‘Geospatial Restructuring of Industrial Trade’: modelling the impact of distance cost change on the UK’s spatial economic structure

keywords: spatial economics; regional economics; energy

November 2014

Contents

List of Figures

1 Introduction

In the UK, millions of business organisations rely on a web of spatial connections to other establishments. Each connection has its own cost to move goods and services across space. It is possible to theorise how a change in costs may affect a single firm, but system-wide effects also occur, which to be seen and understood require a detailed picture of the spatial structure of the UK economy as a whole.

This paper presents a spatial economic model of how changing distance costs may affect the structure of the UK economy, examining which industrial sectors and geographical zones are likely to be most affected. The 'GRIT method' (standing for 'Geospatial Restructuring of Industrial Trade') does this by linking two datasets. First, the Business Structure Database (BSD): this contains information on businesses covering 99% of the UKs total productive output. It describes turnover, employee numbers and location (**?**). It is a relatively new data that has been used successfully to produce novel analyses of the UK economy, including those on measuring business growth and its affect on employment and innovation over time (Anyadike-Danes *et al*. 2009; Mason *et al*. 2009), and an examination of agglomeration of UK production sectors (Simpson 2007). The BSD is managed by the Secure Data Service and must be accessed under secure conditions.

The second - much smaller - data source, the 'domestic use' input-output matrix, contains domestic trade flows describing intermediate demand between Standard-Industrial-Classification (SIC) coded sectors. The SIC code is used to link the two data sources: demand between sectors from the domestic use matrix is distributed geographically by linking to location and SIC from the BSD. A novel form of spatial interaction model is then used to estimate its geographical movement, before the model parameter-sweeps distance costs for the movement of goods to examine how they impact on sectors and geographies. Rather than use all sectors, GRIT concentrates on the 'heavier' sectors that tend to involve more movement of physical goods. The reason for this choice is explained in-depth in section 4.1 and returned to in the conclusion.

The GRIT method aims to demonstrate that linking these two datasets is a useful way to overcome the lack of good spatial economic data in the UK. Without this data, it is difficult to develop insights into the future spatial economy: what will happen as costs change and the UK undergoes an energy transition? Currently, views here are polarised. There are qualitative, 'common-sense' arguments that distance will become more expensive as energy costs increase, and thus localism must be a policy priority (e.g. Glasmeier 2007, North 2010). On the other hand, quantitative analyses all too often ignore the distance factor entirely (cf. Krugman 2010), treating transport as a normal production input, or minimising its importance since the cost of moving goods is a small percentage of overall costs (Glaeser and Kohlhase 2004). Logistics analyses tend to be (understandably) partial, comparing single firm-level supply chains. In one logistics perspective, the shift towards 'commonly held notions' about localism by consumers is seen as counterproductive, with the authors advocating continuing use of established efficiency methods: there is, they argue, a natural "joint positive relationship between economic and environmental efficiency: lean is green" (Oglethorpe and Heron 2010, p.553). Input-output models are able to investigate the multiplier effects of changing a single cost on the whole production matrix (for an energy-cost example see Kerschner and Hubacek 2009) but they do not model what will happen as distance costs between all business organisations change.

This is a particularly important problem to address at this stage in the UK's economic development. Over the coming decades, an energy revolution needs to take place to avoid the worst effects of climate change (Kramer and Haigh 2009, Jefferson 2008). Understanding the impact this revolution will have on the cost of goods and services is vital. Will changing costs mean a substantive spatial reconfiguration of the entire economy or will the changes prove less significant, relative to other economic shifts? As EU carbon emission costs are embedded further into the economy, how will changing production costs amplify tensions in the UK's industrial networks, and drive spatial relocation? This spatial restructuring of the economy has profound implications. These relate to, for example, the location of jobs, the spatial provision of services required by workers and their families, (such as housing and social services), and resource demand and environmental impacts of economic activity.

The GRIT method carries out the following stages. First, money flow amounts spent by each SIC sector from the domestic use matrix are distributed geographically, by linking to location of sectors from the BSD and using turnover per sector as a proxy for geographical spread. Second, these consumption amounts are 'spent' spatially using a form of budget-constrained spatial interaction model (detailed in section 4). This spend is repeated for a range of distance costs, calibrated to bound known values of goods delivery dropoff from Department for Transport data (see next section). This range of distance costs is swept from 'lower than current' to 'higher' and the change in demand amounts per sector and geographical zone is kept at each step. This is used to create rankings of the most affected sectors and geographical zones. Finally, those most consistently affected across all distance cost changes are presented in the results.

Thus, the model concentrates on identifying sectors and geographical zones that may face the most tension from increasing distance costs. It is not a dynamic model; no attempt is made to predict how those sectors or zones will need to adapt. Rather, the aim is to demonstrate the usefulness of this approach in moving away from purely qualitative analyses of the spatial consequences of future energy scenarios, as well as showing that distance is indeed a vital factor to consider when looking at the structure of the whole economy.

The paper is organised as follows. The following section begins by using Department for Transport freight survey data to show exactly how current trade of physical goods within the UK decays over distance. This analysis is then used as the basis for calibrating the range of distance cost change used in the model. Next, more information is given on the 'domestic use' input output data. The GRIT method itself is then explained, including a detailed explanation of how the BSD is used and how it is linked to the intermediate trade flow data. The following two sections outline how the GRIT method is applied to produce a ranking of likely most affected sectors and geographical zones, and how the calibration values for the model are arrived at, in particular how the DfT data is used to do this. The results are then presented, looking first at geographical effects and, second, sectoral effects. The paper finishes with a concluding discussion and suggestions for taking the model approach further.

2 The spatial decay assumption

The Department for Transport (DfT) publishes a series of data tables describing road freight activity for Great Britain[1]. This data comes from a regular rolling questionnaire that feeds into the 'continuing survey of road goods transport'. Hauliers provide detailed information for specific deliveries of vehicles over thirty tonnes; the DfT claims these account for around ninety percent of domestic freight movement. This is a very useful data source: input-output data may not necessarily represent physical goods movement, but there is no ambiguity with this DfT data - each surveyed vehicle is carrying physical material. The analysis in this section shows exactly how good movement decays over distance within Great Britain. This data is then used as the basis for calibrating the range of distance cost change in the GRIT model itself.

Two tables in particular provide a window on the spatial decay of trade in the UK, albeit a rather murky one; the detail in the original sampling is lost, keeping only a very aggregate picture. The first is an origin-destination flow matrix between the former Government Office Regions (GORs), as well as Scotland and Wales[2]. This describes the total tonnage lifted between all of these zones. The second[3] contains a commodity breakdown of goods lifted, with columns giving the number of tonnes lifted to a range of distances from their origin. The number of categories for both sector and distance are small: nineteen commodity groups and eight distance bins, ending in one containing all tonnage moved over 300 kilometres.

Both datasets confirm that tonnage drops off over distance - though with some quirks. Perfect distance decay would produce a perfect negative Spearman's rank correlation of minus one. Figure 1 shows the full set of Spearman's rank correlations for both datasets. Each commodity is ranked against distance, as is each GOR, before each is grouped into its own boxplot.

There are no perfect dropoffs of tonnage over distance, though all but two correlations are below -0.9. Looking more closely at the data, what else can be deduced? To make a comparison between both datasets (and within each Commodity and GOR) the data is adjusted so that tonnage values at a distance of zero are set to one and all others are set proportionally. This is done for each origin region and for each commodity group. The rationale behind this is to identify a rate of spatial decay for each commodity and GOR abstracted from specific tonnage values.

[1] https://www.gov.uk/government/collections/road-freight-domestic-and-international-statistics
[2] Table RFS0138, 'goods lifted by UK-registered HGVs by origin and destination of goods '.
[3] Table RFS0129, 'goods lifted by commodity and length of haul'.

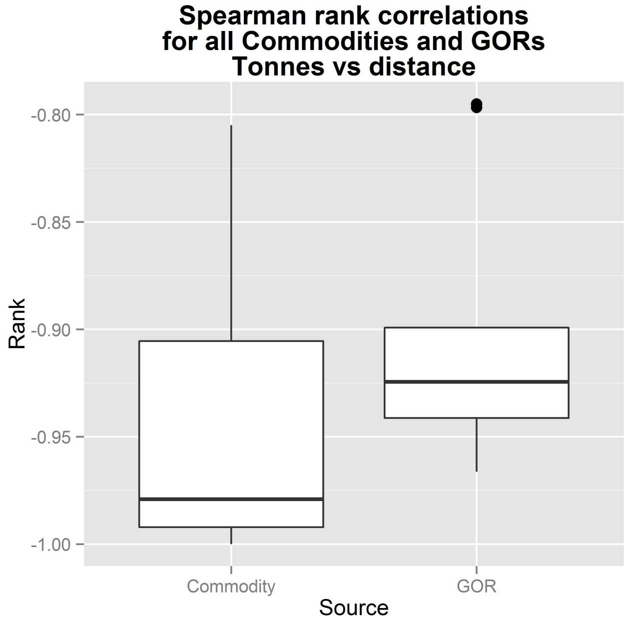

Figure 1: Boxplot: rank correlations on both sets of goods-vs-distance data, 'by Commodity' and 'by GOR'.

Distances are deduced from the origin/destination flows by assuming they move between centroids of the GORs. Only the GORs are used, not Wales or Scotland. The commodity data has the following adjustments made to it. First, its unequal sized distance bins are converted to equal gaps of 25km; values from larger bins are simply split between the smaller ones. It is then assumed that tonnage moved (for example) 'between zero and 25km' ends up at half that distance, 12.5km. This leaves no value at zero to set to one, so the final adjustment is to use a 'loess' local fit model to interpolate where the commodity data would cross the zero line. This value can then be used to correctly adjust the rest to that interpolated unit value.

Figure 2 shows these adjusted values. By-commodity dropoff rates are broken between two sub-figures. Sub-figure 2a shows commodities that drop off after the first datapoint; sub-figure 2b picks out the mentioned 'quirk' in this data - a significant number of commodities deliver more tonnage to the second distance bin before dropping off. These tend to be the heavier goods (though not exclusively). The 'continuing survey of road goods transport' only surveys heavy goods vehicles over 3.5 tonnes. While it covers ninety percent of freight, it is possible that sub-3.5 tonne movements are predominantly over shorter distances.

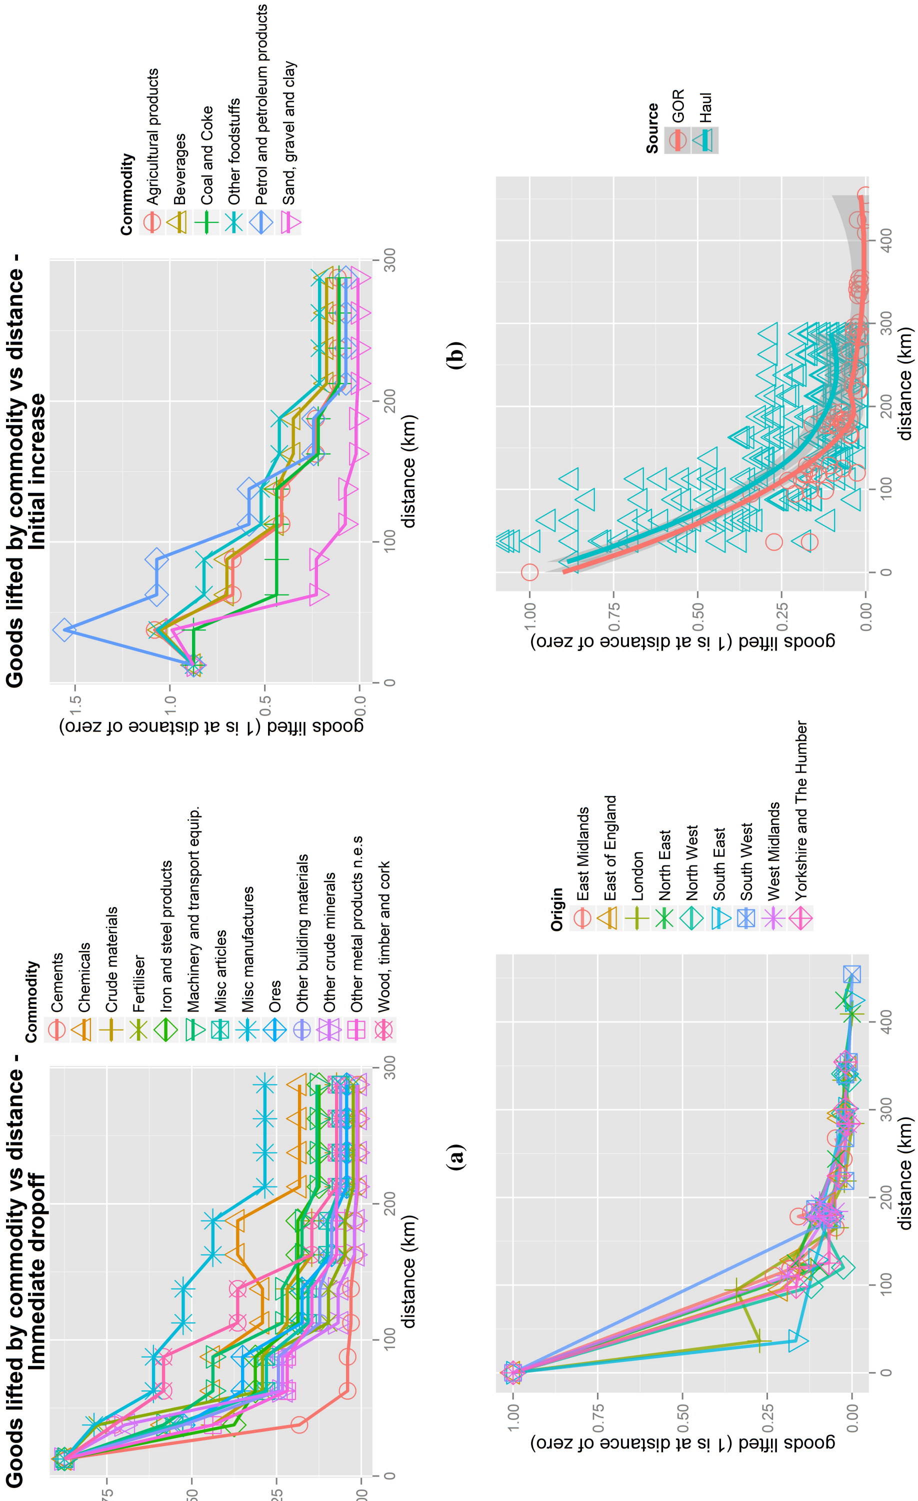

Figure 2: Distance decay in DfT datasets giving 'by-commodity' and 'by-GOR' rates of tonne delivery.

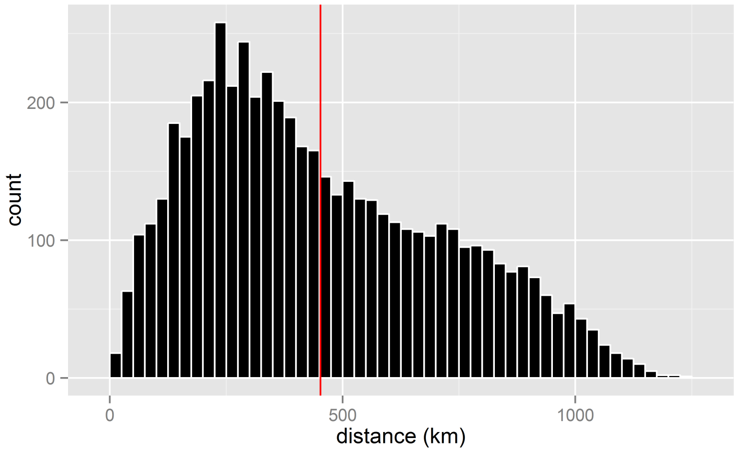

Figure 3: Histogram of the length of road network routes within Great Britain, randomised start and end points, 5500 routes in total. Routes sampled from Google's distance matrix API. Red line is mean value.

Sub-figure 2c shows the dropoff rate for GORs. Despite its variability, there is a consistent rate - though given all of this data is domestic only, there are of course physical limits to how far a delivery can go within England. The final sub-figure (2d) compares both data sources: a locally fitted loess curve is applied to both, replacing the variability of commodities with a single line. Though not a perfect match, the decay rates are close - at least up until 200km. Both are derived from the same survey, though commodity details come from a specific question about distance travelled, where the GOR data is deduced from point of origin and destination, so the sources of distance are different. It is difficult to make any firm conclusions here: the commodities data bin runs from 200 to 300km. As mentioned, this has been spread out evenly over four 25km bins, but the actual dropoff rate is unknown - and cannot be deduced from the final 'more than 300km' figure.

Nevertheless, it is clear that strong spatial decay exists in the movement of domestic freight. An exact null hypothesis for intra-UK trade is tricky to settle on, but assuming it is 'each trade happens between random locations within the island of Great Britain', using the road network, distance distributions look like figure 3[4]. The island's physical limits are clearly imposed but it is far from the clear spatial decay signal from the tonnage data.

As well as showing current spatial decay in the UK, this section's analysis provides the basis for calibrating the GRIT model - the method for doing this is explained in section 6.

[4]Random routes sampled via Google's distance matrix API. See https://github.com/DanOlner/randomNetworkDistancer for code.

3 The intermediate trade flow matrix

The Office of National Statistics (ONS) 'supply and use' data contains a 'domestic use' input-output matrix recording trade flows between 106 sectors in the UK[5]. It does not include imports - only goods and services moving within the UK are included. The domestic use matrix provides a picture of which sectors money moves between during productive activity in the UK. The purpose of the GRIT method is to link these non-spatial flows to the Business Structure Database in order to estimate their geographical origins and destinations, before modelling how they may change as distance costs change[6].

In this paper two sectors are dropped, leaving 104. These are 'services/activities of households as employers of domestic personnel', as this contains no values, and 'imputed rent' - the amount of money imputed to enter into the economy from housing, even if no money is actually exchanged. Out of the resulting 104-squared matrix of possible flows, there are just under nine thousand money flows going from one sector to another. At the time of writing, 2010 is the most recent year for which the domestic use matrix is available. Figure **??** bins these money flows on a log axis to give a picture of their variance.

There are some caveats to using the domestic use matrix at face value. The supply and use data is part of a larger process of balancing different sources against each other to arrive at a consistent gross domestic product (GDP) figure. The supply and use data itself builds on the expertise of specific teams dealing with sets of industries, whose job it is to create a final balanced outcome (see Akers and Clifton-Fearnside 2008 pp.16 for a detailed breakdown of the processing stages and final automated balancing). Finally, it must go through a process of 'balancing and adjustment' (**?**) against the two other legs of the process (income and expenditure) so that all GDP estimates agree.

While the domestic use data is in some ways opaque, then, it is nevertheless an expertly assembled set of numbers. For many of the obviously 'heavy' sectors, it is assumed here that they represent reality well enough to be used for thinking about spatial flows.

Input-output matrices always exactly balance consumption and demand. Each cell is both one sector's spending and another's sales; the former is summed in columns, the latter in rows, so each is double counted to make total consumption and demand identical. However, money changes *distribution between sectors* as it is spent. In the domestic use matrix, the variance of consumption across sectors is higher than the variance of their demand. For the purposes of this paper, this is worth bearing in mind as it has a spatial implication: if spending moves towards lower sectoral variance, it will also become more geographically concentrated.

The domestic use matrix contains 'basic prices' only: "the amount received by the producer

[5]There were 110 sectors in the pre-2013 classification of sectors used in the combined use matrix. From 2013, a number of sectors were consolidated. Most notably, 'construction' became the single largest sector, replacing 'buildings and building construction works', 'construction works for civil engineering' and 'specialised construction works'.

[6]The domestic use matrix is available for download at http://www.ons.gov.uk/ons/rel/input-output/input-output-analytical-tables/2010/index.html.

for a unit of goods or services, minus any taxes payable, and plus any subsidy receivable on that unit" (Akers and Clifton-Fearnside 2008 p.18). Transport costs are, in theory, excluded from these, though they are part of the 'purchaser's price' found in the combined use matrix. They are, however, included if certain sectors buy from the transport sector described in the matrix itself - but the method used to determine how this is assigned is not transparent. In the accounting methods used, there are in fact six different ways to arrange transport costs between basic, producers' and purchasers' prices (**?** p.277), depending on whether firms have in-house freight and a range of other factors.

For the purposes of this paper, a spatial decay function is used to estimate how far goods are delivered over distance within the UK. In order to estimate what impact increasing distance cost has, a sensible range of distance cost values that bound likely values are are swept. Within the set range, plausible transport cost changes should be captured, thus avoiding the need to know in exact detail how the proportion of transport costs breaks down.

4 The GRIT method

Each cell in the domestic use matrix is a number describing the flow of money from one sector (columns) to another (rows). The task of the GRIT method is to take each of these from/to flows and examine how that spending changes spatially for a range of distance costs. The modelling process itself is done in two key stages. Firstly, money 'from' sectors is taken from the (aspatial) domestic use matrix and distributed geographically across the UK, using the method described below. Secondly, this 'from' money (now distributed into geographical zones) is 'spent' on the appropriate 'to' sector across the UK, using a constrained spatial decay function (see below).

Each geographical zone contains up to 104 sectors - though not all zones have a presence from all sectors. In order to keep terminology consistent, sectors in geographical zones will be referred to as a'sector/zone'- for example, 'fabricated metal products' in Wolverhampton is one sector/zone.

The same sectoral SIC coding is used in the both the IO matrix and the BSD. The BSD actually uses five digit SIC codes, categorising businesses into around 700 sectors. In order to create the link with the IO matrix, these are binned into the cruder three-digit categories used in the supply-use data. This link between the two datasets provides the way into estimating where sectoral spending takes place in the UK.

The BSD has records for the vast majority of individual UK businesses, covering around ninety-nine percent of total turnover. The GRIT method relies on two key fields from the BSD for each business unit: their location and their turnover. These two give a way to slice up the non-spatial cell amounts in the domestic use matrix - using turnover as a proxy for the amount sectors spend. 'Travel to Work Area' (TTWA) zones are used - there are 243 of these in the UK, each designed to capture distinct commute zones around centres of employment (Coombes and Bond 2007). This makes them ideal for examining geographical trade between sectors within the UK.

The process of creating the proxy is made a little trickier by the nature of business units in the BSD. These are broken down into ‘enterprise units’ and ‘local units’. The former are central offices and HQs - the head of businesses. The vast majority of records in the BSD are single-unit - these make up about 98

When thinking about the geographical spread of spending by sectors, then, local units are important - they account for around 60

Once firm turnover is split between its local units, turnover is then summed by sector/zone. Since the job of the turnover proxy is only to slice up spending amounts from the domestic use matrix cells, the absolute amounts are not needed here (though absolute turnover values are used in the distance cost calculation; see below). Each matrix cell is split between all TTWAs (or rather, all with a presence from each specific sector) - so all that is kept is a fraction for each TTWA, summing to one across all zones for each SIC. For example, all ‘fabricated metal products’ sector/zones’ turnover proxies sum to one. Each matrix cell can then simply be multiplied by this fraction, giving the estimated money amount that each sector/zone combination will spend across the UK.

Each sector/zone now has an amount it spends on other sectors, set proportional to turnover, constrained so that the total across all zones equals the cell figures in the matrix. Moving on to the second step, the next task is to make each ‘from’ sector/zone buy from selling sectors. To apportion this spend, each sector/zone is treated as a separate economic ‘representative agent with a budget for each other sector they buy from. A budget constraint can then be used to allot that agents spend across the rest of the UK. A Constant Elasticity of Substitution (CES) function is used for this purpose. The CES approach has been used successfully at a more aggregate level to produce analyses of international trade costs (see e.g. Anderson and van Wincoop 2003; Anderson and Wincoop 2004). It allows allocation of demand that easily keeps them within a set ‘budget’ for each sector/zone, and the importance of distance can be controlled through its elasticity parameter and the cost of distance itself (Olner 2013).

The constrained CES function used is given in equation (1):

$$g_j = P_j \frac{P_j^{\frac{1}{\rho-1}} Y}{\sum_{i=1}^{n} P_i^{\frac{\rho}{\rho-1}}} \tag{1}$$

The output g_j is the amount of money received by a single sector/zone. P_j is the cost of buying from that sector; P_i is the cost of all other zone/sector combinations for the SIC being bought from. The CES function constrains the amount being spent (Y) between these. ρ is the elasticity of substitution parameter; because distant buying is more expensive, increasing this has the effect of making the spending more localised. The first P_j - the cost of a unit of output from the selling sector/zone - multiplies the CES output to give the actual money amount flowing to the seller.

The ‘agent’ spending its allocated slice of money faces the cost of each sector/zone combination it is buying from. This cost - P (P_j and P_i in the equation) - must reflect both the relative size of the sector/zone combination (larger ones will receive more money) and the distance between the zones.

The CES function is used solely to find out where the money amount (Y) being spent by one sector/zone ends up. As mentioned, to get these individual money amounts for each sector/zone that receives money, each optimal output (g_j) must be multiplied by the per-unit 'price' to read off to the actual amount spent. The full set of spending thus worked out keeps to the constraint, always summing to Y.

The job of the 'price' itself in this case is two-fold. Firstly, it must reflect that (ceteris paribus) zones with a larger sector presence will receive more money than smaller ones. Secondly, that (ceteris paribus) more distant zones will be less attractive than nearby ones. To make zones with a larger sector presence more attractive, their 'price' has to be cheaper (so more money is spent on it). Turnover amounts per sector/zone is used again for this purpose - though this time, the absolute amounts per sector/zone are used rather than the normalised value. This is because normalising to one affects the demand size of individual sector zones; using the absolute amounts keeps the effect of the full range of sector/zone sizes.

Sector/zones with larger turnovers are subtracted from an 'inverse' to make the 'price' lower for larger sectors. This inverse affects the spatial decay drop-off curve; it has been set to 1.5 times the largest sector/zone turnover amount. (It is larger than this largest value so that the smallest inverse value is still some way larger than zero; a price of zero would make demand infinite for buyers in the same zone.) Note these are not actual prices for goods: it is just a method for using the CES function as a tool for distributing spending spatially.

The 'price' is responsible for allocating the total amount of money spent by a particular sector/zone on those it buys from. As such, as well as the inverse base price representing relative sector size, the distance cost must be proportional to the total consumption taking place between each sector/zone trade. This total is meant to represent all the individual trades taking place over the year, each of which is presumed to have a fixed proportion of distance costs. Though in reality, goods moving between zones would of course have a range of types and distance costs, it is presumed this approach can do a decent job of representing the average aggregate outcome. The 'price' for buying from a particular sector/zone is thus:

$$P = \alpha(i - t) + \beta(d * t) \tag{2}$$

- where t is the sector/zone turnover amount being bought from, i is the sector/zone inverse, d is the distance between buying and selling zones, and α and β are controlling parameters.

The CES function has some useful properties that make it particularly suitable for this kind of spatial job. Primarily, the *rate of spatial decay* is unaffected by the number and position of selling zones. The addition of further zones will affect the distribution of spending of course - but the slope of spatial decay of both spending and goods movement remains unaltered. Thus, ρ controls spatial decay consistently, regardless of the position and number of zones being bought from. It is an open question whether, in reality, the addition of further sites of demand might alter spatial decay but, for use in estimating trade flows, this level of control is useful.

4.1 How good is the proxy?

How effective is turnover as a proxy for spending by sectors in different geographical zones? An imperfect way to check the validity of the proxy is to look at the correlation between total turnover per SIC sector, summed from the BSD, and total consumption per SIC in the domestic use matrix (the total amount each sector buys from all other sectors). Checking against consumption is important: the proxy is being asked to stand in for consumption that takes place by individual sectors in geographical zones of varying sizes (where 'size' is determined by turnover proxy).

Table 1 and 2 show the R^2 values for a set of four correlations, looking firstly at all sectors and then at a subset of the 'heavier' sectors more likely to be dominated by the physical movement of goods. This is the first 58 sectors up to 'construction'. For each, turnover per SIC is summed for all enterprise units - that is, using only the turnover amounts connected to the main business unit, using its SIC code. Turnover is also summed per SIC using the local unit turnover split method described above. Since local units often have different SIC codes to their parent firm, this results in slightly different turnover totals per SIC. The tables also include correlations for both consumption of and demand for each SIC, for comparison.

As the tables show, the strongest correlation ($R^2 = 0.956$) is between local unit turnover totals and consumption by the subset of heavier sectors. Including all sectors drops the same correlation to 0.684. Using local unit totals does better, then - a useful result, given that the extra geographical information the local units provide is a much better match for estimating trade flows. One might expect the correlation between turnover and demand for a sector to be higher than to its consumption, given that demand is the money being received by that sector; as it transpires, this is not so. This is perhaps surprising but also good for the method, since the proxy is being asked to stand in for consumption.

proxy	demand	consumption
Local unit turnover per SIC	0.530811	0.68365
Enterprise turnover per SIC	0.52507	0.667903

Table 1: Proxy correlations, R^2, all sectors

As mentioned, this is far from a perfect comparison. The correlation shows a strong relationship between *per sector* consumption and turnover. However, for this to be a direct comparison, the same relationship would need to hold *within* sectors distributed across each geographical zone. While it is not an ideal comparison, then, this correlation does at least suggest a very strong relationship between turnover levels and sectoral consumption generally.

proxy	demand	consumption
Local unit turnover per SIC	0.926433	0.956081
Enterprise turnover per SIC	0.918931	0.940587

Table 2: Proxy correlations, R^2, 'heavy' bottom 58 sectors, up to 'construction'.

5 Applying the method

How is the above method applied to produce results? It begins by using the proxy approach to distribute domestic use matrix money between all 'from' sector/zones. The CES approach is then used to spend each of those amounts across the UK. This spending is summed to produce the end result: which 'to' sector/zones receive money.

A parameter-sweep approach is used to get the final results: the distribution of spending is repeated for a range of spatial decay values. The rationale is that, rather than attempt to calibrate exactly to a specific spatial decay function, results can be produced from a parameter range that bounds likely values. Other values are fixed (see below) while beta - the spatial decay parameter - is swept from one to twenty in increments of one.

By then looking at the *change* in money flows between parameter steps, it is possible to identify which sectors and zones are most affected by a change in the cost of moving goods across distance. Those most consistently affected, regardless of the actual parameter value, can be picked out of the results.

For the twenty parameter steps taken, there are nineteen differences - i.e. the change in money flow between the first and second parameter value, through to the change between the nineteenth and twentieth. These differences are the basis for the analysis - how much does demand change for sector/zones as distance costs change?

The results presented here concentrate on the percentage differences between distance cost changes to make comparison across sectors and zones more meaningful. For each difference, affected sectors/zones are ranked, from the largest negative to the largest positive difference for each increase in distance costs. It is then possible to observe how the rankings change with each parameter step. Over the full nineteen parameter steps, those sectors and geographical zones most affected become apparent.

This ranking approach is particularly valuable because the sector/zones' response to parameter changes is not linear. Comparing the rank of how affected each sector or zone is between parameter steps keeps the focus on change of sector/zones relative to each other, making comparisons of distance cost affects more clear.

The results look firstly at the most affected geographical zones and, second, the most affected sectors. Each requires a slightly different approach to pulling the numbers out, but they all concentrate on the percentage change in demand between parameter steps.

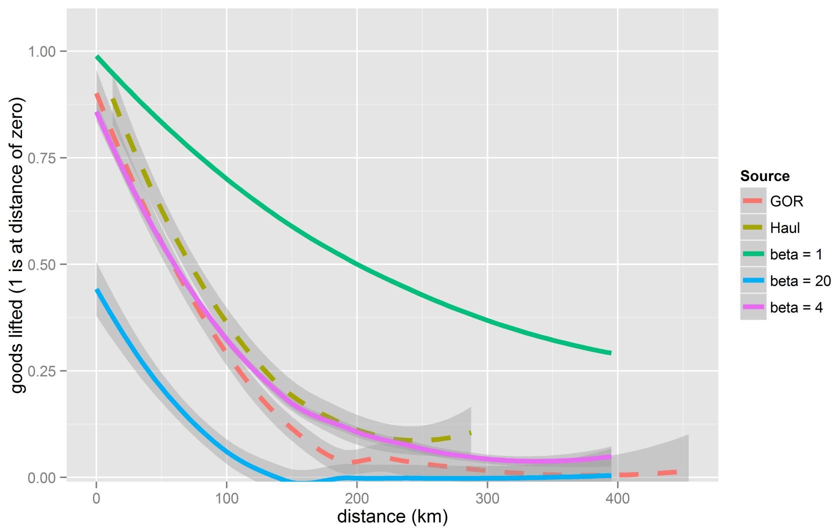

Figure 4: Comparison of Department for Transport 'goods lifted' over distance to the CES model of spatial decay. Bounding values for β are shown (1 and 20) as well as $\beta = 4$, the closest match to data.

6 Selecting calibration values

The method uses spatial decay deduced from Department of Transport tonnage data (section 2) to calibrate. Two different quantities are being compared, however: the tonnage decay rate over distance from the DfT data needs to be compared to *money* decay rate, as the GRIT method models how money flows drop off over distance. The CES function outputs quantity of good bought. The specific units used are not important - the normalised rate of spatial decay can be used to compare model and data.

As explained in the previous section, rather than attempt to exactly calibrate to each good type, the GRIT method instead bounds likely values on each side of the spatial decay rate and parameter sweeps that full range - thus going from low to high distance costs. The full range of calibration values used in the model is shown in figure 4. This compares spatial decay deduced from DfT data (as also shown in figure 2d) and, using a loess curve fit again, compares this to the lowest and highest distance cost values from the model curve, as well as a central value that closely matches the data.

The parameters used in the model are as follows. ρ is set at 0.7. The base cost multiplier α is kept to 1, while the distance cost multiplier β is parameter-swept from 1 to 20, to produce the result range used in the following section. As mentioned, the 'sector inverse' value used to make larger sectors receive more money is set to 1.5 times the value of the largest sector turnover value.

7 Results

7.1 Geographical effects

This section looks at the effects of changing distance costs in the GRIT method on the demand for intermediate goods from specific TTWAs. The analysis concentrates on the ranking of each TTWA, as outlined above. To recap, the amount each TTWA is affected by a change in distance costs is recorded - so e.g. the change between $\beta = 1$ and $\beta = 2$ represents an increase in distance costs. If a TTWA sees a drop in demand for its goods as distance costs increase, this negative change is recorded. The percentage change compared to the previous absolute demand amount is then found for each TTWA. Finally, all TTWAs are ranked, with number 1 being the most negatively affected and 243 being the most positively affected. This ranking process is repeated for each difference between parameter steps.

Figure 5 gives an overview of the result for all TTWAs. This is done by taking an average of each TTWAs rank position across all nineteen differences. By using a straightforward average, if a TTWA drops into a lower ranking, this will push its average up. Those that stay most affected across all parameter sweep steps also become apparent. A number of coastal zones are in the list of most negatively affected (those with an average ranking of 39 or higher). Going anticlockwise from Cornwall, these are: Penzance/Isles of Scilly; Bude/Holsworthy; Bideford; Chelmsford/Braintree; Cromer/Sheringham; a line of zones along the Lincolnshire Cambridgeshire coast (Wisbech, South Holland, Boston, Louth/Horncastle, Scunthorpe); Hull; Scarborough, Middlesborough/Stockton; Newcastle/Durham; Dunfermline and Ayr/Lilmarnock in Scotland; Belfast and Ballymena in Northern Ireland. Inland zones are: a band up through the Welsh border, including Newton/Welshpool; Shrewsbury; Crewe/Northwich and Preston, and finally Derby and Reading/Bracknell.

Figure 6 looks in more detail at ten of the most negatively affected TTWAs, selected based on their average ranking. They are ordered top to bottom in the key: Belfast is most affected, staying within the top three most affected for all parameter steps (and in the top two for all but the first difference). As the figure shows, using the average rank also includes a number of TTWAs that begin relatively unaffected between steps at lower distance costs but rapidly join the other most affected zones. Each of these four is a quite geographically isolated zone: Scarborough, Newton in the centre of Wales, Penzance and South Holland in Lincolnshire.

7.2 Sector effects

This section looks at the SIC sectors most affected by changes in distance costs. As with the analysis of TTWAs, finding the ‘most affected’ is done in stages. First, sectors are ranked, most to least affected in terms of percentage demand change, on each step between distance cost increases. Second, the average of those ranks for each sector across all nineteen parameter change differences is found, with the lowest average as a proxy for ‘most affected’.

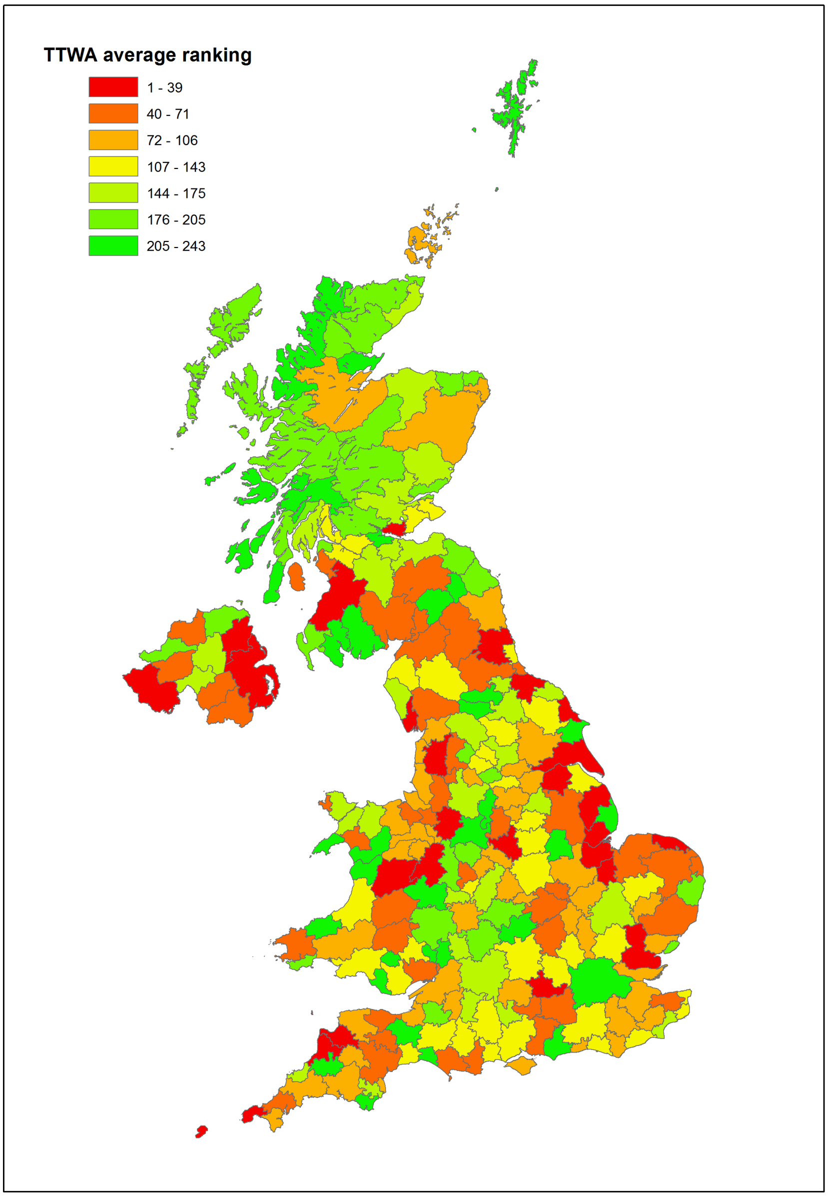

Figure 5: Most affected zones: Travel to Work Areas, average rank over whole parameter sweep.

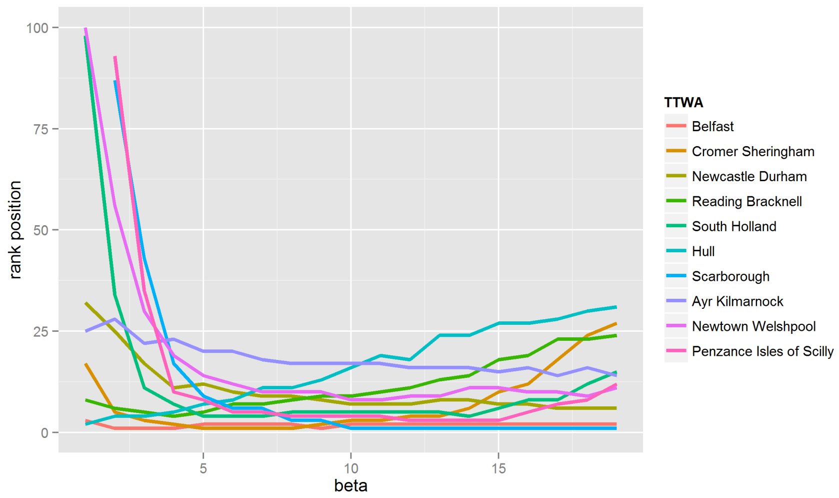

Figure 6: Ten most negatively affected Travel to Work Areas, all rank positions. Red is most negatively affected.

Analysis of sectors is complicated by there being a set amount of money demand going to each sector. The CES method for distributing money amounts must keep to those values from the domestic use matrix - changes in parameter values can only redistribute that fixed quantity of money. The method can only reassign it to different zones - any negative change is balanced somewhere else geographically by a positive. Negative and positive percentage changes, then, can only be read as ‘most negatively affected’ or ‘least negatively affected’ - as a way of ordering the expected sectoral impact.

The analysis separates out those negatively affected sectors (those with only a negative percentage change for an increase in distance costs) and ranks the sectors in order of most negatively affected, with the biggest drop in demand.

Table 3 lists the top ten most affected sectors in order. It includes their average rank position for each parameter step (rounded to the nearest integer), as well as their standard deviation from this rank over the nineteen steps, to give an indication of how far they move from that rank position over the whole parameter sweep (they are quite stable for all changes in distance costs).

The common themes in the top ten are heavy goods (those with very low value density that do not travel far, such as quarry rock and cement), related industries (all construction) and some of the main utilities, including water and gas. There are also two industrial production categories (fabricated metal products and machinery). Note, however, that agriculture is just outside the top ten.

SIC	Sector	Av rank	S.D.
8	Other mining and quarrying products	2	1.27
23.5-6	Cement, lime, plaster, concrete	3	1.55
41, 42 , 43	Construction	3	1.71
38	Waste collection, treatment, disposal; materials recovery	3	0.6
35.1	Electricity transmission and distribution	4	1.87
37	Sewerage	6	1.19
36	Natural water; water treatment and supply services	7	0.5
25OTHER	Fabricated metal products	9	1.29
28	Machinery and equipment	10	0.77
35.2-3	Gas; distribution through mains; steam/air con supply	10	2.22

Table 3: Ten most negatively affected sectors. Average and standard deviation of rank position over the distance cost sweep.

References

Akers, R., Clifton-Fearnside, A. (2008). 'Blue Book 2008: Balanced Estimates of Gross Domestic Product using a Supply and Use Approach'.

Anderson, J., van Wincoop, E. (2003). 'Gravity with Gravitas: A Solution to the Border Puzzle'. *American Economic Review* **93**(1):170–192.

Anderson, J. E., Wincoop, E. v. (2004). 'Trade Costs'. *Journal of Economic Literature* **42**(3):691–751.

Anyadike-Danes, M., Bonner, K., Hart, M., Mason, C. (2009). 'Measuring business growth: high growth firms and their contribution to employment in the UK'. *Tech. rep.* URL: http://strathprints.strath.ac.uk/16124/

Coombes, M., Bond, S. (2007). 'Travel-to-Work Areas: the 2007 review'. *Tech. rep.*, Office for National Statistics.

Glaeser, E., Kohlhase, J. (2004). 'Cities, regions and the decline of transport costs'. *Papers in Regional Science* **83**(1):197–228.

Glasmeier, A. (2007). 'Book Review: Worlds of Food'. *Regional Studies* **41**(5):713–718.

Jefferson, M. (2008). 'Accelerating the transition to sustainable energy systems'. *Energy Policy* **36**(11):41164125.

Kerschner, C., Hubacek, K. (2009). 'Assessing the suitability of Input-Output analysis for enhancing our understanding of potential effects of Peak-Oil'. *Energy* **34**(10):1662–1668.

Kramer, G. J., Haigh, M. (2009). 'No quick switch to low-carbon energy'. *Nature* **462**(7273):568–569.

Krugman, P. (2010). 'The New Economic Geography, Now Middle-Aged'. URL: http://www.princeton.edu/ pkrugman/aag.pdf

Mason, G., Bishop, K., Robinson, C. (2009). 'Business Growth and Innovation. The wider impact of rapidly-growing firms in UK city-regions. Research report October 2009'. *NESTA. London. UK* .

North, P. (2010). 'Eco-localisation as a progressive response to peak oil and climate change A sympathetic critique'. *Geoforum* **41**(4):585–594.

Oglethorpe, D., Heron, G. (2010). 'Sensible operational choices for the climate change agenda'. *International Journal of Logistics Management* **21**(3):538–557.

Olner, D. (2013). *An agent-based modelling approach to spatial economic theory*. Ph.D. thesis, University of Leeds, UK.

Simpson, H. (2007). 'An Analysis of Industrial Clustering in Great Britain'. *Final Report, The Institute for Fiscal Studies, London. In Search of Competitiveness Policies* **27**:5–35.