natural advantage explain agglomeration?" *American Economic Review* 89(2): 311–316.
- [31] Ellison, Glenn D., and Edward L. Glaeser. 1997. "Geographic concentration in U.S. manufacturing industries: A dartboard approach." *Journal of Political Economy* 105(5): 889–927.
- [32] 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.
- [33] Faggio, Giulia, Olmo Silva, and William C. Strange. 2014. "Heterogeneous agglomeration." SERC Discussion Paper #152, Spatial Economic Research Center, London School of Economics, UK.
- [34] Forslid, Rickard, and Toshihiro Okubo. 2014. "Spatial relocation with heterogeneous firms and heterogeneous sectors." *Regional Science and Urban Economics* 46(2): 42–56.
- [35] Fujita, Masahisa, Paul R. Krugman, and Anthony J. Venables. 1999. *The Spatial Economy: Cities, Regions, and International Trade*. MIT Press, Cambridge, MA.
- [36] Fujita, Masahisa, and Jacques François Thisse. 2013. *Economics of Agglomeration: Cities, Industrial Location, and Globalization*, 2nd Edition. Cambridge, MA: Cambridge University Press.
- [37] Friedman, Thomas L. 2005. The World is Flat. Farrar, Straus and Giroux, New York.
- [38] Gaubert, Cécile. 2015. "Firm sorting and agglomeration." *Processed*, University of California, Berkeley.
- [39] Glaeser, Edward L., and Janet E. Kohlhase. 2004. "Cities, regions and the decline of transport costs." *Papers in Regional Science* 83(1): 197–228.

- [40] Globerman, Steven, and Paul Storer. 2010 "Geographic and Temporal Variations in Freight Costs for US Imports from Canada: Measurement and Analysis." Border Policy Research Institute, Western Washington University.
- [41] Head, Keith, and Thierry Mayer. 2013. "What separates us? Sources of resistance to globalization." *Canadian Journal of Economics* 46(4): 1196–1231.
- [42] Head, Keith, and Thierry Mayer. 2014. "Gravity Equations: Workhorse, Toolkit, and Cookbook". In: G. Gopinath, E. Helpman and K. Rogoff (eds.) *Handbook of International Economics*, vol. 4, Elsevier: North-Holland, pp. 131–195.
- [43] Helpman, Elhanan. 1998. "The size of regions." In: Pines, David, E. Sadka, I. Zilcha (eds.) *Topics in Public Economics. Theoretical and Empirical Analysis*. Cambridge University Press, pp. 33–54.
- [44] Hillberry, Russell, and David Hummels. 2008. "Trade responses to geographic frictions: A decomposition using micro-data." *European Economic Review* 52(3): 527–550.
- [45] Hummels, David. 2007. "Transportation costs and international trade in the second era of globalization." *Journal of Economic Perspectives* 21(3): 131–154.
- [46] Irarrazabal, Alfonso, Andreas Moxnes, and Luca David Opromolla. 2015. "The tip of the iceberg: a quantitative framework for estimating trade costs." *Review of Economics and Statistics* 97(4): 777–792.
- [47] Jonkeren, Olaf, Erhan Demirel, Jos van Ommeren, and Piet Rietveld. 2011. "Endogenous transport prices and trade imbalances." *Journal of Economic Geography* 11(3): 509–527.
- [48] Kanemoto, Yoshitsugu. 2013. "Second-best cost-benefit analysis in monopolistic competition models of urban agglomeration." *Journal of Urban Economics* 76: 83–92.
- [49] Kerr, William R. 2008. "Ethnic scientific communities and international technology diffusion." *Review of Economics and Statistics* 90(3): 518–537.
- [50] Krugman, Paul R. 1991. "Increasing returns and economic geography." *Journal of Political Economy* 99(3): 483–499.
- [51] Locklin, Philip D. 1972. *Economics of Transportation*. Homewood, Illinois: Richard D. Irwin Inc.
- [52] Martin, Julien. 2012. "Markups, quality, and transport costs." European Economic Review 56(4): 777–791.

- [53] 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.
- [54] Mori, Tomoya, and Koji Nishikimi. 2002. "Economies of transport density and industrial agglomeration." *Regional Science and Urban Economics* 32(2): 167–200.
- [55] 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.
- [56] Redding, Stephen J., and Matthew A. Turner. 2015. "Transportation costs and the spatial organization of economic activity." In: Duranton, Gilles, J. Vernon Henderson, and William C. Strange (eds.) *Handbook of Regional and Urban Economics, vol. 5.* North-Holland: Elsevier B.V., pp. 1339–1398.
- [57] Rosenthal, Stuart S., and William C. Strange. 2004. "Evidence on the Nature and Sources of Agglomeration Economies". In: J. Vernon Henderson, and Jacques-François Thisse (Eds.) *Handbook of Regional and Urban Economics*, vol. 1, Elsevier: North-Holland, pp. 2119–2171.
- [58] Storeygard, Adam. 2016. "Farther on down the road: transport costs, trade and urban growth." *Review of Economic Studies* 83(3): 1263–1295.
- [59] Tanaka, Kiyoyasu, and Kenmei Tsubota. 2016. "Directional imbalance of freight rates: Evidence from Japanese inter-prefectural data." *Journal of Economic Geography*, forthcoming.
- [60] Train, Kenneth, and Wesley W. Wilson. 2008. "Estimation of stated-preference experiments constructed from revealed-preference choices." *Transportation Research Part B: Methodological* 42(3): 191–203.
- [61] Venables, Anthony J. 2007. "Evaluating urban transport improvements: cost-benefit analysis in the presence of agglomeration and income taxation." *Journal of Transport Economics and Policy* 41(2): 173–188.
- [62] Wicksell, Knut, and F. W. Taussig. 1918. "International freights and prices." *Quarterly Journal of Economics* 32(2): 404–414.

Appendix

This set of appendices is structure as follows. Appendix **A** contains a detailed description of our data and variables. Appendix **B** contains additional tables and results.

Appendix A: Data

A.1. Transport costs

Trucking Commodity Origin Destination Survey. Ad valorem rates are estimated using Statistics Canada's Trucking Commodity Origin-Destination Survey (TCOD). The TCOD is a for-hire carrier-based survey that collects data on a per shipment basis, including the origin and destination, (network) distance shipped, revenue to the carrier, tonnage, and the commodity of the shipment. In order to calculate ad valorem rates, the value of the shipment is also needed. However, the TCOD does not report the value of goods shipped. Hence, value per tonne estimates by 6-digit Hs commodity from an 'experiment export trade file' produced in 2008 is used to estimate the value of the shipments. Commodity export price indices are used to project the value per tonne estimates through time (see Brown, 2015, for details). These commodity-year value per tonne estimates are used to estimate the value of shipments.

This augmented TCOD file is the basis from which two files are constructed. The first is a file derived directly from the TCOD and is used to estimate ad valorem trucking rates by industry and year as outlined in Section 2.2 and used in Section 5. This particular analysis required a long time period to improve the accuracy of the predicted 500km rates. Furthermore, these estimates are based on survey weights that ensure trucking rates are representative of the population of carriers and so is more appropriate when assessing their impact on the location of plants across the country.

The second file was specially constructed to estimate regional trade flows that are used in Section 3. The Surface Transportation File (STF) is build from the TCOD and a census of waybills from the railways. The STF, which covers the period from 2004 to 2012, uses a set of benchmark weights to ensure regional trade flows adds to known inter-provincial trade totals (see Bemrose, Brown and Tweedle, 2016) for details on the construction of the benchmark weights.)⁴² This file also provides estimates of ad valorem rates that correspond the the benchmarked trade flow estimates. Only truck-based flows and rates are used in the analysis.

⁴²Becasue the TCOD only includes for-hire carriers, it excludes private trucking. One of the benefits of benchmarking to provincial trade totals is that it addresses the underestimation of flows that may result from this.

us price indices. We use detailed year-by-year NAICS 6-digit price indices from the NBER-CES Manufacturing Productivity Databas (http://nber.org/data/nberces5809.html) to construct instruments for Canadian industry-level transportation costs.

A.2. Geography

Plant-level dataset. Our plant-level data comes from the *Scott's National All Business Directories* database, which contains information on plants operating in Canada. The database is based on the Business Register and has an extensive coverage of the manufacturing sector. Our data span the years 2001 to 2013, in two-year intervals (2001, 2003, 2005, 2007, 2009, 2011, and 2013). For every establishment, we have information on its primary 6-digit NAICS code and up to four secondary 6-digit NAICS codes; its employment; whether or not it is an exporter; and its 6-digit postal code. The Scott's database constitutes probably the best alternative to Statistics Canada's proprietary *Annual Survey of Manufacturers Longitudinal Microdata File* or the micro-level *Canadian Business Patterns*. Although the dataset is only a large sample and not the universe of manufacturing plants, it has a very wide (85%–90%) and similar coverage. It contains most of the large plants and many small plants.⁴³ Behrens and Bougna (2015, Appendix A) provide detailed information on the data quality and its representativeness — both in terms of provinces and industries — of the manufacturing portion of the database.

We consider that a plant is a manufacturer if it reports a manufacturing industry (NAICS 31–33) as its primary sector of activity. Scott's assigns primary NAICS codes based on the main line of business of the establishment. Our data span four different industrial classifications: NAICS1997, NAICS2002, NAICS2007, and NAICS2012. We have concorded those classifications to a stable set of 242 manufacturing industries. The manufacturing classification remained fairly stable over time. There is only one change at the 4-digit level, where we loose NAICS 3391 which is aggregated with other industries. We hence have a stable set of 85 4-digit industries which we use in our analysis.

We geocode all plants by latitude and longitude using their 6-digit postal code centroids obtained from Statistics Canada's Postal Code Conversion Files (PCCF). The latter associate each postal code with different Standard Geographical Classifications (sGC) used for reporting census data. We match plant-level postal code information with geographic coordinates from the PCCF, using the postal code data for the next year in order to consider the fact that there

⁴³There is no 'sampling frame' strictly speaking (though Scott's uses the *Canadian Business Patterns* – which contains the universe of entities – to contact the different establishments in a systematic way to include them into their database). There are some selection and updating biases, since establishments are contacted to sign up but are of course free to not do so. Also, small/new establishments may appear in the database with a lag only (and establishments may exit with a lag only). We do not think that these issues create sustantial biases in our analysis.

is a six months delay in the updating of postal codes. For example, the census geography of 1996 and the postal codes as of May 2002 (818,907 unique postal codes) were associated with the 2001 Scott's data.

To summarize, our manufacturing data are very similar to those of the *Annual Survey of Manufacturers* or the *Canadian Business Patterns* in terms of coverage and both province- and industry-level breakdown of plants and, therefore, provide a fairly accurate and representative picture of the overall manufacturing structure in Canada. Furthermore, since postal codes are very fine grained in Canada — especially in the more populated areas — our data are as good as geocoded.

Input-output shares. We use three-year lagged input-output matrices (1998, 2000, 2002, 2004, 2006, 2008, and 2010), which we concord to our stable set of 242 manufacturing industries (and 864 industries for the matrices comprising the whole economy). Since the finest public release of the input-output matrices is at the L-level (link level), which is between NAICS 3 and 4, we disaggregated those matrices further to the W-level (NAICS 6) using either sales or employment data as sectoral weights.⁴⁴ We use input-output tables at buyers' prices. For each manufacturing industry, i, we allocate inputs purchased or outputs sold in the L-level matrix (at the 3- or 4-digit level) to the corresponding NAICS 6-digit subsectors. To do so, we allocate the total sales of each sector to all subsectors in proportion to those sectors' sales in the total sales to obtain a 242× 242 matrix of NAICS 6-digit inputs and outputs for manufacturing only. We use these matrices to compute the shares that sectors buys from and sell to each other.

We compute three different versions of the input and output shares: (i) the share taking into account all industries (including services and primary industries); (ii) the shares for the manufacturing submatrices, rescaled to sum to unity; and (iii) the shares for the manufacturing submatrices excluding within-industry coefficients, rescaled to sum to unity. We also use these matrices to compute for each industry: (i) the share of inputs bought from and sold to service industries (NAICS 51–53); and (ii) the share of inputs bought from and sold to primary industries (NAICS 11–22).

To make our measures symmetric, for the industry pairs ij and ji we take the maximum of the respective coefficients. Hence, in our coagglomeration regressions the input coefficient for industries ij is the maximum of the two input coefficients ij and ji. We do the same for the output coefficients. Finally, in most of our specifications, we take the maximum of the input and output coefficients jointly (see Ellison et al., 2010, for additional details and discussion).

⁴⁴Due to confidentiality reasons, we did not use the finer *W*-level matrices which are internally available at Statistics Canada. However, tests ran in Behrens et al. (2015) using those matrices yielded similar results to those using the matrices constructed by our methodology.

To address endogeneity issues associated with input-ouput links, we also construct instruments based on the U.S. input-output benchmark tables from the BEA. Using the detailed 6-digit tables for 1997, 2002, 2007, and 2012 we construct the same input-output shares as explained above, using U.S. data. We again work with the whole input-output tables, including services and primary industries and excluding private consumption, government items, and imports/exports. We aggregate the data to the 4-digit level, which is identical to the Canadian NAICS that we use. Then, we compute the three shares (i), (ii), and (iii) as explained above. Note that the input-output benchmark tables are only available every five years. Hence, when required we assign twice the matrix to two consecutive years in our Scott's data.

A.3. Controls

The remaining control variables are constructed as follows:

Occupational employment similarity. We compute measures of occupational employment similarity of the workforce in the different industries. To this end, we use Occupational Employment Survey (OES) data from the *Bureau of Labor Statistics* for 2002, 2003, 2005, 2007, 2009, 2011, and 2013 to compute the share of each of 554 occupations in each 4-digit NAICS industry. We use 2002 data for the 2001 plant sample, and then data for each year t for the plant sample in year t. Using 2002 as the starting year for the OES data allows us to avoid a concordance from SITC to NAICS, and a concordance between the old and the new occupational classifications. Our measure of occupational employment similarity is computed as the correlation between the vectors of occupational shares of industries i and j. By construction, this measure is symmetric in ij and ji.

Patent citation patterns. We construct proxies for 'knowledge spillovers' or 'knowledge sharing' by using the NBER Patent Citation database and by following previous work by Kerr (2008). Our proxy for knowledge flows is the maximum of the shares of patents that industry i (or j) manufacture and which originate from the other industry j (or i). We take the maximum of the shares ij and ji to obtain a symmetric measure for the pairs ij and ji.

Trade data. The industry-level trade data come from Innovation, Science and Economic Development Canada 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

⁴⁵There are 808 occupations in the OES data. We only use occupations for which there is at least some employment in manufacturing (e.g., there are no 'Surgeons' in manufacturing industries, hence we exclude them completely from our data).

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. For each group of countries, we split exports and imports in all 4-digit industries. For each industry pair ij, we take the average share of industries i and j (either exports or imports) to make the measure symmetric.

Appendix B: Additional results

As a simple test of the effect of the number of bilateral trips and the balance of those trips on rates, we regress the sum of trips in both directions and the balance of trips on trucking rates (revenue per tonne-km). Trip and their balance are calculated across carrier types (Truckload, Less-than-truck-load, and Specialized) under the assumption that they serve different markets. Also included in the model is distance and distance squared and a set of year fixed effects. As expected, trip balance always has a positive effect on rates — see columns (1) and (2) — while the number of trips is positively related to rates, but only when it is entered in its quadric form, and for trade flows above 46 kilometres. Hence, the correlative evidence at hand suggests greater trade is more likely to raise rates and so the endogeneity of τ_{fb} likely results in an underestimate of its coefficient.

Table 13: Effect of bilateral trips and trip balance on the rate per tonne-km.

	(1)	(2)
Trips	0.002	-0.043 ^a
	(0.003)	(0.005)
Trips squared		0.006^{a}
		(0.001)
Balance	0.055^{a}	0.051^{a}
	(0.006)	(0.006)
Distance	0.715^{a}	0.843^{a}
	(0.062)	(0.063)
Distance squared	0.049^{a}	0.041^{a}
	(0.005)	(0.005)
Observations	66,698	66,698
R-squared	0.835	0.835

Notes: The dependent variable is $\ln(r_{m,k})$ measured on a per tonne-km basis. All variables in logs. Huber-White robust standard errors are in parentheses. Coefficients significant at: a 1%; b 5%; and c 10%.