

The Puzzling Persistence of the Distance Effect on Bilateral Trade*

Anne-Célia Disdier[†] Keith Head[‡]

April 8, 2008

Abstract

One of the best established empirical results in international economics is that bilateral trade decreases with distance. Although well-known, this result has not been systematically analyzed before. We examine 1467 distance effects estimated in 103 papers. Information collected on each estimate allows us to test hypotheses about the causes of variation in the estimates. Our most interesting finding is that the estimated negative impact of distance on trade rose around the middle of the century and has remained persistently high since then. This result holds even after controlling for many important differences in samples and methods.

JEL classification: F10, C10.

Keywords: gravity, trade costs, meta-analysis

*Previously titled: “Exaggerated Reports of the Death of Distance: Lessons from a Meta-analysis.” We thank the two anonymous referees, Deborah Swenson, Silvana Tenreyro, Thierry Mayer, and participants at seminars at the University of Toulouse and Geneva for helpful comments. We acknowledge financial support from the SSHRC-INE program and the ACI “Espaces et Territoires” (2004).

[†]INRA, UMR 210 Economie Publique INRA AgroParisTech, 16 rue Claude Bernard, 75005 Paris, France. Tel: (33) 1 44 08 72 72, Fax: (33) 1 44 08 16 63, Email: anne-celia.disdier@agroparistech.fr

[‡]Corresponding author: Sauder School of Business, University of British Columbia, 2053 Main Mall, Vancouver, BC, V6T1Z2, Canada. Tel: (604)822-8492, Fax: (604)822-8477, Email: keith.head@ubc.ca

1 Introduction

The answers to many important economic questions depend in large part on how much distance affects trade. In the factor proportions model, there is stronger pressure for factor price equalization if spatial separation does not raise trade costs. In models with increasing returns to scale, the magnitude of distance effects sets the penalty for geographic isolation. Regions that are distant from markets tend to have lower wages and/or less ability to attract foot-loose manufacturing. Outsourcing to remote suppliers may not be profitable if distance costs outweigh savings in production costs. Distance-based trade costs also amplify the gains from regional liberalization because proximity between the members make them “natural” trading partners.¹

Fortunately, there is no shortage of estimates of the effect of distance on trade flows. A huge number of papers have examined the determinants of bilateral trade flows and they almost invariably control for distance. Despite the abundance of such studies, assessments of the distance effect tend to be anecdotal. Leamer and Levinsohn (1995) mention two early studies and then claim “These and many subsequent studies have found a distance elasticity of about -0.6 .” Overman *et al.* (2003) state “the elasticity of trade volumes with respect to distance is usually estimated to be in the interval -0.9 to -1.5 .” They provide references to three studies as examples. The fact that Leamer and Levinsohn’s point estimate lies outside the interval proposed by Overman *et al.* reinforces the desirability of identifying a “typical” distance effect based on a larger sample of estimates.

This paper conducts a comprehensive, quantitative analysis of the magnitude of the distance effect and the factors that explain variation in the estimates. After using article search engines to construct a database of 1467 estimates from 103 papers, we find that the mean effect is about 0.9, with 90% of estimates lying between 0.28 and 1.55. On average, then, a 10% increase in distance lowers bilateral trade by about 9%.

Distance effects are found to be persistent in two senses: (1) they hold up in a very wide range of samples and methodologies, (2) they are not declining in studies employing

¹Overman *et al.* (2003) remark that “trade costs prevent goods price equalisation from occurring, and hence also prevent factor price equalisation.” They consider the effects of geography on factor incomes and the location of production. Krugman (1991) sketched the reasoning for expecting higher trade creation benefits relative to the costs of trade diversion for customs unions between neighboring countries.

more recent data. Distance effects of this magnitude pose an important unsolved puzzle. They almost certainly do not arise solely from transports costs. Using US data, Glaeser and Kohlhase (2004) report that “80% of all shipments (again by value) occur in industries where transport costs are less than 4% of total value.” Grossman (1998) performs a simple calculation showing that estimated distance effects are about an order of magnitude too large to be explained by shipping costs. He speculates that the reason why distance matters so much is lack of familiarity or cultural differences. This hypothesis is consistent with Blum and Goldfarb’s (2006) finding of a distance effect of 1.1 for “digital goods” consumed over the internet. Our paper does not resolve the puzzle of persistently high distance effects; rather, it establishes the central tendency and main causes of dispersion for distance effects found in the prior literature. These findings provide a reference point for subsequent investigations.

We envision several uses for the results contained in this paper. Authors such as Grossman (1998) wanting to refer to prior findings of large distance effects for trade in goods to motivate their observations can point to our mean from 1467 estimates rather than to an estimate plucked from a single study. Calibrated models where distance plays an important role, such as Alvarez and Lucas (2005), could use our mean or explore sensitivity of the simulation results across the range of common estimates. Finally, techniques for estimating bilateral trade equations improve, it should prove useful to compare new estimates to those obtained with prior methodologies.

The remainder of the paper is structured as follows. We briefly discuss the rationale for meta-analysis and some frequent criticisms of the approach in Section 2. We describe our sample in Section 3 and display the substantial variation in estimated distance effects. To permit interpretation of this variance, Section 4 describes three important sources of differences in results and proposes ways to investigate the contribution of each source. In Section 5 we present “meta-regression” results on the causes of distance effect variation. Our most striking finding is that, after slightly decreasing in the first half of the century, the distance effect begins to rise around 1950. We consider several studies that approached this issue using original data and discuss possible explanations for this pattern. Section 6 subjects our sample of distance effect estimates to two tests for “publication bias.” We conclude in Section 7.

2 The rationale for meta-analysis

Glass (1976) coined the term meta-analysis to “refer to the statistical analysis of a large collection of analysis results from individual studies for the purpose of integrating the findings.” In some medical meta-analyses, where estimates all come from randomized trials of comparable sets of subjects, the presumption is that the variation in estimates arises just from sampling error. In those cases the primary interest is to extract a single measure of the strength of some treatment. This contrasts with fields such epidemiology and social sciences that rely upon statistical analyses that are heterogeneous in terms of sample characteristics and estimation methodology. In those cases “integrating” past findings means more than just finding the best way to pool prior results. In the first illustration of meta-analysis, Smith and Glass (1977) related differences in the efficacy of psycho-therapy outcomes to the characteristics of the underlying studies (for example, the age of the clients and whether the study used randomization or not to create the control sample). As with this original meta-analysis, our goal is to determine the central tendency of the prior results as well as the determinants of variation.

The estimates that constitute our sample differ in the use of various econometric specifications, explanatory variables, degree of regional and product aggregation, and, of particular interest to us, time periods. As is common in meta-analyses going back to Smith and Glass (1977), we estimate “meta-regressions” in which the variance in the individual estimates is explained by a set of right-hand-side variables that quantify the different attributes of each study. This is useful because it reveals what aspects of the underlying gravity equation estimation really matter for determining the magnitude of the distance effect. The meta-regression can tell us, for example, whether more recent samples yield smaller distance effects, *holding other study characteristics constant*.

Economists have increasingly employed meta-analytical approaches.² Topics of investigation include the employment effects of minimum wages (Card and Krueger, 1995), the union-nonunion wage gap (Jarrell and Stanley, 1990), Ricardian equivalence (Stanley, 1998 and 2001), gender wage discrimination (Stanley and Jarrell, 1998), taxes and foreign direct investment (de Mooij and Ederveen, 2003), productivity spillovers (Görg and Strobl, 2001), the effects of

²For an introduction to the use of meta-analysis in economics, see Stanley and Jarrell (1989).

currency unions on trade (Rose and Stanley, 2005), the relationship between location decisions of firms and environmental regulations (Jeppesen *et al.*, 2002), and the rank-size relationship for cities (Nitsch, 2005).

Several criticisms have been raised against meta-analysis. The obvious critique is that when we take the mean of the effects estimated in multiple studies, we are averaging a possibly large number of poor studies with a relatively small number of high quality studies. Whereas the qualitative reviewer assigns little or no importance to estimates he does not trust, the meta-analyst may give undue weight to shoddy studies. Furthermore, there may be advances in estimating techniques that cast skepticism on the results obtained with techniques now deemed to be obsolete. Baldwin's (2006, Section 2.7) review of the literature on currency union effects exemplifies this viewpoint. He argues that since most estimates omit country-level fixed effects, they should be ignored, rather than weighted equally with methodologically defensible estimates. We view the averaging across results as a starting point to indicate objectively the central tendency in the prior literature. Our meta-regressions go a step further, investigating the magnitude and significance of different methods on results. We can not quantify all differences in methods. However, to the extent that papers published in better journals use systematically better methods, we are able to assess whether these unobserved quality differences yield significantly different distance effects.

A second criticism is that meta-analysis may combine estimates that are not actually comparable. Examples would include the combination of studies with different outcome or explanatory variables. Our study considers only estimates of the effects of geographic distance on bilateral trade. Even if studies measure the same relationships, they cannot be averaged if they measure effects using different units. Our study benefits from the near universal use of the gravity equation which estimates a units-free elasticity. Nevertheless, different populations—countries versus provinces, old versus recent data, industries versus aggregate trade—may have different distance parameters. The meta-regression method also allows us to explore whether different settings yield systematically different distance effects.

A third possible weakness of meta-analysis occurs in the presence of publication bias. The tendency among editors of academic journals to publish results that are statistically significant could bias the results of the meta-analysis. This problem also occurs for qualitative literature

reviews. We employ standard meta-analysis methods to show that there is almost no sign of publication bias in the distance effect estimation.

3 The distance effects data set

Distance effects are estimated as a parameter in the gravity equation. Gravity equations take a variety of forms in empirical implementations. Almost all of them can be represented in the following equation for the expected value of x_{ij} , the exports from country i to country j :

$$E[x_{ij}] = A_i^X A_j^M D_{ij}^{-\theta} \exp(\boldsymbol{\lambda} \mathbf{L}_{ij}). \quad (1)$$

In this equation A_i^X and A_j^M are indexes of the attributes of exporter i and importer j , D_{ij} is the distance between them, and \mathbf{L}_{ij} is a vector of bilateral indicators of the “linkages” between the two countries. We define θ as the “distance effect”, *the negative of the elasticity of bilateral trade with respect to distance*. The vector $\boldsymbol{\lambda}$ represents the coefficients on the linkage variables, which typically include sharing a common border (adjacency), common language, or colonial history. The name gravity refers to the similarity with Newton’s law, in which the A_i^X and A_j^M correspond to the masses of two objects, $\theta = 2$, and $\exp(\boldsymbol{\lambda} \mathbf{L}_{ij})$ is replaced with G , the gravitational constant. Unlike the equation from physics, the economic applications of gravity are not expected to fit the data perfectly. The standard approach is to assume $x_{ij} = E[x_{ij}] \eta_{ij}$, where η_{ij} has a conditional expectation of one. Then taking logs, one obtains

$$\ln x_{ij} = \ln A_i^X + \ln A_j^M - \theta \ln D_{ij} + \boldsymbol{\lambda} \mathbf{L}_{ij} + \ln \eta_{ij}. \quad (2)$$

In many empirical applications, exporter and importer indexes are implicitly assumed to be given by $A_i^X = Y_i^{\alpha_1^X} y_i^{\alpha_2^X}$ and $A_j^M = Y_j^{\alpha_1^M} y_j^{\alpha_2^M}$, where Y represents GDP and y is GDP per capita.³ Being linear in the parameters, equation (2) can be estimated with ordinary least squares.

There are four current exceptions to the general practice of estimating with OLS. First, Anderson and van Wincoop (2003) derive from first principles a symmetric version of equation

³Some specifications omit y since it does not emerge from any well-known theoretical derivation. Also papers that use the log of total bilateral trade, $\ln(x_{ij} + x_{ji})$, as the dependent variable need to impose symmetry in the A^X and A^M terms, restricting $\alpha_1^X = \alpha_1^M$.

(1) in which $A_i^X = A_i^M = Y_i P_i^{\sigma-1}$, where P_i is a country's "multilateral resistance index" and depends on the Y_j , $D_{ij}^{-\theta}$, $\lambda \mathbf{L}_{ij}$, for all countries. This approach makes (2) non-linear in the parameters and they therefore estimate it via non-linear least squares. Feenstra (2004) suggests that in many applications, an easier way to incorporate the insight of Anderson and van Wincoop (2003) is to estimate the $\ln A_i^X$ and $\ln A_j^M$ terms as fixed effects in linear equation.⁴ Second, a number of authors have recognized that OLS with $\ln x_{ij}$ as the dependent variable discards the $x_{ij} = 0$ observations. A variety of more or less *ad hoc* non-linear methods have been utilized to incorporate zeros or correct for their absence. Helpman *et al.* (2006) derive a method for handling zeros that builds on a fully specified model of firm-level heterogeneity in productivity. Due to non-linearity in the parameters, it is estimated via maximum likelihood.

The third departure from linear estimation is Coe *et al.* (2002), who assume $x_{ij} = E[x_{ij}] + \eta_{ij}$ and estimate with non-linear least squares. Finally, Santos Silva and Tenreyro (2006) show that heteroskedasticity in η_{ij} can cause the OLS and NLS methods to yield biased estimates. They argue that the most robust estimation method for multiplicative equations like (1) is Poisson pseudo-maximum likelihood.

The first step of the meta-analysis is to construct a database of estimates of θ for bilateral flows of trade in goods. Hundreds of empirical papers are based on the gravity equation. Therefore, it appears practically impossible to include all papers in our database. However, this abundance makes the construction of a representative sample of this literature easier. To maximize the replicability and objectivity of our meta-analysis, we built our base sample from English language papers listed in the *Econlit* database. We searched for "gravity equation," "gravity and distance and trade," and "gravity and history and trade" in the keywords or the title. As gravity equations are used in the recent empirical literature on border effects, we also searched for the keywords "border and distance and trade" and "home bias and trade."

The final sample based on *Econlit* keywords comprises 78 papers, 60 of which are published in academic journals, 4 are chapters in books, and 14 are working papers. This sample omits many well-known papers estimating distance effects. Also the sample draws so heavily on lesser journals that it would be difficult to estimate the impact, if any, on the distance effect of being

⁴The importer and exporter fixed effects approach is increasingly popular. In addition to simplicity, it has the advantage of being consistent with other, asymmetric, derivations of the gravity equation such as Eaton and Kortum (2002) and Chaney (2006).

published in a high-quality outlet. Therefore, we augmented our sample by a searching in JSTOR for papers published in the *American Economic Review* and the *Review of Economics and Statistics*. We also searched the website of the *Journal of International Economics*. This search identified 25 additional papers.

The next step involved deletion of estimates that were not in the form of elasticities. We eliminated 5 estimates where trade was entered in levels, 10 where distance was entered in levels, and 9 where both trade and distance were entered in levels. These estimates would not be comparable to elasticity estimates from the standard linear-in-logs specification.⁵

Extreme deviations from the main sample of estimates would be problematic for both our graphical and statistical analyses. In the former, they compress the variation of the rest of the sample and in the latter, they are likely to lead to fragile findings. Hence, we deleted some extreme outliers using the Grubbs test (NIST/SEMATECH, 2004) as our criterion. This test calculates, G , the maximum deviation from the sample mean, \bar{x} , divided by the standard deviation, s , calculated including that observation. If $G = \max\{|x_i - \bar{x}|/s\} > G^*$, the observation is deleted. The critical value for G for confidence level α is given by $G^* = (N-1)\sqrt{t^2/[N(N-2+t^2)]}$, where t is the critical value of a t-distribution with $N-2$ degrees of freedom and a confidence interval of $\alpha/(2N)$. We start with $N = 1475$ observations and set $\alpha = .05$, yielding $G^* = 4.13$. The Grubbs procedure eliminates one outlier at a time, recalculating \bar{x} and s with each iteration. Application of this procedure led to the removal of 8 distance coefficients.⁶ This caused the standard deviation of $\hat{\theta}_i$ to decline from 1.57 to 0.40 while the mean hardly changed (0.94 vs. 0.91).

After the above deletions, the 103 studies provide 1467 usable observations. An online appendix, available at http://strategy.sauder.ubc.ca/head/papers/meta_papers.pdf, lists the full sample of 103 papers, including the number and range of estimates from each paper. They span a relatively large period going from 1870 to 2001, including 188 pre-1970 sample estimates.

The estimated distance coefficients ($-\hat{\theta}_i$) range from 0.04 to -2.33, with 1466 estimates—all

⁵The Poisson pseudo maximum likelihood estimates of Santos Silva and Tenreyro (2006) are included because they also generate elasticities of trade with respect to distance.

⁶The eliminated observations were 51.71, -26.68, 7.28, 6.53, 2.84, 2.8, 2.63, and 2.62. With a 1% confidence level, 2.63 and 2.62 would have been retained. Use of a 10% confidence level deletes the same values as the 5% level. A Stata program is available from the authors.

except one—of a negative effect of distance on bilateral trade. The mean distance effect is 0.91 and the median is 0.87.⁷ Taking a simple mean does not make use of any information on the precision of each estimate. A minimum variance estimate of the mean weights each individual estimate by the inverse of its variance. For 1257 estimates with reported standard errors, the weight given to each estimate $\hat{\theta}_i$ is $\omega_i = 1/\text{s.e.}(\hat{\theta}_i)^2$. The resulting minimum variance estimate is 0.86. A second possible way to weight observations is to use degrees of freedom. Other things equal, the standard error of $\hat{\theta}$ should be inversely proportional to the square root of the degrees of freedom. One over the variance would therefore be proportional to the degrees of freedom. This alternative weighted mean—which is less vulnerable to econometric methods that underestimate the standard errors—is 1.07.

Figure 1 provides the kernel density estimates. The vertical lines show that both arithmetic mean and the inverse-variance weighted mean are quite close to the mode of the distribution.⁸ The “best” estimates from each paper—as measured by the R^2 of the corresponding regression—are depicted graphically in figure 1 using short vertical lines arrayed along the top side of the figure. These estimates average 0.79. We will return to these “best” estimates when we consider the possibility of publication bias in section 6.

The figure illustrates the large amount of variance in the estimated distance effect. The fifth percentile is 0.28 and the 95th is 1.55. The standard deviation is 0.39 and the interquartile range is 0.52. The next section categorizes the main causes of this variation.

4 Why distance effects vary

We should expect different studies and different specifications within studies to generate different estimates for three main reasons.

sampling error: chance errors in estimating a population parameter arising from the finite sample drawn from that population.

“structural” heterogeneity: differences in parameters across sub-populations of the data.

⁷Thus more than half of the sample lies below the interval suggested by Overman *et al.* (2003).

⁸An earlier version of this paper mistakenly included 10 estimates that were based on distance measured in levels, rather than logs. Their very low estimated standard errors pushed the inverse-variance weighted mean down to 0.5, illustrating the sensitivity introduced by these weights.

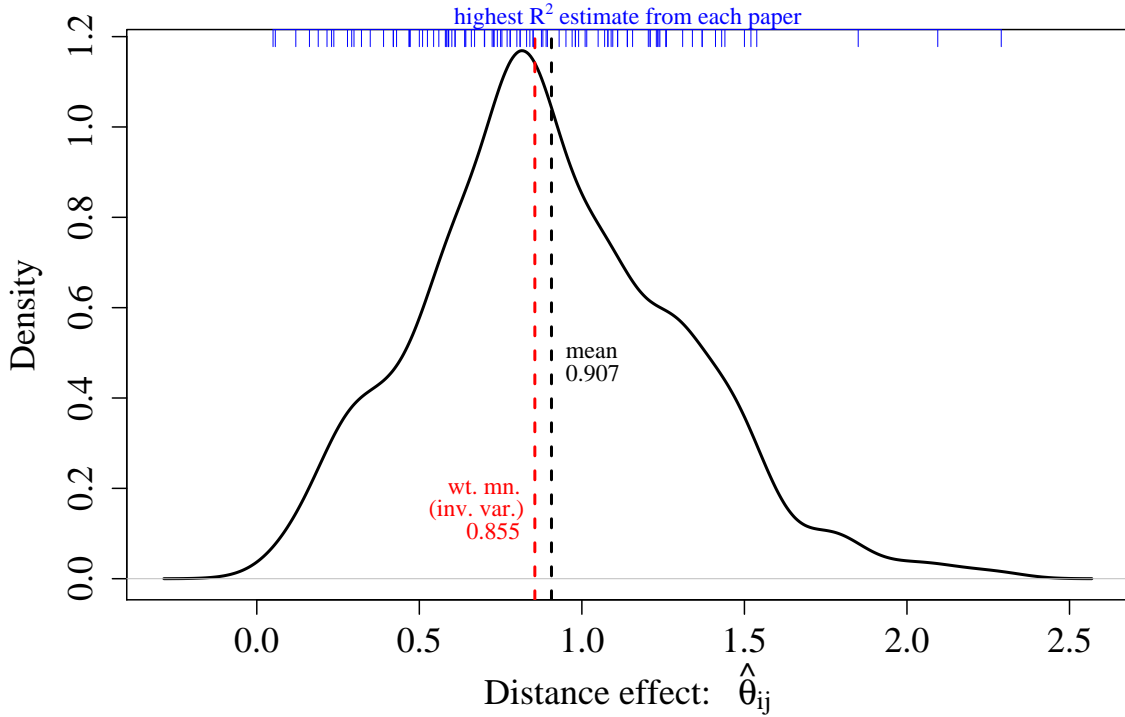


Figure 1: Estimated density of the 1467 $\hat{\theta}$ estimates

“method” heterogeneity: differences in statistical technique lead to different estimates.

These include biases away from the true value caused by mis-measurement of the explanatory variable or omission of important control variables.

As we shall describe below, we introduce a number of explanatory variables to quantify the importance of each source of heterogeneity in estimated distance effects.

4.1 Sampling error

Distance coefficients are usually based on a sample of countries and years. Even if all samples were drawn from a population with the same underlying distance effect, regression estimates of the distance effect would differ from the true population mean by an amount referred to as sampling error. In areas like medicine, a key purpose of meta-analysis is to reduce sampling error by combining estimates from many studies, each of which had small samples (usually due to high costs of clinical trials). In economics, sampling error is expected to play less of a role. Indeed, the “sample” may include the entire relevant population of trading partners in a gravity equation.

We investigate the role of sampling error in gravity equations using a graphical display and the I^2 statistic. Let $\hat{\theta}_i$ represent an individual estimate of the distance effect and $\bar{\theta}$ be an estimate of the population mean. Define $z_i \equiv (\hat{\theta}_i - \bar{\theta})/\text{s.e.}(\hat{\theta}_i)$. Under the null of a single population mean, z_i should follow a t distribution with $n_i - k_i$ degrees of freedom. In our database, $n_i - k_i$ is always 29 or higher, with a median of 1035. With these degrees of freedom, the t -distribution closely resembles the standard normal. Hence we graph z_i with the standard normal as the benchmark for the case of a common population parameter.

Figure 2 reveals that sampling error can only explain a small portion of the variance in the $\hat{\theta}_i$. More formally, we estimate a Higgins *et al.* (2003) I^2 statistic of 98.2%, which indicates “the percentage of total variation across studies that is due to heterogeneity rather than chance.”⁹ The graph and I^2 statistic make a clear case for investigating the causes of *heterogeneity* in the distance effects.

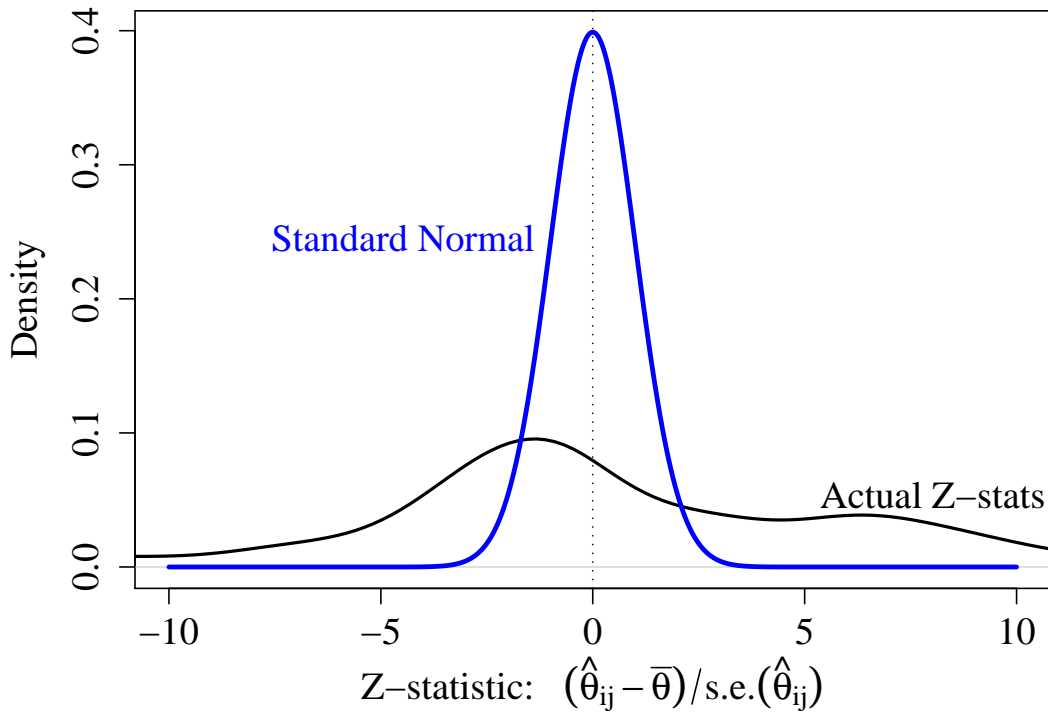


Figure 2: Excess dispersion of z_i relative to the standard normal.

⁹We made the calculation using the `metagen()` function of R, www.r-project.org.

4.2 Structural heterogeneity

Gravity equations have been estimated for heterogeneous sets of countries and industries. It seems likely that different “sub-populations” will have different distance effects. To understand why heterogeneity in the distance effect is likely to be important we now consider the structural parameters underlying the distance effect, θ , estimated in equation (2). Here we draw on the formal derivations of gravity equations by Eaton and Kortum (2002), Anderson and van Wincoop (2003), and Chaney (2006).

In general terms we can think of θ as *the product of two elasticities*: (a) the elasticity of trade costs with respect to distance, and (b) the elasticity of trade with respect to trade costs. Anderson and van Wincoop (2004) denote factor (a) as ρ . In derivations involving differentiated products and representative firms, factor (b) is given by $\sigma - 1$ where σ is the elasticity of substitution between varieties in monopolistic competition or between regional products in Anderson and van Wincoop (2003). Since $\theta = \rho(\sigma - 1)$, cross-sample variation in either the trade cost parameter or the preference parameter can lead to structural heterogeneity in the distance effect. Eaton and Kortum (2002) show that the gravity equation can also arise in a model of homogeneous goods with heterogeneity across nations in their production technologies. In that model factor (b) is given by an inverse measure of international productivity dispersion. Chaney (2006) maintains differentiated products but allows for firms with heterogeneous productivity. He also allows distance to determine the fixed costs of bilateral trade. One interpretation would be distance impedes information flows, increasing upfront search costs, as emphasized by Rauch (1999). In Chaney’s model, $\theta = \rho_1(\sigma - 1)$ for individual firms, where ρ_1 is the elasticity of *variable* trade costs with respect to distance. However, for aggregate trade flows, endogenizing the entry process leads to very different results: $\theta = \rho_1\gamma + \rho_2[\gamma/(\sigma - 1) - 1]$, where γ is the inverse of the standard deviation of log productivity, and ρ_2 is the elasticity of *fixed* trade costs with respect to distance.

The takeaway from this discussion is that whenever a gravity equation is estimated on data with different degrees of substitutability between goods (σ), different productivity dispersion (γ), or different responsiveness of trade costs to distance (ρ), we should expect different distance effects on trade.

Our database includes samples of world trade in which many country pairs are separated

by oceans. It also includes studies of trade between regions within a single continent, where most goods trade probably takes place along land routes. Therefore we introduce a dummy “single continent” to capture land versus ocean differences in ρ_1 . We also include dummies for whether the estimation includes only developed economies, only developing and/or transition countries, or a mix of both groups. We expect lower distance effects for sets of rich countries because of superior transport infrastructure leading to low ρ_1 . This would not have to be the case in the Chaney model, since it is not clear how γ and σ vary with the level of development. To take into account likely differences in σ and γ for more disaggregated definitions of goods, we include a dummy for industry and product-level trade data.

4.3 Method heterogeneity

Meta-analysis of distance effects is facilitated by the widespread use of the gravity equation. The linear-in-logs functional form makes it possible to compare estimated elasticities directly. Within the gravity equation framework, there remain a large number of minor and major differences in econometric methods.

One way that studies differ is in measurement of trade flows and distance. Some studies sum imports and exports for a given country pair whereas other studies focus on directional trade. We take this into account by including a dummy to control for whether the dependent variable is total bilateral trade flows or only bilateral import or export flows. Most of the papers calculate distances using the great-circle formula. This is appealing because it requires only the latitude and longitude of principal cities for each country. However, great-circle routes often differ substantially from actual cargo routes (especially when the former cross over the poles). A small number of authors have collected actual distances traveled by road or sea. We use a dummy to test the hypothesis that this improvement in measurement should increase the estimated distance effect.

Studies also differ in their selection of control variables. Some regression specifications omit variables that are (i) important determinants of trade and (ii) correlated with distance. This induces omitted variable bias in the estimated distance effects from those studies. We consider four dummies that each takes a value of 1 if the underlying estimation controls respectively for adjacency, the sharing of a language, the belonging to the same preferential

trade agreement, and the inclusion of a measure of remoteness. For cases where a common language or membership of a preferential trade area is irrelevant due to the nature of the sample (e.g. trade between states in the US), we specify those dummies with a value of one, because we wish to interpret a zero as reflecting a failure to control for an important variable *when necessary*.

Anderson and van Wincoop (2003) point out that many of the remoteness variables employed in the gravity literature do not use functional forms that correspond to the underlying theory. A specification consistent with theory involves the use of fixed effects for each importer and exporter. This specification overcomes what Baldwin (2006) calls the “gold medal” mistake in gravity equations—the failure to consider relative prices. We therefore add a dummy for whether the estimation includes countries fixed effects or not.

While most large countries trade with each other, there are some country pairs with zero trade, especially in disaggregated industries. Discarding these zeros would result in selection bias that again might have an effect on distance effects. We code a dummy which takes a value 1 if the estimation incorporates or corrects for the zero flows. Some samples do not include zero flows (e.g. aggregate trade flows between OECD countries). We therefore code an additional dummy equals to 1 if there are no zero flows in the sample and hence no need to devise a method to incorporate or correct for them.

The problem of excluding zeros arises because of the linear in logs specification conventionally used in gravity estimation. Santos Silva and Tenreyro (2006) challenge this approach, arguing that it is particularly vulnerable to heteroskedasticity. They advocate the use of poisson pseudo maximum likelihood (PPML). We include a PPML dummy for their estimates using this method.

Another econometric problem researchers occasionally address in gravity equations is the possible endogeneity of the GDP terms. This could cause simultaneity bias that might feed into bias in distance effects. We include a dummy that equals one if the authors control for the potential correlation between the GDP regressors and the errors terms.

Finally, we control for the quality of the publication by adding a dummy equals to 1 if the study was published in the *American Economic Review*, the *Journal of International Economics*, or the *Review of Economics and Statistics*. The idea is that the standard for

methodological rigor could be higher at these journals and this could correct biases present in other estimates.

Table 1 summarizes the meta-independent variables and presents the mean for each sub-period. A few trends seem noteworthy. First, the distance effect is much smaller before 1970. Second, studies looking at trade within a continent tend to use more recent data. It seems likely that this accounts for the rise in the share of developed countries only in later samples. Third, authors using recent data are more likely to control for remoteness. Finally, studies of recent sample periods tend to distinguish between export and import flows rather than summing them.

Table 1: Descriptive statistics

Name	Mean			
	≤ 1969	1970s	1980s	≥ 1990
Dependent variable				
Distance Effect (– elasticity of trade w.r.t. dist.)	0.63	0.90	0.96	0.95
Structural variables				
Average Year (Midpoint of estimation period)	1937.59	1974.30	1984.94	1993.46
Single Continent	0.15	0.20	0.29	0.28
Developed Economies Only	0.28	0.45	0.42	0.42
No Developed Economies	0.00	0.11	0.04	0.03
Disaggregated Data	0.18	0.42	0.21	0.33
Method Variables				
Total Bil. Trade (sum of two-way trade flows)	0.56	0.31	0.34	0.20
Road/Sea distance	0.13	0.28	0.11	0.13
Adjacency Control	0.80	0.75	0.47	0.50
Common Lang. Control	0.34	0.46	0.56	0.45
Trade Agreements Control	0.91	0.91	0.64	0.62
Remoteness Control	0.01	0.19	0.20	0.21
Country Fixed Effects	0.14	0.01	0.07	0.12
Incorporates Zero Flows	0.38	0.06	0.16	0.12
No Zero flows to be incorporated	0.28	0.53	0.43	0.29
Poisson PML	0.00	0.00	0.00	0.01
GDP Endogeneity Correction	0.07	0.01	0.13	0.08
High Quality Review	0.36	0.50	0.25	0.25

5 Meta-regressions results

In this section, we present the empirical results for the meta-regression analysis. All studies in our sample, except one, report more than one estimate. Different observations from the same paper are not likely to be independent. Concern over lack of independence causes some authors to reduce their sample to a single observation per paper. We do not adopt the one-estimate-per-study approach for three reasons. First, it is inefficient to discard information. Second, it is not clear which estimate one should use.¹⁰ Third—and most importantly for our purposes—different estimates often differ in terms of sample period, method, etc. and therefore within-study variation among distance effect estimates can be used to assess the importance of such variables. We consider estimates from the same study as distinct—but possibly correlated—observations in our meta-regressions.

To retain the useful information contained in the multiple estimates from each paper while dealing with the problem of dependence between these estimates, we adopt the random effects panel specification suggested by Jeppesen *et al.* (2002). The estimated model can be expressed as

$$\hat{\theta}_{ij} = u_i + \beta \mathbf{X}_{ij} + e_{ij},$$

where $\hat{\theta}_{ij}$ is the j th distance coefficient reported in study i , \mathbf{X}_{ij} is the matrix of the meta-independent variables included to explain the variation of the distance elasticities, and β is the vector of meta-regression coefficients. The u_i are the random paper effects.¹¹ To test the robustness of our random-effects results, we also report OLS coefficients with paper-clustered standard errors. Unlike random effects, the OLS regression gives as much weight to *between*-paper variation as it does to *within*-paper variation.

Results are reported in table 2. The plan of the table is to take a first pass at establishing the time trend in distance coefficients (columns 1 and 2 and Figure 3). Then we will add “structural” variables (column 3) and method variables (column 4) to assess whether these controls change the results. The last column focuses on the question of whether the OLS

¹⁰Some meta-analyses (e.g. Card and Krueger, 1995) identify a “preferred” estimate, others use averages or medians of the estimates from each paper and some even randomly select one estimate (Stanley, 2001).

¹¹In Jeppesen *et al.* (2002) the u_i correspond to random author effects. As our database includes multiple papers written by the same authors, we decided to use random paper effects. We also experimented with a nested random effects model with author and paper effects. The results were very similar to the specification with just paper effects.

cluster regression provides different results.

Table 3 presents measures of the overall fit of the regressions. For panel specifications, there are two dimensions of variation, *between* and *within* studies. Recall that the error associated with the former is u_i and the latter is e_{ij} . The diagnostics reveal that the majority of the variation in the distance effect cannot be explained by the meta-regression variables.

Specification (1) regresses the distance estimates solely on the midyear of each sample. We subtract 1870 from this variable so that the constant can be interpreted as the distance effect for the earliest observation of the data set. The estimated coefficient on the average year is positive and significant at the 1% level. This implies that the negative impact of distance on trade seems to be increasing over time.

The linear trend specification of column (1) is a strong assumption to impose on estimates spanning 130 years. Figure 3 graphs the estimates against time and fits a kernel smoother through the data. The highest R^2 estimate of each paper is shown with solid circle. The lighter lowess smoother line is associated with these estimates. The darker line is a smoother through the entire set of estimates. The relationship between distance effects and sample period seems fairly flat until the 1950s. We therefore estimate the rest of our regressions with a more flexible form based on dummies for 4 period ranges: before 1970, the 70s, the 80s, and since 1990.

Column (2) provides the results for this specification. It shows a significant increase in the distance effect in the post-1970 data. Distance impedes trade by 37% more—(0.25/0.68)—since 1990 than it did from 1870 to 1969.

An important question is whether distance really matters more in later periods or whether there are systematic differences in the attributes of studies that cause upward bias in the estimates for later data. Columns (3) and (4) control for aspects of the estimates that could matter.

The first set of controls—shown in column (3)—are what we term “structural” variables. The coefficient on the variable “single continent” is positive and significant at the 1% level. Given that intra-continental trade is much more likely to involve land transport, this suggests a higher transport elasticity for land trade. Hummels (2001a) shows that distance has a larger effect on freight costs for rail and truck shipments than for ocean shipments (although he finds that all three have lower distance elasticities than air freight). We find a significant and

Table 2: Meta-regressions using 1467 distance coefficients

Specification:	(1)	(2)	(3)	(4)	(5)
Method:	R.E.	R.E.	R.E.	R.E.	OLS
Intercept	0.54 ^a (0.09)	0.68 ^a (0.05)	0.60 ^a (0.05)	0.67 ^a (0.06)	0.54 ^a (0.11)
Midyear of sample - 1870	0.003 ^a (0.001)				
1970 ≤ Av. Year ≤ 1979		0.18 ^a (0.04)	0.17 ^a (0.04)	0.17 ^a (0.04)	0.25 ^a (0.08)
1980 ≤ Av. Year ≤ 1989		0.24 ^a (0.04)	0.23 ^a (0.04)	0.21 ^a (0.04)	0.33 ^a (0.08)
Av. Year ≥ 1990		0.25 ^a (0.04)	0.24 ^a (0.04)	0.23 ^a (0.04)	0.29 ^a (0.07)
Single Continent			0.33 ^a (0.06)	0.24 ^a (0.06)	-0.01 (0.08)
Developed economies only			-0.07 (0.05)	-0.06 (0.05)	0.19 ^b (0.08)
No developed economies			0.42 ^a (0.06)	0.44 ^a (0.06)	0.55 ^b (0.22)
Disaggregated Data			0.11 ^b (0.05)	0.09 ^c (0.05)	-0.07 (0.11)
Total Bilateral Trade				-0.06 ^c (0.04)	-0.15 ^b (0.06)
Road/Sea Distance				0.00 (0.07)	0.11 (0.08)
Adjacency Control				-0.15 ^a (0.03)	-0.08 (0.05)
Common Language Control				0.10 ^a (0.03)	0.12 ^b (0.06)
Trade Agreements Control				-0.01 (0.03)	0.09 (0.06)
Remoteness Control				0.03 (0.03)	-0.04 (0.07)
Country fixed effects				0.14 ^a (0.03)	0.26 ^a (0.09)
Incorporates zero flows				0.08 ^b (0.03)	0.13 ^c (0.07)
No Zero flows				-0.08 ^c (0.05)	-0.27 ^a (0.06)
Poisson pseudo-ML				-0.35 ^a (0.13)	-0.22 ^b (0.10)
Corrects for GDP Endogeneity				-0.02 (0.03)	0.03 (0.05)
High quality review				0.09 (0.06)	0.15 ^b (0.06)

Note: Standard errors (clustered in specification 5) in parentheses with ^a, ^b and ^c respectively denoting significance at the 1%, 5% and 10% levels.

Table 3: Meta-Regression Diagnostics: Random “Paper” Effects

Specification:	(1)	(2)	(3)	(4)	(5)
Method:	R.E.	R.E.	R.E.	R.E.	OLS
R^2 overall	0.048	0.073	0.134	0.202	0.366
R^2 between	0.051	0.069	0.139	0.255	n/a
R^2 within	0.010	0.019	0.090	0.119	n/a
Std. error of u_i	0.316	0.316	0.308	0.261	n/a
Std. error of e_{ij}	0.221	0.220	0.212	0.209	0.316

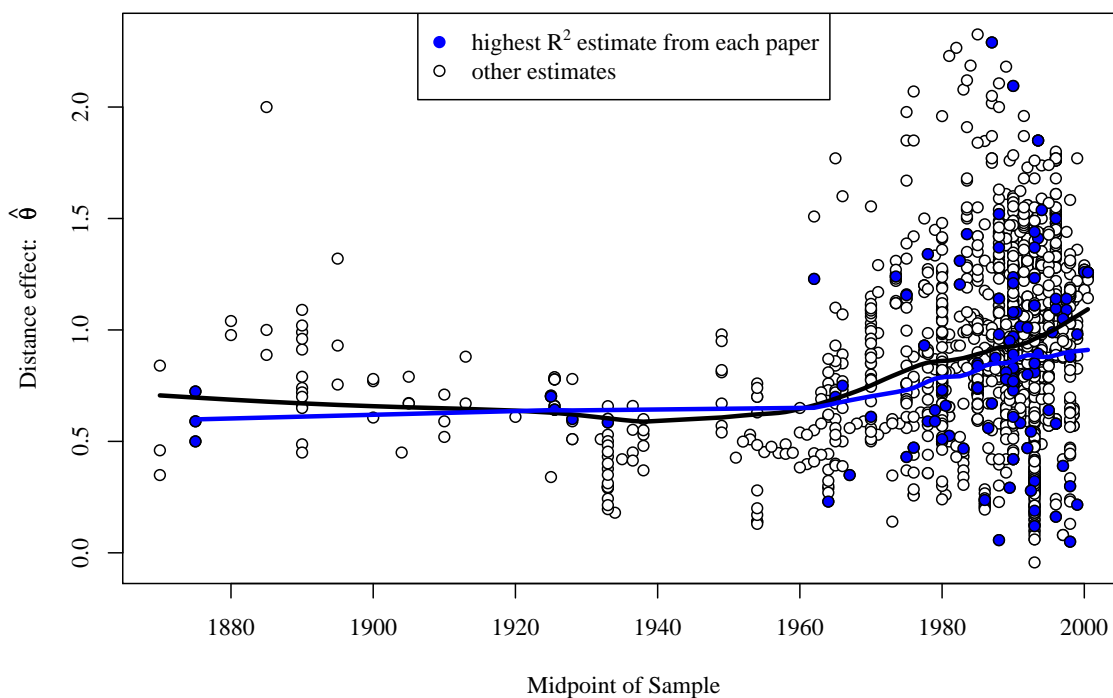


Figure 3: The variation of $\hat{\theta}$ graphed relative to the mid-period of the data sample.

positive coefficient on the dummy for samples that do not include any developed countries, while the coefficient on samples of only developed countries is negative. A higher distance effect for samples of poorer economies conforms to expected differences in the quality of transport infrastructure.

The coefficient on disaggregated data shows that estimations based on industry or product-level data tend to obtain higher distance effects than estimations conducted with aggregate data. This makes sense in the representative firm model (see subsection 4.2) if one thinks of disaggregated industries as having larger elasticities of substitution (σ). Hummels (2001a) supports this presumption, finding that a move from 1 to 2 to 3 digits of disaggregation, raises estimates of σ from 4.8 to 5.6 to 6.9. In the Chaney (2006) formulation—where higher σ lowers the distance effect—the higher θ could be interpreted as the consequence of less productivity heterogeneity in disaggregated sectors.

Column (4) introduces twelve “method” variables that consider issues such as measurement of the key variables, the set of controls, the set of econometric “corrections”, and journal quality. Estimates that sum (or average) bilateral exports and imports before taking logs obtain smaller distance effects than studies that use log exports as the dependent variable. Thus, what Baldwin calls the “silver medal” mistake appears to cause a borderline significant negative bias. The distance coefficient is hardly affected by the use of distances by road or by sea. This is unexpected since great circle distances disregard so much about the actual geography of transportation. Perhaps it is reassuring to all who have relied on the great circle distances because they wanted to avoid the task of collecting actual distances traveled.

Our next results suggest that omitted variable bias can have significant impacts on estimated distance effects. In particular a failure to include a dummy for adjacent country pairs seems to cause an overestimate of the distance effect. This makes sense since adjacency is likely to be negatively correlated with distance, leading to upward omitted variable bias (on the distance effect—the bias on the negative distance coefficient would be downward).

Another important control is for a common language. Here the correlation with distance is not obvious. Some pairs like Belgium and France, Ireland and the U.K., are relatively proximate whereas country pairs that share a language due to colonization patterns (U.K. and Australia, say) are very far apart. The results suggest that the latter set of countries dominate:

the inclusion of the common language control significantly raises the distance effect.

Two other controls that one would expect to matter have a negligible impact. Controlling for membership of a preferential trade agreement has a small and insignificant effect. Distance effects on trade also seem to be insensitive to the introduction of a “remoteness” control variables. As mentioned before, this might be because many of the remoteness variables do not use proper functional forms. Our result shows that the use of fixed effects instead of a-theoretical remoteness variables increases the distance coefficient.

Using methods that incorporate or correct for zero trade flows seems to raise the estimated coefficients. On the other hand, samples that do not have zero flows tend to obtain smaller distance coefficients. However, this result is significant only at the 10% level. In unreported results, we investigated whether the particular method for dealing with zeros matters. Tobit and Heckman methods tend to yield considerably larger estimates, corroborating Overman *et al.*'s (2003) observation that “The difference in estimated [distance] coefficients arises, at least in part, because of the treatment of zeros. Tobit estimation typically yields larger coefficients.” The standard errors on these method indicators are large: only the Tobit procedure makes a statistically significant difference.¹²

The Poisson PML method advocated by Santos Silva and Tenreyro (2006) leads to much smaller distance effects estimates. This is based on just 4 estimates in one paper for one year of data, 1990. It seems worthwhile to investigate the PPML method for alternative samples and time periods.

Using instruments to control for the endogeneity of GDP has no discernable impact on the distance effect. Finally, the distance effects in high quality journals do not differ significantly from the rest of the sample.¹³

Recall that we constructed our sample by combining estimates from papers found through an Econlit search with papers found through a more focussed search within specific journals. The Econlit sample is more objective because we exercised more discretion in selecting the

¹²The Helpman *et al.* (2006) paper does not enter our sample because, at the time of writing, it was a mimeo and therefore not listed in Econlit. The use of their ML method capturing the heterogeneity effect on trade partner selection reduces the distance effect implied for a firm by 0.4 (from 1.2 to 0.8). A Heckman correction alone slightly raises the distance effect.

¹³An alternative proxy for improved econometric method is the year of publication. In unreported regressions we experimented with time trends and period dummies based on publication year but found small and insignificant effects.

remaining papers. The time effects for that sample (unreported) are slightly lower than for the whole sample but there are no other noteworthy differences in the results.

The random effects method places greater emphasis on within-paper variation than cross-paper variation. We report results based on the OLS in column (5) of Table 2. In this specification, we deal with the correlation between coefficients in the same paper just by clustering standard errors at the paper level. The increasing distance effect after 1970 remains quite pronounced, although it now peaks in the 1980s, instead of the final sub-period.

Estimated coefficients on other variables reveal some notable differences. Dummies for single continent, disaggregated data, and adjacency are no longer significant. In addition, country fixed effects have a bigger (positive) impact on the distance estimates and samples without zero flows obtain a much smaller distance coefficient.

After controlling between and within-study differences in sample composition, controls, and methods, we find that the basic message of figure 3 remains intact: estimated distance effects are not diminishing over time and in most specifications they seem to be rising. This finding raises the puzzle of how to reconcile technology driven reductions in trade costs with a non-shrinking effect of distance.

A recent literature examines the evolution of the distance effect and generates mixed results. Brun *et al.* (2005) and Coe *et al.* (2002) conduct panel estimation of gravity equations from 1962–1996 and 1975–2000 using the IMF DoTS data set. For standard gravity specifications, both studies find rising distance effects. Brun *et al.* are able to find a declining trend for distance effects only when they estimate an “augmented” specification confined to the sample of rich countries. Coe *et al.* find declining distance effects when they re-specify the gravity equation with an additive error term and estimate it using non-linear least squares.¹⁴ Combes *et al.* (2006) estimate distance coefficients year by year, using fixed effects for exporters and importers. They find a pattern of rising coefficients since the 1950s. Instead of estimating gravity equations, Carrère and Schiff (2005) weight bilateral distances by bilateral trade flows to calculate an “average distance of trade” by country from 1962–2000. They find that these distances are falling over time, corroborating the gravity equation results of larger distance effects.

¹⁴Santos Silva and Tenreyro (2006) report monte-carlo simulations showing “extremely poor performance” of non-linear least squares estimator in the presence of some forms of heteroscedasticity.

One possible cause of rising distance effects for the aggregate bilateral trade flows used in the preceding studies would be a shift in composition towards industries with relatively high distance effects. Berthelon and Freund's (2004) study of industry-level trade finds that 75 percent of industries do not exhibit significant changes in the distance effect. The significant changes are almost all in the direction of a larger distance effect over the 1985–2000 period.

Felbermayr and Kohler (forthcoming) suggest an econometric resolution to the puzzle of rising distance effects. They start by noting that there has been a dramatic expansion in the *extensive* margin of international trade: a much higher share of country pairs trade positive amounts now than did in 1950. OLS on the sample of positive traders can yield inconsistent distance estimates but the bias should disappear gradually as the share of positive traders rises. This account implies that the early estimates of distance effects are biased downwards and the more recent estimates are closer to the “true” values. The authors find that going from OLS to Tobit switches the sign of the time trend interacted with distance to be negative instead of positive.

Even if Felbermayr and Kohler (forthcoming) are correct about the cause of rising *estimated* distance effects, it remains an open question why the distance effect persists in being so large. In a recent book, Frances Cairncross (1997) announces “the death of distance” due to advances in electronic communication. With slightly less hyperbole, Glaeser and Kohlhase (2004) comment “Certainly it is an exaggeration to claim that moving goods is free, but it is becoming an increasingly apt assumption.” Why hasn't technological progress dramatically *reduced* the distance effect?

Three explanations for the persistent effect of distance strike us as worthy of investigation. First, technological progress may have been smaller or less ubiquitous than casual empiricism would suggest. In particular, advances like email and teleconferencing may not radically alter the marginal cost of distance for trade in goods. Second, Hummels (2001b) and Deardorff (2003) suggest that the influence of time on trade is increasing. Greater use of just-in-time as well as a simple income-driven increase in the value of time could raise distance costs. Third, changes in the composition of trade might be biased towards goods with high distance costs. Berthelon and Freund (2004) find that industry-level compositional changes had almost no impact on the distance effect.

Nevertheless, Duranton and Storper (2006) provide a model in which *within industries*, falling transport costs prompt firms to trade more sophisticated goods with higher transaction costs. This mechanism could endogenously maintain high distance effects in the face of falling transport costs.

6 Publication bias

A persistent concern in all literature reviews, is that the publication process may have influenced the set of findings available to be assessed. To the extent that referees and editors of academic journals insist upon statistically significant results, the published sample of results will differ systematically from the full set of estimates. Fortunately, researchers using meta-analysis have developed tests to uncover the presence of publication bias.

We use two methods proposed by Card and Krueger (1995) in a meta-analysis of the employment effects of minimum wages. Since all the estimates in a given paper are either published or not, it does not make much sense to consider publication bias at the level of individual estimates. For this reason we follow Card and Krueger’s (1995) practice of reducing the sample to one $\hat{\theta}_i$ estimate per paper. Card and Krueger selected a “preferred estimate” for each of the papers in their meta-analysis. In our sample it is often infeasible to determine a single estimate preferred by the authors. We opt instead to select the “best” estimate from each paper using a quantitative criterion: the highest R^2 for the corresponding regression. For the two papers that do not report R^2 , we use the last estimate of each paper.

The first method consists of a regression of the log of the t -statistic on the distance coefficient on the log of the square root of the degrees of freedom. Suppose that the size of the sample is determined exogenously by data availability. Then, sampling theory predicts that the absolute value of the t -statistic should be proportional to the square root of the degrees of freedom. In the absence of publication bias we therefore expect a unit value on the estimated elasticity. Some papers which do not report standard errors or t -statistics, and/or sample size are excluded. This leaves us with a sample of 87 papers, 71 of which are published in journals. Using just the latter we find a coefficient of 0.65 with a standard error of 0.09. Although this is significantly less than one, one should not interpret this as strong evidence for publication bias. The key finding is that increasing sample size does have the large positive effect on

significance that sampling theory predicts. In sharp contrast, Card and Krueger (1995) find a *negative* relationship in their meta-analysis. Görg and Strobl (2001) find a slightly negative correlation in their study of productivity spillovers from multinational corporations (MNCs).

Figure 4 illustrates the 87 studies reporting enough information to calculate t -statistics and degrees of freedom. Note that the 16 book chapters and working papers (depicted using hollow circles) do not appear to have a markedly different pattern than the papers published in journals (shown with solid circles).

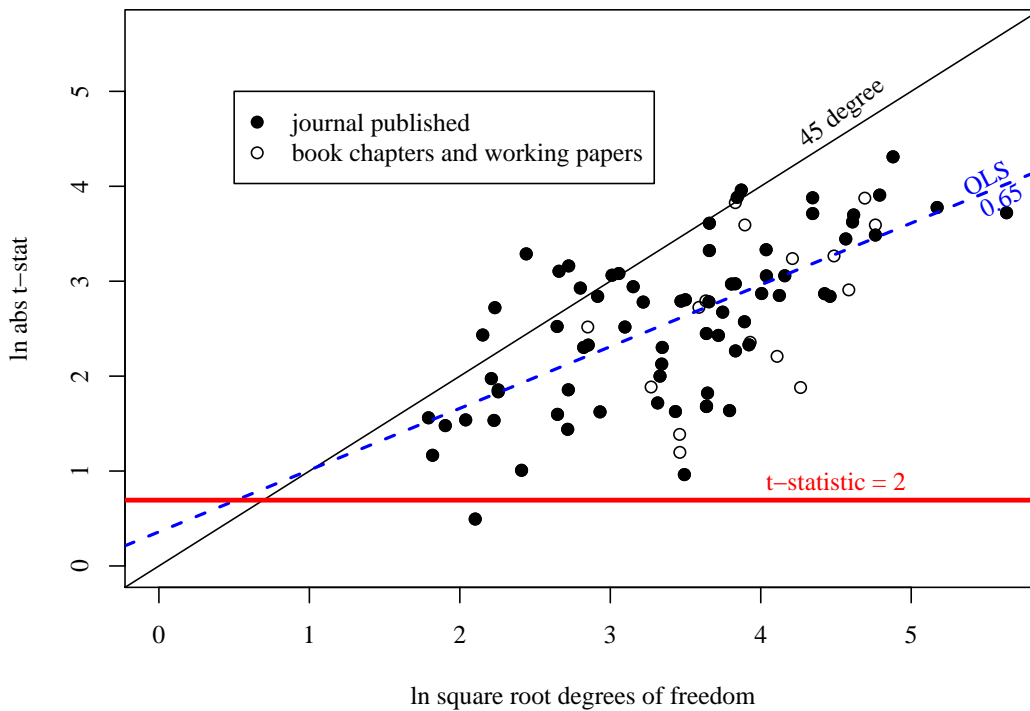


Figure 4: Absolute values of t -statistics rise with the square root of degrees of freedom, but not proportionately

The second method relates the distance coefficients to their standard errors. In the absence of publication bias, we have no reason to expect a relationship between the *strength* of the distance effect and the *precision* with which it is measured. However, if the best estimate from a study needs to be statistically significant at conventional levels to be published, then a process of specification searches might lead to coefficients that cluster at or slightly above twice their standard error.

Figure 5 provides a line corresponding to t -statistics of two. We see that most of the data

lie well above this line and that there is little apparent relationship between estimates and their standard errors. The correlation for published papers in journals is 0.06. This figure displays seven estimates that were not shown in figure 4 because the papers did not report degrees of freedom. Six of them are published in journals and the last one is a working paper. All of them except two have t -statistics over two. These results contrast sharply with the meta-analyses of Card and Krueger (1995) and Görg and Strobl (2001). The former observe that a line through the origin with a slope of two “fits the data rather well.” The latter estimate the relationship, obtaining a coefficient of about three. In our data the OLS regression on the published papers yields a slope of 0.43 with standard error of 0.88 and an R^2 of 0.003.

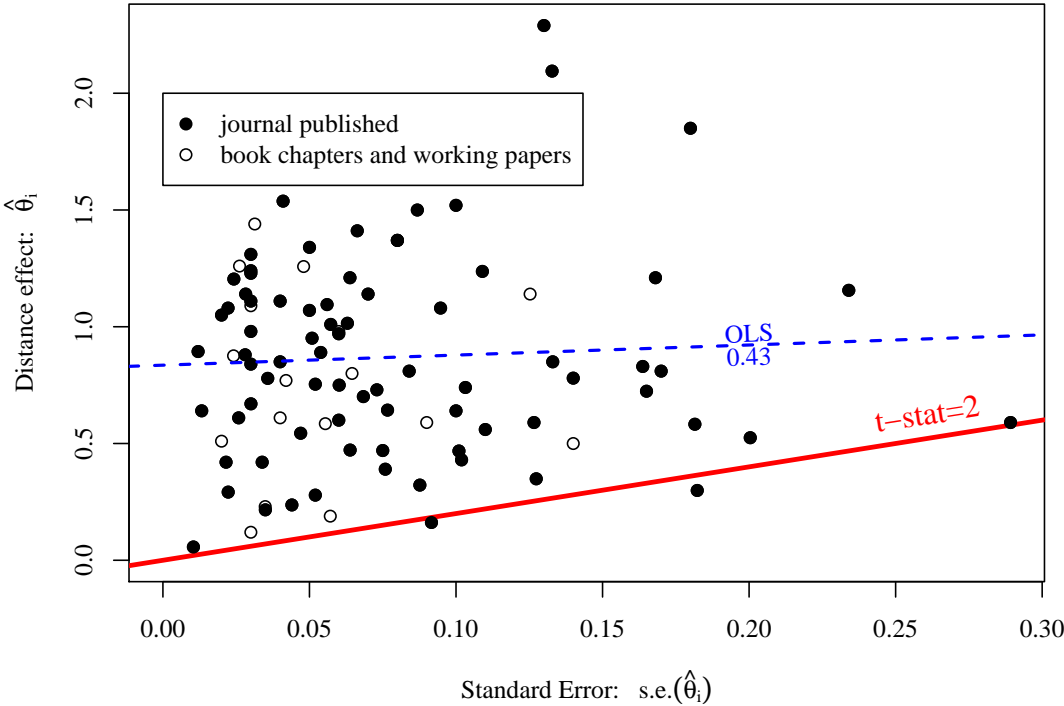


Figure 5: The (non) relationship between coefficients and standard errors

In contrast to earlier studies, we find only very weak evidence of publication bias. Perhaps this should not be very surprising since distance is usually just a control variable in studies based on the gravity equation. Hence, publication pressure may have less effect on the distance coefficient than it has on variables of substantial policy interest like minimum wages and MNC spillovers.

7 Conclusion

Leamer (2006) remarks that the distance effect on international commerce is “possibly the only important finding that has fully withstood the scrutiny of time and the onslaught of economic technique.” Our paper quantitatively supports this claim with a systematic analysis of 1467 estimates of the distance effect. We find a mean elasticity of 0.9, indicating that on average bilateral trade is nearly inversely proportionate to distance.

We explore the great variation in estimated distance effects and show that only 2% of it can be explained by mere sampling error. We attribute the remaining variation to heterogeneity in data sets and econometric methods. Meta-regressions show which differences have the most important impacts on estimated distance effects. One of the most significant explanatory variables is the time period of the data used in the estimation. Using estimates spanning well over a century, we show that distance effects decreased slightly between 1870 and 1950 and then began to rise. The use of a large number of “meta-variables” to control for relevant differences in the regressions producing our estimates does not cause a notable change in the increase in the distance effect. These findings represent a challenge for those who believe that technological change has revolutionized the world economy causing the impact of spatial separation to decline or disappear.

References

- Alvarez, F., and R. E. Lucas Jr., “General Equilibrium Analysis of the Eaton-Kortum Model of International Trade,” National Bureau of Economic Research working paper no. 11764 (2005).
- Anderson, J.E., and E. van Wincoop, “Gravity With Gravitas: A Solution to the Border Puzzle,” *American Economic Review* 93:1 (2003), 170–192.
- Anderson, J.E., and E. van Wincoop, “Trade Costs,” *Journal of Economic Literature* 42:3 (2004), 691-751.
- Baldwin, R., “The Euro’s Trade Effects,” European Central Bank working paper no. 594 (2006).
- Berthelon, M., and C.L. Freund, “On the Conservation of Distance in International Trade,” World Bank policy research working paper no. 3293 (2004).
- Blum, B., and A. Goldfarb, “Does the Internet defy the Law of Gravity?,” *Journal of International Economics* (2006), forthcoming.

- Brun, J-F., C. Carrère, J. de Melo, and P. Guillaumont, “Has Distance Died? Evidence from a Panel Gravity Model,” *World Bank Economic Review* 19:1 (2005), 99–120.
- Cairncross, F., *The Death of Distance: How the Communications Revolution will change our Lives* (London, U.K.: Orion Business Books, 1997).
- Card, D., and A.B. Krueger, “Time-Series Minimum-Wage Studies: a Meta-analysis,” *American Economic Review* 85:2 (1995), 238–243.
- Carrère C., and M. Schiff, “On the Geography of Trade: Distance is Alive and Well,” *Revue Economique* 56:6 (2005), 1249–1274.
- Chaney, T., “Distorted gravity: Heterogeneous Firms, Market Structure and the Geography of International Trade,” University of Chicago, manuscript (2006).
- Coe D.T., A. Subramanian, N.T. Tamirisa, and R. Bhavnani, “The Missing Globalization Puzzle,” International Monetary Fund working paper no. 171 (2002).
- Combes, P-P., T. Mayer, and J-F. Thisse, “*Economie géographique*,” Chapter 5 (2006). Available at <http://www.vcharite.univ-mrs.fr/PP/combes/Chapitre5.pdf>.
- Deardorff, A., “Time and Trade: The Role of Time in Determining the Structure and Effects of International Trade, with an Application to Japan” (pp. 63–76), in R.M. Stern (Ed.), *Japan’s Economic Recovery: Commercial Policy, Monetary Policy, and Corporate Governance* (Cheltenham, U.K. and Northampton, MA: Edward Elgar Publishing Inc., 2003).
- Duranton, G., and M. Storper, “Rising trade costs? Agglomeration and trade with endogenous transaction costs,” University of Toronto, manuscript (2006).
- Eaton, J., and S. Kortum, “Technology, Geography, and Trade,” *Econometrica* 70:5 (2002), 1741–1779.
- Feenstra, R.C., *Advanced International Trade*, (Princeton, NJ: Princeton University Press, 2004).
- Felbermayr, G., and W. Kohler, “Exploring the Intensive and Extensive Margins of World Trade,” *Review of World Economics*, forthcoming.
- Glaeser, E.L., and J.E. Kohlhase, “Cities, Regions and the Decline of Transport Costs,” *Papers in Regional Science* 83:1 (2004), 197–228.
- Glass, G.V., “Primary, Secondary, and Meta-Analysis of Research,” *Educational Researcher* 5:10 (1976), 3–8.
- Görg, H., and E. Strobl, “Multinational Companies and Productivity Spillovers: A Meta-Analysis,” *The Economic Journal* 111:475 (2001), 723–739.
- Grossman, G. M., “Comment” (pp. 29–31), in J.A. Frankel (Ed.), *The Regionalization of the World Economy* (Chicago, IL: The University of Chicago Press, National Bureau of Economic Research Project Report, 1998).
- Helpman, E., M. Melitz, and Y. Rubinstein, “Trading Partners and Trading Volumes,” Harvard University, manuscript (2006).

- Higgins, J.P., S.G. Thompson, D.J. Deeks, and D.G. Altman, "Measuring inconsistency in meta-analyses," *BMJ* 327:7414 (2003), 557–60.
- Hummels, D., "Toward a Geography of Trade Costs," Purdue University, manuscript (2001a).
- Hummels, D., "Time as a trade barrier," Purdue University, manuscript (2001b).
- Jarrell, S.B., and T.D. Stanley, "A Meta-Analysis of the Union-Nonunion Wage Gap," *Industrial and Labor Relations Review* 44:1 (1990), 54–67.
- Jeppesen, T., J.A. List, and H. Folmer, "Environmental Regulations and New Plant Location Decisions: Evidence from a Meta-Analysis," *Journal of Regional Science* 42:1 (2002), 19–49.
- Krugman, P., "The Move to Free Trade Zones" (pp. 7–41), in *Policy Implications of Trade and Currency Zones*, a symposium sponsored by the Federal Reserve Bank of Kansas City (1991).
- Leamer, E.E., "A Flat World, a Level Playing Field, a Small World After All, or None of the Above?" UCLA Anderson Forecast Book Review (2006).
- Leamer, E.E., and J. Levinsohn, "International Trade Theory: The Evidence" (pp. 1339–1394), in G. Grossman and K. Rogoff (Eds.), *Handbook of International Economics Volume 3* (Amsterdam: North-Holland, Elsevier, 1995).
- de Mooij, R., and S. Ederveen, "Taxation and foreign direct investment. A synthesis of empirical research," *International Tax and Public Finance* 10:6 (2003), 673–693.
- NIST/SEMATECH, *e-Handbook of Statistical Methods*, Chapter 1.3.5.17
 < <http://www.itl.nist.gov/div898/handbook/> >, September 2004 download.
- Nitsch, V., "Zipf Zipped," *Journal of Urban Economics* 57:1 (2005), 86–100.
- Overman H.G., S. Redding, and A.J. Venables, "The Economic Geography of Trade, Production, and Income: A Survey of Empirics" (pp. 353–387), in E. Kwan-Choi and J. Harrigan (Eds.), *Handbook of International Trade* (Oxford, U.K.: Blackwell Publishing, 2003).
- Rauch J.E., "Networks Versus Markets in International Trade," *Journal of International Economics*, 48:1 (1999), 7–35.
- Rose, A.K., and T.D. Stanley, "A Meta-Analysis of the Effect of Common Currencies on International Trade," *Journal of Economic Surveys* 19:3 (2005), 347–365.
- Santos Silva, J.M.C., and S. Tenreyro, "The Log of Gravity," *Review of Economics and Statistics* 88:4 (2006), forthcoming.
- Smith, M.L., and Glass, G.V., "Meta-analysis of psychotherapy outcome studies," *American Psychologist*, 32:9 (1977), 752–60.
- Stanley, T.D., "New Wine in Old Bottles: A Meta-Analysis of Ricardian Equivalence," *Southern Economic Journal* 64:3 (1998), 713–727.
- Stanley, T.D., "Wheat From Chaff: A Meta-Analysis as Quantitative Literature Review," *Journal of Economic Perspectives* 15:3 (2001), 131–150.

Stanley, T.D., and S.B. Jarrell, "Meta-Regression Analysis: A Quantitative Method of Literature Surveys," *Journal of Economic Surveys* 3:2 (1989), 161–170.

Stanley, T.D., and S.B. Jarrell, "Gender Wage Discrimination Bias? A Meta-Regression Analysis," *Journal of Human Resources* 33:4 (1998), 947–973.