Exploring the Distribution of Electricity Consumption in the United Kingdom using Statistical Mechanics

Dan Chester

This thesis is submitted in part fulfilment of the requirements for the M.Sci. in Physical Geography degree at Lancaster University.

January 2019

Exploring the Distribution of Electricity Consumption in the United Kingdom using Statistical Mechanics

Abstract: New data for electricity consumption per capita at a resolution of approximately one square kilometre are studied in order to describe the distribution across the United Kingdom and trends between urban and rural areas. A two-class distribution is found, with around 99% of the data in the lower class, characterised by an exponential Boltzmann-Gibbs distribution, and 1% in the upper class, demonstrating a Pareto power-law distribution. This is shown to apply even at the smallest spatial scales, confirming the econophysics literature which predicts that this distribution will invariably apply to economic data as a result of entropy maximisation under the second law of thermodynamics. Electricity consumption per capita is found to be systematically higher in rural areas and lower in urban areas in the United Kingdom, with each containing similar amounts of inequality. London, however, is shown to be considerably more unequal, and a much higher proportion of the data (15%) is found to belong to the upper class.

Keywords: electricity consumption, distribution, inequality, econophysics, statistical mechanics, Yakovenko

Word count: 9791

Contents

Section 1

Introduction

Climate change, energy security and inequality are among the greatest challenges of our time. This makes the consumption of electricity and its distribution a topic of vital importance: electricity consumption is one of the key drivers of climate change; its distribution describes access to and demand for energy resources, a key determinant of quality of life and an increasingly contested commodity; and is inextricably linked to inequality in income and wealth. It also provides a valuable metric for measuring economic activity and how it varies in space and time, which is imperative to informing policy both on energy and the wider economy.

Emerging theories from the field of econophysics discuss the possibility that the distribution of money in the economy is universal across all societies and tends towards a Boltzmann-Gibbs probability distribution, maximising entropy through pairwise interactions in the same way as the energy of atoms in a gas in a closed system (Dragulescu and Yakovenko, 2000). It has been shown that this distribution fits the lower 95-99% of society, while the upper section, which accumulates its wealth from capital rather than wages, demonstrates a Pareto power-law distribution (Dragulescu and Yakovenko, 2001b). These have been extensively tested against income and wealth data, but have only recently been applied to energy consumption, and rarely below a national level (Banerjee and Yakovenko, 2010; Lawrence et al., 2013) or to just electricity consumption. Access has been granted to new data mapping global electricity consumption at the scale of approximately one square kilometre (Jarvis, 2018), allowing these methods from statistical physics to be applied to electricity consumption on an unprecedented scale, and for the trends and inequalities in the distribution of electricity consumption in the United Kingdom to be studied at a high level of detail.

1.1 Measuring and understanding inequality

The persistence of economic inequality in most societies has been perhaps one of the most studied economic and public policy phenomena over at least the last few centuries, and the desire to understand it empirically and mathematically has fuelled the ever-growing field of econometrics, which applies statistical methods to economic relationships. Methods for examining distributions of economic data, including, but rarely applied to, electricity consumption, have developed into conventional forms which will be discussed here, but the theory behind what causes these distributions remains hotly contested.

To solve some of this contestation, the relatively recent field of econophysics goes further and applies models and methods from physics to explain economic phenomena. The term was coined by Stanley et al. (1996) who discussed "the possibility that behavior of large numbers of humans (as measured, e.g., by economic indices) might conform to analogs of the scaling laws that have proved useful in describing systems composed of large numbers of inanimate objects". Theories in this field seek to show that distributions of economic data, such as income, wealth or electricity consumption, are determined by factors inherent to physics.

1.1.1 Early models of income distribution: Pareto and Lorenz

An early pioneer in the mathematics of inequality was Vilfredo Pareto. Amid rising concerns about inequality and the spread of Marxist thinking in the late nineteenth and early twentieth century, Pareto began experimenting with mathematical models to gain a better understanding of real economic data, starting with tax records from European cities and states (Rodd, 1995). In his first work, Pareto (1897) showed that the distribution of wealth was roughly the same in all countries and societies. Plotting various cumulative income distributions on a log-log scale, he observed that each distribution gave a straight line with roughly the same gradient regardless of the study area (Fig. 1.1). He concluded that the shape of income distribution is deterministic and can be summarised as an inverse power law: the cumulative probability of a person having income or wealth of at least w varies as $P(w) \propto w^{-\alpha}$, where α is a constant gradient of the log-log graph (Fig. 1.1).

An early critic of this presentation of this data was Lorenz (1905), who claimed that logarithmic curves are misleading as they can easily be mistaken as linear by casual observers, and fail to take account of changes in population. He proposed the now-widespread Lorenz curve (Fig. 1.2) with axes displaying cumulative amounts of the population and cumulative amounts of income (Lorenz, 1905; Persky, 1992). Gini (1912) developed the Gini coefficient based on the Lorenz curve, defining inequality as $G = \frac{A}{A+B}$ where A is the area between the line of equality and the Lorenz curve and B is area below the Lorenz curve, giving an output between 0 (perfect equality) and 100 (perfect inequality) which remains one of the most widely used measures of inequality.

Despite this and much other criticism, including on the validity of his empirical work (Pigou, 1912, cited in Persky, 1992; Gini, 1936 cited in Persky, 1992), Pareto's assertion that the distribution of wealth is fundamental and unchanging gathered considerable

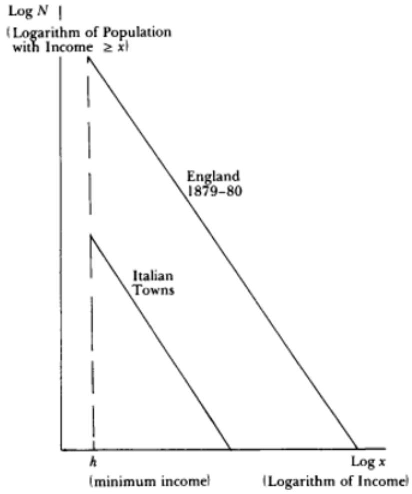

Figure 1.1: Income distributions plotted on a log-log scale give a straight line with the same gradient, regardless of where they are from (Pareto, 1897, cited in Persky, 1992).

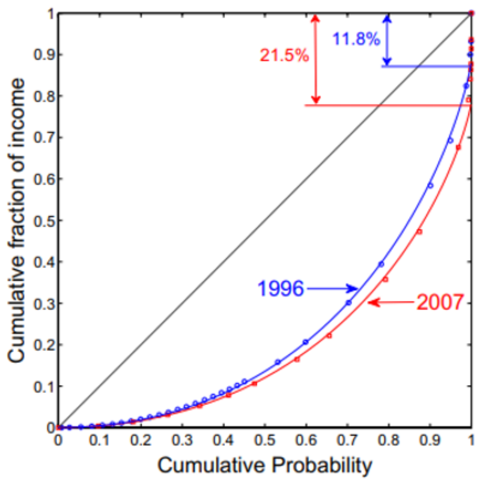

Figure 1.2: An example of Lorenz curves from Banerjee and Yakovenko (2010), showing income distributions for the United States in 1996 and 2007 against a line of perfect equality (black).

interest and economists and mathematicians set out to explain it. Numerous models were conceived to demonstrate Pareto's law using stochastic processes: considering each economic agent (people in the economy) as individual entities with probabilistic income and wealth (Champernowne, 1953; Gibrat, 1931; Kalecki, 1945). From this research it emerged that Pareto's law is observed in high incomes but "not even approximately obeyed" for low incomes and thus a "two-tailed distribution" is necessary (Champernowne, 1953). Mandelbrot (1960) proposed this as a "weak Pareto law", applying to only the top 20% of incomes and often only asymptotically rather than directly.

1.1.2 Turning to physics: Yakovenko

Around the year 2000, several econophysicists observed that the probability distribution of money resembles that of identical molecules in a gas. These molecules spontaneously form unequal energy distributions due to energy being transferred randomly in collision, tending towards an exponential distribution of energy between molecules at maximum entropy. They began to propose models based on economic agents similarly performing pairwise economic transactions, rather than having isolated individual fluctuations (Chakraborti and Chakrabarti, 2000; Dragulescu and Yakovenko, 2000; Krapivsky and Redner, 1998; Mimkes, 2000). They all follow the same broad principles, so this discussion will focus on the most famous of these, proposed by (Dragulescu and Yakovenko, 2000).

Based on the crude idea that the total amount of money in a closed economic system (a system from which there is no input or output of money) is conserved, Dragulescu and Yakovenko (2000) theorised that the distribution of money must reflect the Boltzmann-Gibbs law. This is a fundamental law of statistical mechanics that states that energy in a closed physical system, ε, will naturally form an exponential probability distribution $P(\varepsilon) \propto e^{-\varepsilon/T}$ where T is the temperature. They then demonstrated both theoretically and empirically – using computer simulations of random money transfers between economic agents which are each given the same initial amount of money – that most of the distribution of money, m, is similarly modelled by an exponential function:

$$P(m) \propto e^{-m/T_m}$$

With just one parameter, the "money temperature", T_m (Dragulescu and Yakovenko, 2001a). Just as the temperature of a gas measures the average energy of the molecules, the "temperature" of money corresponds to the average amount of money in the system.

Dragulescu and Yakovenko (2001b, 2003) subsequently showed that this function only fits the lower section of the income distribution, and a Pareto power-law function fits the upper end of incomes. Shown by plotting real economic data on a log-log scale (Fig. 1.3), the linear upper end demonstrates the power law, and on a log-linear scale (inset) the lower end is proved to be exponential by producing a straight line. Indeed, this combination of distributions has a counterpart in physics: under certain conditions, plasma produces a power-law distribution above a critical energy level and an exponential distribution below – known as 'thermal' and 'superthermal' parts (Hasegawa et al., 1985).

The explanation given for this two-part distribution evoked more conventional economics: the two distinct sections of the distribution reflect the "two-class structure" of the USA, with the intersection of the lines reflecting the boundary between the upper and lower class (Silva and Yakovenko, 2005). Silva and Yakovenko (2005) found that for the 1997 US income data, this boundary was at an annual income of around 120 k$, and 1-3% of

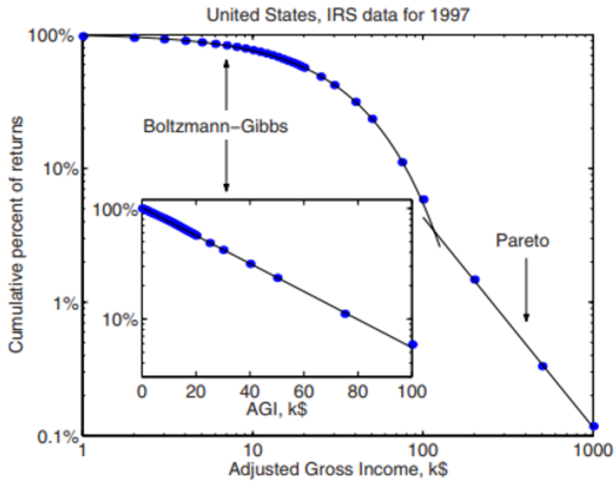

Figure 1.3: Cumulative probability distribution of annual income in the United States in 1997 from Internal Revenue Service data, with lines fitted to the Boltzmann-Gibbs exponential and Pareto power-law functions (Dragulescu and Yakovenko, 2003).

the population belonged to the upper class compared to 97-99% in the lower. By way of an explanation for the sharp boundary between the distribution patterns, Levy and Levy (2003) suggested that the upper tail is predominantly made up by those who own capital and generate the majority of their wealth from investment assets, while income in the lower tail is almost entirely from wages and salaries.

1.1.3 Applying statistical mechanics to energy consumption

While these methods proved useful in characterising and understanding wealth and income inequality within countries, Banerjee and Yakovenko (2010) identified that they are difficult to apply on a global scale due to the varying value of different currencies. They instead chose to measure energy consumption per capita – a physical quantity which is measured in the universal unit kW and therefore comparable across the world – as a proxy for living standards and therefore wealth. As the vast majority of global energy production is based on fossil fuel extraction, they claimed it can be considered a limited resource that is divided among the population and therefore, as with money, the distribution should tend towards maximum entropy rather than equality and thus show an exponential Boltzmann-Gibbs distribution:

$$P(\varepsilon) \propto e^{-\varepsilon/T_\varepsilon}$$

where ε is energy consumption per capita and the "temperature" T_ε is the average energy consumption per capita (Banerjee and Yakovenko, 2010; Lawrence et al., 2013).

Banerjee and Yakovenko (2010) studied empirical energy consumption data from the World Resources Institute (WRI) for 1990, 2000 and 2005, dividing total consumption for each country by its population and constructing a cumulative probability distribution. This work was replicated and improved by Lawrence et al. (2013) using Energy Information Administration (EIA) data, which included considerably more countries and data up to 2010. Both find qualitative evidence of a correlation between the theoretical exponential function above, shown by the dotted black line, and the empirical data (Fig. 1.4a), proving that energy consumption per capita approximately follows the Boltzmann-Gibbs distribution. They also show that the distribution of energy consumption has gradually become more equal between 1980 and 2010 and more exponential over that time (Fig. 1.4b), which Lawrence et al. (2013) attribute to economic globalisation facilitating greater interconnectivity and bringing the world closer to maximum 'entropy'.

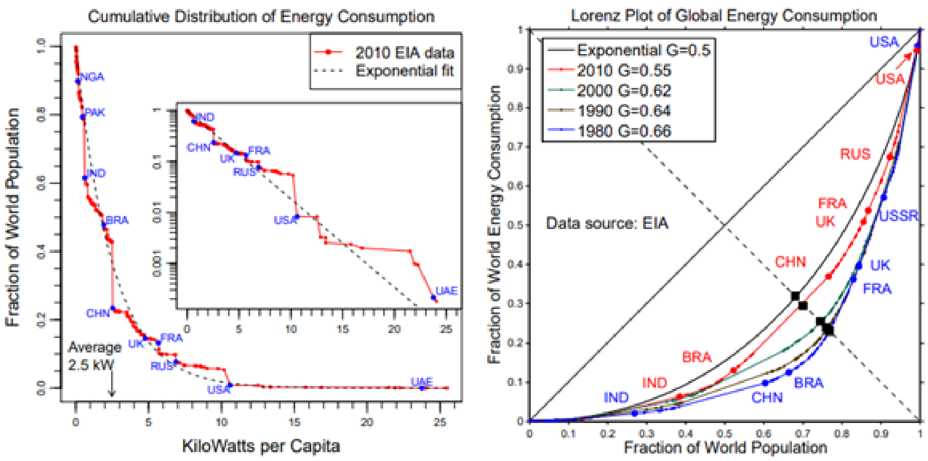

Figure 1.4: (a) Cumulative probability distribution of energy consumption per capita for a selection of countries in 2010 from EIA data, against a simple exponential function based on the average energy consumption per capita. The inset graph shows the same functions on a log-linear graph to prove their exponential distribution (Lawrence et al., 2013).
(b) Lorenz plots of energy consumption per capita for the same countries from 1980-2010 against a similar exponential function to Figure 5a, showing the trend towards more exponential (less unequal) distributions (Lawrence et al., 2013).

However, as Banerjee and Yakovenko (2010) particularly acknowledge, this analysis relies on the very crude assumption that the residents of each country each have identical energy consumption, as the only data widely available at the time of the research was total consumption per country. This leads to sharp discontinuities in the data (Figure 6), removing the possibility of a true exponential distribution, and allowing only a shallow level of analysis. As much of the work cited earlier shows, there is vast inequality *within* as well as between countries, and income distributions within countries show a more complex two-tailed structure – whether such a structure exists for the distribution of energy consumption remains to be seen.

With higher resolution data, it would be possible to determine with more accuracy whether an exponential Boltzmann-Gibbs distribution exists for energy or electricity

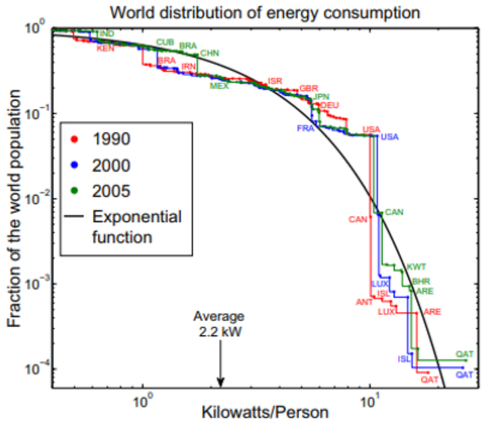

Figure 1.5: The functions from Figure 5 presented on a log-log scale, showing an exponential distribution analogous to the lower section of cumulative income distribution (Fig. 1.3) (Banerjee and Yakovenko, 2010).

consumption and test it both qualitatively and quantitatively (neither Banerjee and Yakovenko (2010) nor Lawrence et al. (2013) consider their data suitable enough to quantitatively test). There is also opportunity to examine whether there exists an upper tail of the population exhibiting a Pareto power-law distribution, which ought to follow if the distribution of energy or electricity consumption correlates with that of income.

1.2 Electricity consumption in the United Kingdom

As well as the potential for investigating whether such a two-class structure exists, disaggregated data for energy consumption within countries would allow analysis of inequality in electricity consumption on a sub-national level, something which is insufficiently studied in the literature (Jacobson et al., 2005). Such analysis would be able to investigate trends in the spatial distribution of electricity consumption within countries, reflecting regional inequalities, economic trends and the effects of public policy.

This is as important in the UK as anywhere, which has a growing gap between those on high and low incomes (Hood and Waters, 2017) and growing inequalities between regions (Ebell, 2017; Koop et al., 2018; Raikes et al., 2018) which are increasingly thought to be polarising politics (Zoega, 2016). Average household wealth is more than twice as high in the South East of England compared to the North East, for example, and wealth is growing at under half the rate in the North of England that it is in London (Raikes et al., 2018). Not only is this inequality a problem of economic and moral justice, and not only is it harmful to the economy as a whole (Wilkinson and Pickett, 2009), but it is increasingly recognised that successful strategies to move towards a sustainable economy must take into account the heterogeneity of the economy and the effect this has on policy (Lawrence et al., 2013; Wu et al., 2012).

For example, in researching the link between income and energy consumption, Cayla et al. (2011) sugested that a policy such as a universal carbon tax could be ineffective in reducing emissions, as poorer households do not have the necessary access to capital to make improvements that would reduce their energy bills. Even if given the capital, they argued, there would be potential for a rebound effect as they use the improvements to increase their consumption in order to catch up with middle class living standards. Indeed, the consumption of electricity is a cause as well as a product of economic inequality: over 11% of households in England are in fuel poverty (defined as households requiring fuel costs above the national median, or where fuel costs leave the household with a residual income below the official poverty line) (BEIS, 2018a). This too is, unsurprisingly, unevenly distributed regionally, with 13.8% of households in the North East in fuel poverty and 9.0% in the South East (BEIS, 2018a).

1.2.1 The urban-rural divide in electricity consumption

A consideration that cannot be ignored in spatial analysis of the UK and the consumption of electricity is the urban form and the rural-urban divide. The UK is one of only ten countries which has over 5% of its total land area covered by cities (Angel et al., 2011), and it has over 80% of its population and growing living in urban areas (ONS, 2014). With urban settlements responsible for 76% of global energy consumption (IEA, 2013) and 71% of energy-related carbon dioxide emissions (WEC, 2013), understanding the size, distribution and consumption patterns of urban and rural areas and their relation to electricity is highly relevant.

Although much of the literature is focused outside the UK, the effect of the urban form on electricity consumption in general is an area that is studied with keen interest. Norman et al. (2006) studied energy use and greenhouse gas emissions in the building operation, construction and transportation sectors across residential areas of low and high density in Toronto, and discovered substantially lower per capita energy use and emissions in high density urban core development. O'Brien et al. (2010) confirmed this, finding reduced net energy use with increasing density of housing. As electricity makes up a substantial part of energy consumption, particularly on a domestic level (BEIS, 2018b), trends in energy consumption can broadly be assumed to apply to electricity and therefore urban areas in the UK might be expected to display lower electricity consumption per capita than rural areas. Some research has suggested, however, that although direct energy use is generally lower for households in urban environments, they have larger total energy use profiles associated with their wider consumption of goods and services (Wiedenhofer et al., 2013) – a factor frequently neglected in this area of research due to the complexity of tracing the inputs to everyday household consumption.

Hui (2001) quotes a number of reasons that may suggest why urban environments demonstrate lower household energy consumption: (a) the building stock, as it is more compact and in higher densities, consumes energy more efficiently; (b) the improved (and in some cases reduced need for) transport and communication reduces energy use; (c) new and more efficient energy technologies are more easily implemented in cities; and (d) greater potential for mixed land use leads to improved efficiency. Significant energy savings through reduced heat demand in high density housing are consistently shown in studies of urban morphology and energy use (Rode et al., 2014), including studies of thermodynamics that show denser housing typically has lower surface area-to-volume ratios and therefore reduced heat losses (Mohajeri et al., 2015).

However, not all studies give a direct negative correlation between density of urban environment and energy use. Rather, Minx et al. (2013) studied the carbon dioxide emissions of urban and rural areas of the UK at a high resolution and found that "carbon footprints can be comparatively high or low across density gradients depending on the location-specific socio-demographic, infrastructural and geographic characteristics of the area under consideration". Baiocchi et al. (2015), moreover, reject any attempts to explain energy use and emissions using simple correlation analysis and suggest instead that consumption is driven by "unique, place-specific combinations of emission drivers" including density, income, household size, access to central heating and the number of houses in poor condition. Their study of the UK on a finer spatial resolution did, however, find that density and income were dominant factors in energy use. Both studies note the importance of granularity on the accuracy of this research and impacts on policy: not only does finer-scaled data allow for better analysis of how energy consumption is distributed and therefore better generalised policy, but also better understanding of heterogeneity within cities and regions and therefore specialisation of policy (Baiocchi et al., 2015).

The exact relationship between the urban form and electricity consumption, and its importance in the distribution of electricity consumption in the UK, is therefore a topic that needs more research with high resolution data. These could maximise the use of the methods discussed earlier: if urban electricity use per capita is systematically different to that in rural areas, a cumulative probability curve for urban areas would be shifted horizontally, and if all other factors are equal the curve would remain the same shape.

1.3 Aims and hypotheses

The distribution of electricity consumption in the UK is a very broad topic, and the data that has been made available (Jarvis, 2018) has yet to be applied to it, allowing a broad scope for potential areas of study. As resources are limited, however, this research will focus on two main aims. First, to study the extent to which electricity consumption in the UK follows the Boltzmann-Gibbs and Pareto laws, as the literature would suggest. If this distribution is truly universal and entropy is always maximised, it should follow this distribution at any scale, anywhere in space - few studies have been able to compare the extent to which it applies at very small scales. Second, to explore how the distribution of electricity consumption varies across of the UK, with a particular focus on comparing urban and rural areas. To support these aims, three null hypotheses will be tested:

(1) Electricity consumption per capita in the UK does not demonstrate a two-class distribution that fits the Boltzmann-Gibbs and Pareto laws.

(2) Electricity consumption per capita in the UK does not demonstrate a two-class distribution fitting the Boltzmann-Gibbs and Pareto laws at all spatial scales; if it applies at the national level, it ceases to apply at other scales.

(3) Electricity consumption per capita in the UK demonstrates the same distribution in urban and rural areas.

Section 2

Method

2.1 Sourcing and preparation of datasets

The Global Gridded Electricity (GGE) dataset contains raster files in GeoTIFF format mapping annual electricity consumption in kilowatt-hours per square kilometre (kWh/km^2) with a 30 square arcsecond (approximately one square kilometre) spatial resolution (Jarvis, 2018). These are provided for 2000, 2005, 2010 and 2015, data which could be used for temporal studies, but for the purposes of this research only the most recent (2015) was used.

The 2015 GGE raster was imported into ArcGIS, which provides a wide range of tools for manipulation of spatial data. Administrative boundaries for the UK in shapefile format were sourced from the Database of Global Administrative Areas (Hijmans et al., 2018), and, using the *Clip* tool in ArcToolbox, the GGE raster was clipped to the boundary for the UK to create a UK gridded electricity map.

A global gridded population raster for 2015 was taken from the Gridded Population of the World version 4 dataset (GPWv4), which similarly provides yearly global population data for 2000, 2005, 2010 and 2015 in 30 square arcsecond cells (CIESIN, 2016). The *Raster Calculator* tool in ArcGIS's Spatial Analyst extension was then used to divide the electricity consumption by the population count for each cell to produce a map of UK gridded electricity per capita.

This map was exported as a raster and imported into R using the *raster()* function, and to extract the data in a format suitable for analysis was then converted into a data frame using the *as.data.frame()* function and from that into a vector using *unlist()*.

2.2 Hypothesis one: fitting a two-class distribution

The electricity consumption per capita map was investigated to assess the extent to which its data demonstrated Dragulescu and Yakovenko (2001b)'s observed two-class distribution. Per capita data removes the effect of uneven population distribution and allows areas to be compared like for like, and it is a reasonable predictor of incomes (de Rezende Francisco et al., 2007), hence it is the chosen metric in the literature. Cumulative probability graphs were produced by iterating through the values and calculating the probability of a square kilometre having the particular electricity consumption n, $P(\varepsilon)$, as $1 - \frac{n}{N}$ where N is the total number of values. This was plotted on a log y axis against electricity consumption on a log x axis, and repeated for electricity consumption per capita. If the null hypothesis is true, a two-class distribution should be evident on these graphs, with an exponential shape on the left and a linear upper tail on the right. If such a distribution was qualitatively visible, the inflection point between the exponential and linear curves was located by eye and the data either side of it separated for quantitative testing.

To test the extent to which the lower section of the plot is exponential, a further plot was made on which the natural log is taken of just the electricity consumption (x axis) data, which would be expected to produce a perfectly linear curve if the log-log graph is truly exponential. The *lm()* function was used to perform linear regression, fitting a straight line to the data and testing the variance from the line, giving an adjusted R^2 value as a simple test of the fit and as such a test of whether the distribution matches the expected exponential shape. A similar linear regression was then performed on the data above the inflection point with the natural log performed on both axes, to test the extent to which it is linear, reflecting a power-law relationship.

Splitting the distribution into sections by eye from a graph is an imprecise method and therefore the R^2 values will have a substantial amount of uncertainty, but as this hypothesis is only concerned with proving the presence of the expected distribution shape at all and the R^2 values will not be used for any further quantitative analysis, it is acceptable. They will be presented with two significant figures, reduced from the four produced by the *lm()* function, to convey the low amount of precision in their production.

2.3 Hypothesis two: fitting at all spatial scales

The second hypothesis required this relationship to be tested across space and across scale to see if there were areas of the country where the distribution does not apply, or spatial scales at which it is no longer true. In the interests of the time and processing power available, it was necessary to test a relatively small sample of locations and scales. To ensure that the sample is representative of the country, and with the third hypothesis in mind, a mix of rural and urban areas were selected.

Five cities in the UK were chosen: Manchester, Newcastle, Glasgow, Birmingham and Bristol. These were selected because they were among the ten largest built up areas in the dataset by area (ONS, 2017a) and were geographically diverse, representing at least the North, Midlands and South of England, and Scotland. London was also included, but considered separately from other cities in case 'megacities' obey a different relationship. Shapefiles for these cities were provided by ONS (2017a), and centre points for them were calculated using ArcMaps's *Calculate Geometry* tool. Five rural local authorities were also chosen. They were selected to be a geographically diverse sample of the local authorities with the lowest population density (ONS, 2013): Argyll & Bute and Highland in Scotland, Eden and Ryedale in England, and Powys in Wales. Shapefiles for local authorities were included in the data sourced from Hijmans et al. (2018), and their centre points were calculated using ArcMaps's *Calculate Geometry* tool.

As a macro-scale fit will have been investigated for the first hypothesis, three spatial scales from regional level and below were studied for each location. Circles of 75 km, 25 km and 13 km diameter were extracted from the data: 75 km was a size that includes the entirety of each city (excluding London), 25 km was likely to include just urban area, and 13 km is an extreme close-up of the data to see whether the distribution is observed even at very small scales. The circles were drawn from the centre point of the area using ArcMap's *Editor* toolbar and extracted using the *Clip* tool in ArcToolbox (Fig. 2.1).

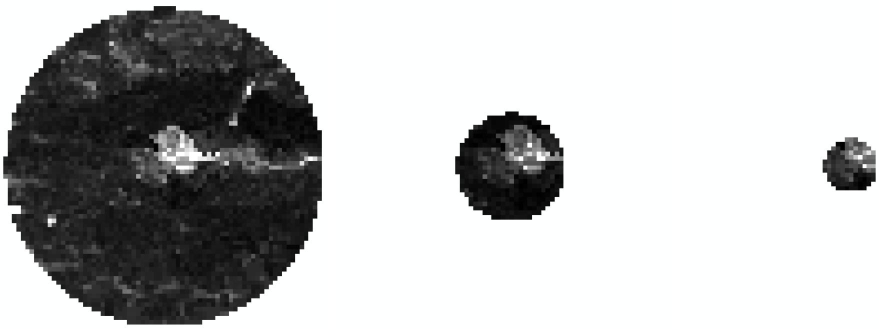

(a) 75km diameter (b) 25km diameter (c) 13km diameter

Figure 2.1: An example of one location examined: a raster image of 75 km, 25 km, and 13 km diameter circles of electricity consumption per capita in square kilometre cells from the centre of London. The River Thames and City of London are visible as areas of high per capita consumption.

For each scale at each location, the raster was imported into R and converted into vectors as in section 2.1, and the cumulative probability calculated as in section 2.2. If there was a two-class distribution evident in a log-log graph, the inflection point was located by eye, and a linear function was fitted to a log-linear plot of the lower section and another fitted to a log-log plot of the upper section as in section 2.2. A large R^2 value for both would indicate that each section of the distribution is exponential and power-law respectively, and disprove the null hypothesis.

Again, as the second hypothesis is only concerned with in a binary test of whether the expected distribution is found at any place and scale or not, splitting the distribution into sections by eye from a graph is an acceptable method. It follows, however, that the results from this are imprecise and care should be taken in comparing them. The values will again be reduced two significant figures for presentation.

2.4 Hypothesis three: comparing urban and rural

To explore the effect that different types of environment have on the distribution of per capita electricity consumption, a large sample of both urban and rural data were aggregated into single datasets and these were compared to the total UK distribution. A 'cities' dataset was created from a combined raster of the top 10 cities by population after London, produced using the *Clip* tool in ArcToolbox on selected built up area shapefiles from ONS (2017a). A 'rural' dataset was produced from a similar combined raster of all wards given the most rural classification, E2, in the UK Government's Rural Urban Classification (ONS, 2017b).

Density plots and box plots were then produced using the *ggplot* package and *boxplot()* fuction in R respectively for each of the cities, urban, UK and London datasets; the latter again included separate from other UK cities to investigate the effect of its substantially greater size. These plots demonstrate the general qualities of the distribution as compared to the overall UK distribution, such as their comparative averages and range. A log-log cumulative probability graph was then produced with the cities, rural and London data on the same plot to investigate whether the distribution shape, if it fits the two-class distribution predicted by the literature, changes based on the environment. If the null hypothesis is true, the graph for each should overlap perfectly or at least show no systematic difference.

Additionally, to assess the inequality within each type of environment, a Lorenz curve was produced along with Gini coefficients using the *Lc()* function and the *ineq* package respectively. If the null hypothesis is true each environment should demonstrate the same inequality.

Section 3

Results

3.1 Hypothesis one: fitting a two-class distribution

The distribution of electricity consumption per capita for each of the United Kingdom's 485,271 square kilometre cells is heavily skewed towards low values, with a peak at around 12 kWh/p/km^2 (Fig. 3.1).

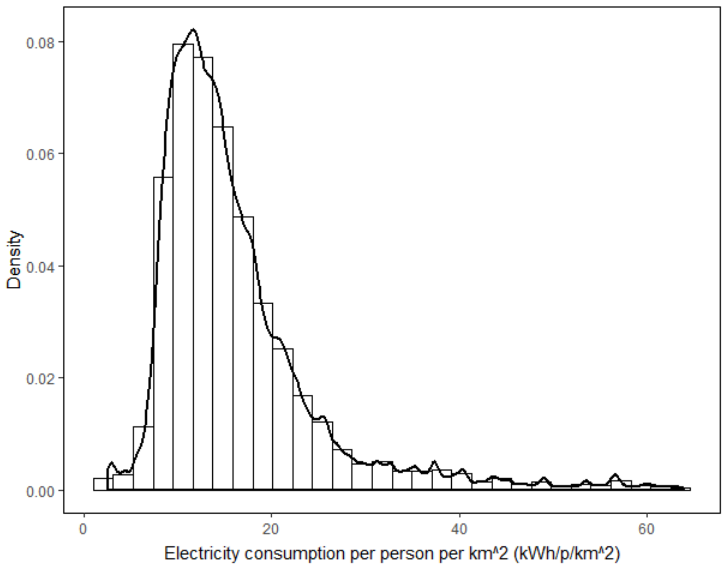

Figure 3.1: Density plot of electricity consumption per capita per square kilometre in the United Kingdom, demonstrating a strong positive skew.

On a log-log cumulative probability plot, similar to that used by Dragulescu and Yakovenko (2001b), the data appear to show both an exponential and linear section (Fig. 3.2). Splitting the data for each section into separate graphs makes this visually clearer (Fig. 3.3).

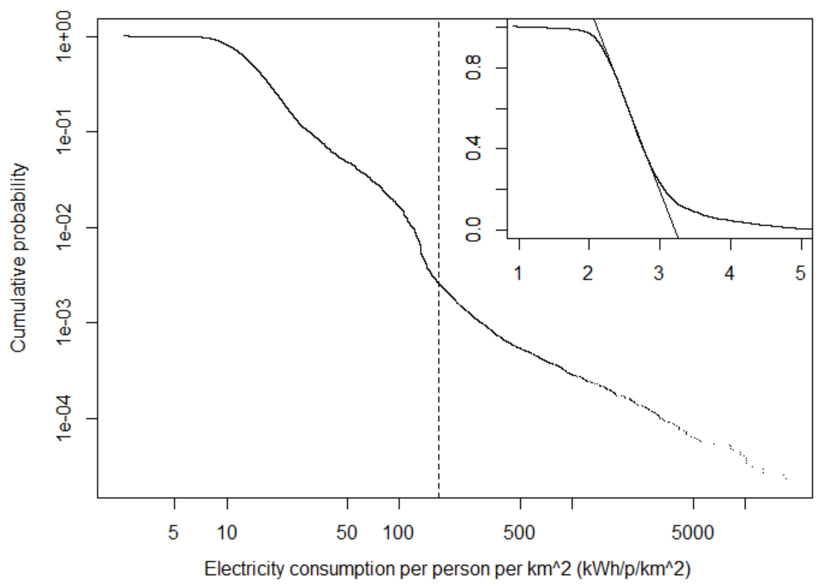

Figure 3.2: Cumulative probability graph showing electricity consumption per capita per square kilometre in the United Kingdom with log-log axes. The dashed line indicates the suggested divide between a lower exponential (Boltzmann-Gibbs) section and upper liner (Pareto) section.
Inset: A log-linear plot of the same data (the x axis showing the natural log of electricity consumption per capita per square kilometre), with a linear function fitted to the data.

The lower section is confirmed to be exponential by fitting a linear function to the natural log of the x axis data (Fig. 3.2, inset), giving an R^2 value of 0.99 (where $p<.001$, indicating that the result is statistically significant), and a similar fit to the natural log of both axes in the upper section gives an R^2 value of 0.99 ($p<.001$), confirming it to be linear.

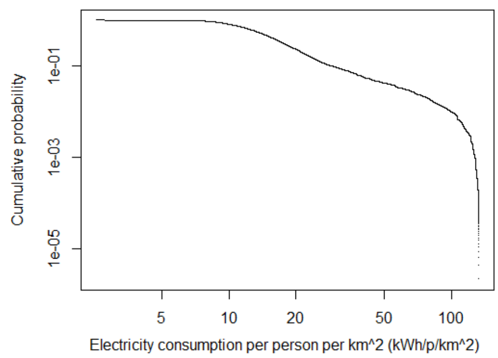

(a) Lower exponential section

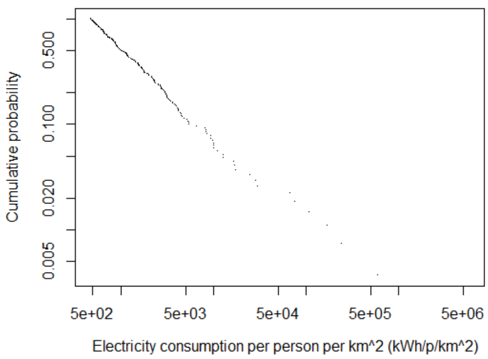

(b) Upper linear section

Figure 3.3: The electricity consumption per capita cumulative probability data from Fig. 3.2 on log-log axes, showing just the data below and above the dashed line respectively. This shows the distinct shape of the two distributions.

The boundary between the lower and upper section is at around 139 kWh/p/km^2; 99% of the data is in the lower section.

3.2 Hypothesis two: fitting at all spatial scales

Linear functions fitted to each part of the distribution (log-linear to test for an exponential relationship in the lower section and log-log to test for a power-law relationship in the upper) give good measures of fit for almost all locations and spatial scales tested (Table 3.1). An example provided below shows the reduction in quantity of data at smaller scales but the persistence of the two-class distribution (Fig. 3.4).

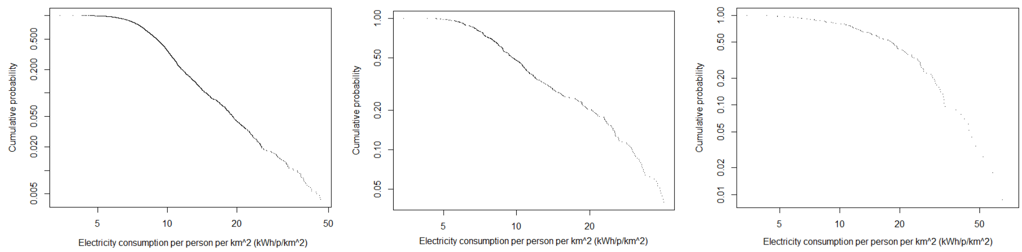

(a) 75km diameter (b) 25km diameter (c) 13km diameter

Figure 3.4: An example of one location tested and the distribution at each spatial scale. Electricity consumption per capita in square kilometre cells within a 75 km, 25 km, and 13 km diameter circle around London.

Table 3.1: Adjusted R squared values for straight lines fitted to log-linear plots of the lower (Boltzmann-Gibbs) section, and log-log plots of the upper (Pareto) section, of electricity consumption per capita data in square kilometre cells within a 75 km, 25 km, and 13 km diameter circle around London, five other urban areas, and five rural areas. With a small number of exceptions, a good correlation is demonstrated, suggesting that the expected two-class distribution applies universally. For all values $p<.001$.

Area	Diameter	Lower R^2	Upper R^2
London	75km	0.99	0.99
	25km	0.99	0.90
	13km	0.99	0.98
Manchester	75km	0.99	0.80
	25km	0.99	0.98
	13km	0.99	0.97
Tyne & Wear	75km	0.99	0.79
	25km	0.99	0.92
	13km	0.99	0.89
Bristol	75km	0.99	0.80
	25km	0.99	0.97
	13km	0.98	0.90
Glasgow	75km	0.99	0.93
	25km	0.99	0.96
	13km	0.99	0.95
West Mids	75km	0.99	0.97
	25km	0.99	0.98
	13km	0.99	0.82
Argyll & Bute	75km	0.98	0.86
	25km	0.97	0.95
	13km	0.98	0.94
Highland	75km	0.95	0.95
	25km	0.91	0.93
	13km	0.79	0.48
Eden	75km	0.99	0.90
	25km	0.95	0.92
	13km	0.83	0.89
Ryedale	75km	0.99	0.98
	25km	0.99	0.87
	13km	0.96	0.94
Powys	75km	0.99	0.96
	25km	0.98	0.97
	13km	0.97	0.43

3.3 Hypothesis three: comparing urban and rural

Aggregated data for the most urban and most rural environments show a clear trend: electricity consumption per capita is higher than the UK average in rural areas and lower than the UK average in urban areas (Figs. 3.5, 3.6, 3.7). Urban areas also show a substantially more concentrated distribution with a smaller range of values than either rural distributions or that of the whole country (Figs. 3.5, 3.6). London is broadly similar to other cities in both of these regards (Figs. 3.5, 3.6).

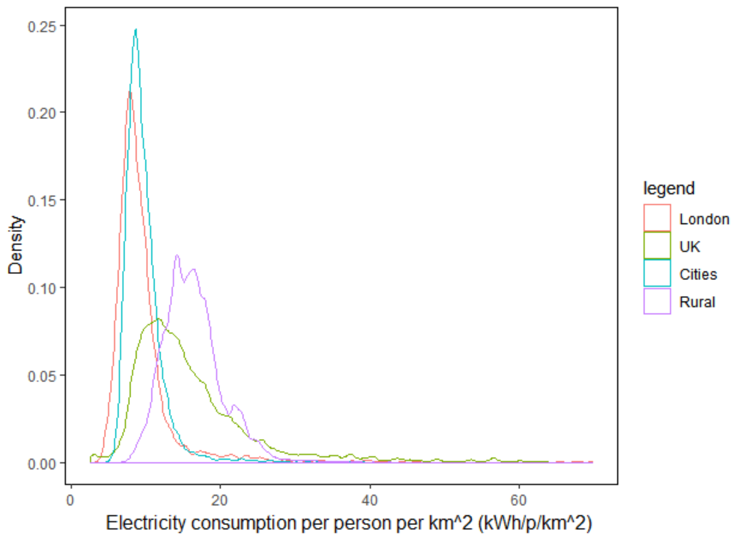

Figure 3.5: Density plots for electricity consumption per capita data in square kilometre cells for the ten largest UK cities after London combined, all ward areas classified as E2 (the most rural) in the Government rural-urban classification combined, London, and the United Kingdom.

These trends are confirmed by log-log cumulative probability plots: rural areas displace the distribution to the right, representing uniformly increased electricity consumption per capita across the whole distribution (Fig. 3.7). London's distribution is distinctly different from other cities however, with slightly lower values in the lower section but an upper tail that makes up a much larger proportion of the distribution. To investigate this further, the proportion of the population of each study area in the lower section was recorded, and this demonstrated that London's upper tail makes up a proportion 10% higher than in either typical urban or rural environments (Fig. 3.2).

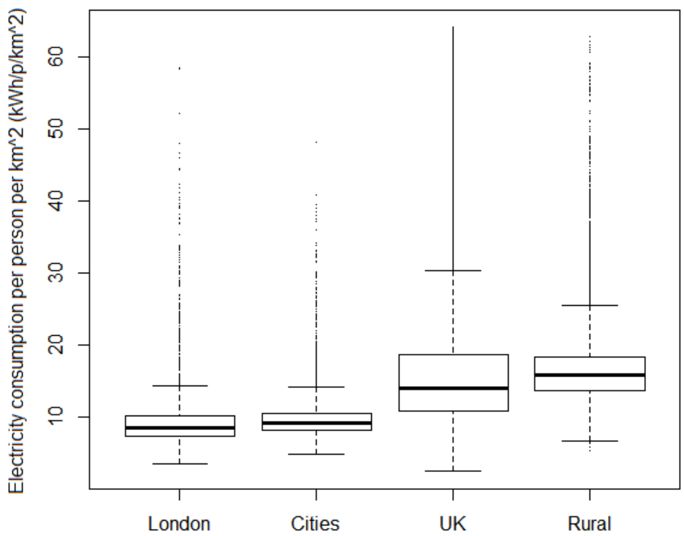

Figure 3.6: Box plots for electricity consumption per capita data in square kilometre cells for the ten largest UK cities after London combined, all ward areas classified as E2 (the most rural) in the Government rural-urban classification combined, London, and the United Kingdom.

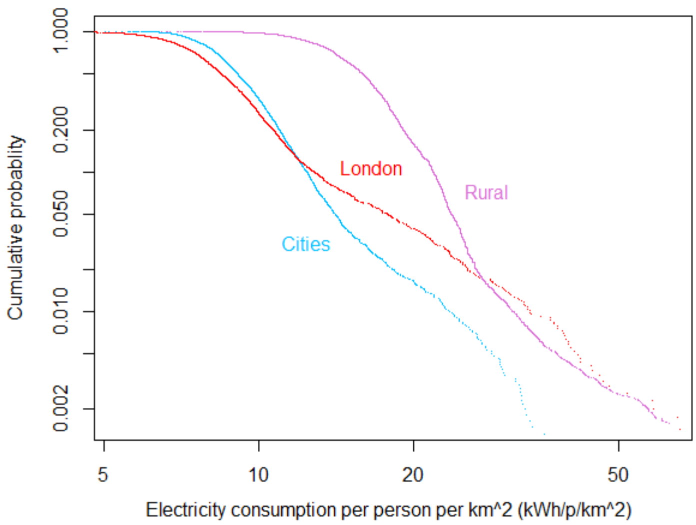

Figure 3.7: Cumulative probability graph showing electricity consumption per capita in square kilometre cells with log-log axes for London, the ten largest UK cities after London combined, and all ward areas classified as E2 (the most rural) in the Government rural-urban classification combined, showing how the distribution varies between different settings.

Finally, Lorenz curves and Gini coefficients give simple measures, qualitatively and quantitatively respectively, of how pure inequality varies between each of the environments. Both rural and urban settings have a considerably less unequal distribution than the United Kingdom as a whole, but London is more unequal than both (Figs. 3.8).

Table 3.2: The proportion of the distribution of electricity consumption per capita displaying an exponential Boltzmann-Gibbs distribution rather than a power-law Pareto distribution, and therefore belonging to the lower of two societal classes, for the UK, London, the ten largest UK cities after London combined, and all ward areas classified as E2 (the most rural) in the Government rural-urban classification combined.

Study area	Proportion of data points in lower class
UK	99%
London	85%
Rural	98%
Cities	96%

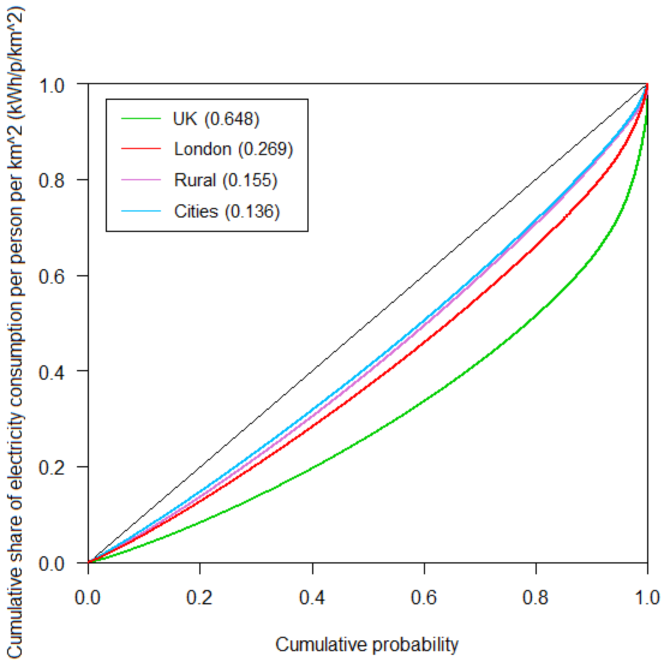

Figure 3.8: Lorenz curve demonstrating inequality in the electricity consumption per capita in square kilometre cells for the UK, London, the ten largest UK cities after London combined, and all ward areas classified as E2 (the most rural) in the Government rural-urban classification combined. The straight black line represents perfect equality, with increasing distance from it representing greater inequality in the study area. Gini coefficients are provided in brackets, with higher values up to 1 representing greater inequality.

Section 4

Discussion

4.1 Hypothesis one: fitting a two-class distribution

The basic distribution shape of electricity consumption per capita (Fig. 3.1) is as expected, with most of the data points concentrated at the lower end of the distribution and the density of values exponentially tailing off towards the higher end of electricity consumption. This substantial inequality in electricity consumption fits with the literature on energy consumption (and indeed also on wealth and incomes) which states, for example, that the top third of the world population consumes two thirds of the world's energy (Lawrence et al., 2013). Most notably, as discussed in section 3.1, on a log-log cumulative frequency plot the distribution fits the two-class structure that Dragulescu and Yakovenko (2001b) found to fit money in the economy, with what appears to be an exponential Boltzmann-Gibbs lower section and a linear Pareto upper tail (Fig. 3.3). These are confirmed quantitatively by R^2 tests which give very strong measures of fit to the two expected shapes and are both statistically significant. Although the lower section does not appear to fit an ideal exponential shape on log-log axes, on a plot with the natural log performed on just one axis (Fig. 3.2, inset) an almost perfectly straight line is formed, giving reasonable evidence that it is broadly exponential. The deviations this straight line at the top and bottom are likely to be the effect of the distribution below the peak in Fig. 3.1, which is not reflected in the log-log cumulative plot, and the power-law upper tail respectively.

Despite the imprecise methods used (discussed in section 4.4), the evidence, particularly the quantitative tests, is strong enough that the null hypothesis can be considered to be disproved: a two-class exponential and power-law distribution is evident. This is, to the best of our knowledge, the first time that such a distribution has been proven for electricity consumption, at least in the United Kingdom. This lends validity to the argument that the distribution of electricity consumption in a closed system, as with income and wealth, is predisposed to this distribution by physical processes, and naturally forms an unequal distribution as a result of the second law of thermodynamics (Dragulescu and Yakovenko, 2001b; Lawrence et al., 2013).

This has substantial and far-reaching implications: for example, that the distribution is near-exponential suggests that it is close to maximum entropy. Lawrence et al. (2013) propose that as the distribution of energy consumption reaches maximum entropy (as they suggest is now beginning to happen on a global level), the general trend they identify over recent decades towards reductions in inequality will slow down and inequality will stabilise at the current level; in other words, the distribution shown represents the "natural" inequality. This inequality, it is claimed, is "virtually unavoidable" and difficult to fight against (Lawrence et al., 2013). If the current inequality in electricity consumption (as well as wealth and income, according to Dragulescu and Yakovenko (2001b)) across the UK is inevitable and fixed at roughly its current state, this is hugely significant to the efforts of much economic policy. One may conclude, as Pareto did over a century ago, that it is an inevitable reflection of human nature that a small number of "parasites" will always hold much of society's wealth (Rodd, 1995), and that to increase the wealth of the lower classes it is necessary to raise the wealth of the whole country (Persky, 1992).

Lawrence et al. (2013) point out, however, that their analysis of energy consumption - the assumptions from which are replicated in this study - relies on the idea that energy consumption is a 'closed system' and therefore, just as in the Boltzmann-Gibbs distribution of energy in gas particles, the total quantity of it in the system is finite. As discussed in section 1.1.3 this is only assumed to be true of energy consumption because the majority of it is based on naturally-finite fossil fuels; with the advent of widespread renewable energy there is potential for this to no longer be true. A wholesale reconfiguration of the economy towards one based around renewables could potentially therefore not only substantially reduce greenhouse gas emissions, but also reduce inequality in a way that is simply impossible in the current economy. Lawrence et al. (2013) also note that decentralised energy, produced and consumed locally, may not be subject to the principles of maximising entropy that cause this 'inevitable' distribution as they are not able to be distributed and exchanged between most of the population.

The boundary between the classes, at around 99% of the population in the lower class, is high in comparison to the 95% which Dragulescu and Yakovenko (2001b) found for income in the UK. Given the uncertainty involved in the methodology (see section 4.4) and the lack of any other similarly specific research into the statistical mechanics of inequality in the UK, it is perhaps unwise to draw too much analysis from the apparent 4% difference. It can at least be concluded that electricity consumption is broadly indicative, if not perfectly representative, of income in the UK. This may suggest another conclusion about the impact of these findings: the problem of inequality in electricity consumption is an inevitability under the current economic system because it is strongly linked to the distribution of money, which is inevitably unequal. Rather than simply decoupling this system from fossil fuels (with the unrealistic necessity of ever-increasing, perhaps infinite, electricity production in order to allow redistribution) reducing this inequality may necessitate severing the connection between the money system and electricity altogether.

4.2 Hypothesis two: fitting at all spatial scales

The second null hypothesis stated that this specific two-class distribution shape would not apply at all scales, and suggested that as the scale of the study area got smaller there is likely to be a point at which it ceases to apply. Replicating the analysis used to test the first hypothesis for circles of increasingly small radii around a geographically diverse set of locations, however, gave strong (and statistically significant) R^2 measurements for fit to this distribution even at the smallest scales tested (below which there is not enough data to meaningfully analyse) (Table 3.1). This is strong evidence for disproving the null hypothesis.

Only two R^2 values can be considered too low to prove a reasonable fit (all other values are 0.79 or above): the upper power-law sections in the smallest scale (13 km diameter) samples taken in the Highland and Powys local authorities give 0.48 and 0.43 respectively (Table 3.1). This appears to be a problem exclusive to the upper sections - the comparative lower sections demonstrate good measures of fit - and the smallest scale - at larger diameters both sections give good measures of fit - and it is only found in rural areas. A likely explanation is that the upper Pareto sections are simply much smaller in rural areas (see Table 3.2) and these are among the most rural areas in the UK; at a 13-cell diameter it is likely that they are made up of a very small number of data points, giving a high potential for random variation to distort the overall shape of the distribution. The strong exponential shape for the lower section and strong power-law shape at other scales suggest that entropy is being maximised in the expected pattern, but the upper tail is simply too small to be reliably measured. With this in mind, and considering the overwhelming evidence from the rest of the measurements, this null hypothesis can be considered disproved: not only does a two-class exponential and power-law distribution apply to electricity consumption in the UK, but it applies across the country at all spatial scales.

4.3 Hypothesis three: comparing urban and rural

The third null hypothesis proposed that there is no systematic difference between the distribution of electricity consumption in urban and rural areas, in response to contested literature about the effect of the urban form on electricity consumption and inequality. The results, however, demonstrate substantial difference between urban and rural electricity consumption, most obviously that it is consistently higher than the UK average in rural areas and lower in cities (Figs. 3.5, 3.6). Both demonstrate a two-class exponential and power law distribution as expected from disproving the second hypothesis, but the rural distribution is uniformly displaced to the right (representing higher electricity consumption) and urban to the left (Fig. 3.7), suggesting systematically increased electricity consumption in rural areas and reduced in cities. This firmly refutes the null hypothesis and gives credence to research that suggests a correlation between building density, or "urbanness", and electricity consumption (Hui, 2001; Norman et al., 2006; O'Brien et al., 2010).

Cities also demonstrate a more homogeneous distribution (Figs. 3.5, 3.6), with an interquartile range, for example, nearly half that of rural areas (Fig. 3.6). This might suggest that the effects which produce lower average electricity consumption in cities, such as reduced heat demand in higher density housing (Mohajeri et al., 2015; Rode et al., 2014) or more efficient transport (Hui, 2001), also have the effect of unifying the amount of electricity consumed across the city. Rural areas, on the other hand, are often very heterogeneous and the demand for heating, transport and other uses of electricity is more varied, perhaps explaining the wider spread of values. However, this homogenisation is not reflected in measures of inequality within each environment: cities are slightly more equal than rural areas, with the exception of London, but the difference is relatively minor with less than 0.02 difference in Gini coefficient (Fig. 3.8). Therefore while the distribution of electricity consumption in cities mostly covers a smaller range of values than rural areas, the data within that range follow a similar unequal distribution, which as demonstrated is mostly exponential with a proportionally small power-law upper tail (Fig. 3.7).

Although the distribution is a similar shape for both the UK's rural areas and most of its cities, London is an exception. It has a strong two-class distribution (Table 3.1, Fig. 3.7) and the lower section is similar to that of the other cities tested, but its upper tail makes up a considerably larger proportion of the data than either the other cities or rural areas (Table 3.2, Fig. 3.7). Unlike the next ten largest UK cities, which have 4% of their combined population in the upper class (itself high compared to the UK average of 1%), London has a staggering 15%. London therefore has, according to Levy and Levy (2003), a disproportionately large part of society that generates its wealth from owning capital rather than earning wages (assuming that the two-class distribution of electricity consumption is produced by the same factors as that for income and wealth). Nonetheless, the inequality within London is also greater than that of cities or rural areas (Fig. 3.8), suggesting a heavily divided city. The dominance of capital and the considerable inequality in London requires an explanation greater than this study can afford, but is a growing affliction of globally-connected cities in the twenty-first century (Florida, 2017).

Figure 3.8 ostensibly shows, however, that inequality present within the UK as a whole is a lot higher than the sample taken of its cities, its rural areas and even London. This is likely to be a result of the vastly greater quantity of data for the whole UK - Lorenz curves and Gini coefficients are disproportionately affected by outliers (Cowell and Victoria-Feser, 1996; Prendergast and Staudte, 2015) and the UK data has a very large number of these (Fig. 3.6). OECD (2018) found an income Gini coefficient for the UK of

0.351 in 2016, and have not found an income Gini coefficient for any country over the last five years which is higher than 0.459, so a value of 0.648 for electricity consumption in the UK is unlikely. It seems reasonable to conclude that this value is probably unreliable. Those for London, other cities and rural areas do not appear to be unduly affected in this way however, indeed London has a smaller number of data points (3,473) than the other cities and rural areas (6,056 and 49,578 respectively, by comparison the UK has 485,271) but a higher Gini cofficient. A reasonable amount of uncertainty should nonetheless be assumed in the values for these, but the broad comparisons made in this section appear acceptable.

4.4 Assumptions, limitations and uncertainties

Naturally for a study on this scale, a number of assumptions have been made that will affect the interpretations made from the research, and the methodology involved produced a number of sources of uncertainty in the findings. Firstly, electricity consumption per capita was averaged over each square kilometre as that was the highest resolution of the available data: in any given square kilometre, every person was said to have the same consumption. This is a major improvement from past attempts, such as those of Lawrence et al. (2013) which averaged per capita data over whole countries, and is to the best of our knowledge the highest resolution on which these methods have been applied to energy or electricity consumption. Nonetheless the inequality present within these square kilometre areas and any extremes within them may be neglected, and this is likely to disproportionately affect urban areas where there are high population densities.

Furthermore, the electricity consumption data provided for each square kilometre is total consumption, including all forms of residential, industrial and commercial electricity demand, rather than just household consumption (Jarvis, 2018). This is appropriate as the literature theorised that *total* energy or electricity in the system will follow this distribution if it is finite, but may produce unexpected results when compared to household income data, for example, which is used extensively in the literature as a proxy for the total distribution of money (Dragulescu and Yakovenko, 2001a). Industrial areas with large amounts of electricity consumption but low permanent population will produce unexpectedly high electricity consumption per capita figures, which may appear to distort the distribution in some local areas and make geographic comparisons to household data difficult. Care was taken not to compare the electricity consumption per capita and household income data except at the level of national inequality.

The method used to test for the expected distribution also produces uncertainties. R^2 tests are a very simple and useful measure of fit to linear functions, but can be misleading: they can give inappropriately high values for a small number of data points (resolved by using adjusted R^2), and they can demonstrate a good level of fit despite the data potentially being fit better by another function. For example, the log-linear plots created to prove an exponential distribution by showing a straight line gave a non-linear result that would qualitatively appear to fit a cubic distribution better than a linear one (Fig. 3.2, inset), but nonetheless gave a high adjusted R^2 value to a linear function as there are a roughly equal number of data points on either side. In this instance, however, as discussed in section 4.1 the middle section was treated as the expected straight line, and the deviations at either end explained as unrelated parts of the distribution and removed from the test. All the data tested against linear functions appeared linear in a graphical plot, such as in Fig. 3.3b.

As discussed in section 2.2, locating the transition between the exponential and power-law parts of the distribution by estimating based on appearance on a graph is an imprecise method, which will have produced uncertainty in the R^2 values (Table 3.1) and percentages of the proportion in each section (Table 3.2). While more advanced mathematical techniques could have been used to test for the change in gradient and locate the point that gave maximum R^2 values for each section, as these figures were not used for any further quantitative testing and only sought to show a high measure of fit was possible for each section (as was successful), the simpler method was deemed appropriate. All results produced from this method were presented to two significant figures, however, and further analysis or manipulation of them beyond that done here is not advised.

The very limited sample of study sites selected to test the second hypothesis is a clear limitation to this study. Although nearly all of the results were high values (Table 3.1), there were two exceptions representing two out of the five rural locations tested - testing a greater range of rural locations would help to explain whether they were truly exceptions or a feature common to some rural areas. Only testing three scales - 75 km, 25 km and 13 km diameter - further limits the extent of the analysis. It is unclear whether even smaller scales could have produced positive results, and while the trend is generally towards lower measures of fit at smaller scales, at some locations (particularly in the upper Pareto section) this is reversed, so tests at larger scales may also have been informative. Moreover, for this and the third hypothesis, the analysis may have been limited by only studying the extremes of the rural/urban divide. Around 45% of England's urban population live in local authorities that do not belong to 'major urban' conurbations (Pateman, 2011), so the largest cities are not necessarily representative of all urban areas. With more study sites, the distribution among towns (large and small) could have been investigated, along with a sample that is more representative of the full breadth of the UK, including of Scotland, Wales and Northern Ireland.

Finally, the analysis of the third hypothesis is very simplistic. In reality, there are many more factors affecting electricity consumption than the extent to which an area is rural or urban, and this dichotomy can also often be too simple, ignoring considerable heterogeneity in rural areas (Hoggart, 1990). By studying the effect of only one variable, the analysis above risks oversimplifying the trends in electricity consumption in the UK or, at worst, conflating correlation with causation and masking other, more important factors. As mentioned in section 1.2.1, Baiocchi et al. (2015) warned about this type of analysis, arguing that factors including income, household size, housing quality and access to central heating, are too diverse and place-specific to allow conclusions to be drawn from simple correlation analysis on a large scale. The results and conclusions must therefore be treated with considerable caution, but despite these concerns the strength and clarity of the correlation give sufficient evidence to disprove the third null hypothesis.

Section 5

Conclusions

The data for electricity consumption per capita in the United Kingdom were found to fit a two-class distribution, the majority of which is exponential, with a power-law distribution demonstrated by roughly the top 1%. This fits with the literature on the distribution of wealth and income (nationally and globally) and global energy consumption per capita, but is to the best of our knowledge the first time this has been shown to apply to electricity consumption per capita in the United Kingdom. Moreover, this distribution was found to exist at all locations tested across the country, and at all spatial scales. This confirms theories that suggest that inequality in electricity consumption, as with that in money, is an inevitable result of entropy maximisation. The socio-economic implications of this finding are therefore highly significant, and suggestions for how this inevitability may be overcome, such as shifting the economy towards renewables, decentralised power generation or a complete decoupling of money and electricity, will be important areas for future research.

Electricity consumption per capita was shown to be systematically higher in rural areas as compared to urban areas, with a two-class distribution and similar measures of inequality found in both. Based on the literature this is predicted to be a result of factors such as cities having higher density housing, which increases the efficiency of heating, and more efficient (and less need for) transport. London was found to have much higher inequality present within it than other cities, with around 15% of its data demonstrating a power-law distribution compared to an average of 4% for a selection of the UK's other large cities, which is theorised to represent a much higher proportion of the population who generate their wealth from capital rather than wages.

Further research may seek to investigate whether the same trends exist for energy consumption per capita, or may use the data for 2000, 2005 and 2010 to see whether, as Lawrence et al. (2013) predict, entropy has maximised over time and the two-class distribution has become more pronounced. This would also allow analysis of how the rural-urban divide and spatial patterns of inequality have changed between 2000 and 2015. Further work may also address the issues identified in section 4.4, such as studying a wider range of locations and scales to add more rigour to the analysis of the second hypothesis, or testing the electricity consumption per capita data against other variables (such as GDP per capita, which is commonly available in a similar format) that are likely predictors of it to build on upon the analysis of spatial trends for the third hypothesis. Comparing the distribution in London to other megacities around the world to investigate whether it is unique to London, or a megacity phenomenon, would also be highly informative.

Section 6

References

Angel, S., Parent, J., Civco, D., Blei, A., Potere, D. (2011) The dimensions of global urban expansion: estimates and projections for all countries, 2000–2050. *Progress in Planning* 75(2), 53-107.

Baiocchi, G., Creutzig, F., Minx, J., Pichler, P.P. (2015) A spatial typology of human settlements and their CO2 emissions in England. *Global Environmental Change* 34, 13–21.

Banerjee, A. and Yakovenko, V. (2010) Universal patters of inequality. *New Journal of Physics* 12(7).

Cayla, J.M., Maizi, N., Marchand, C. (2011) The role of income in energy consumption behaviour: Evidence from French households data. *Energy Policy* 39(12), 7874-7883.

Center for International Earth Science Information Network (CIESIN) (2016) *Gridded Population of the World, Version 4 (GPWv4): Population Count.* http://sedac.ciesin.columbia.edu/data/collection/gpw-v4 [6.12.18]

Chakraborti, A. and Chakrabarti, B. (2000) Statistical mechanics of money: how saving propensity affects its distribution. *European Physical Journal B* 17, 167–170.

Champernowne, D. (1953) A Model of Income Distribution. *The Economic Journal* 63(250), 318-351.

Cowell, F. and Victoria-Feser, M.P. (1996) Robustness properties of inequality measures. *Econometrica: Journal of the Econometric Society* 64(1), 77-101.

de Rezende Francisco, E., Aranha, F., Zambaldi, F., Goldszmidt, R. (2007) Electricity Consumption as a Predictor of Household Income: a Spatial Statistics Approach. In: Davis, C., Monteiro, A. (eds) *Advances in Geoinformatics.* Berlin: Springer.

Department of Business, Energy and Industrial Strategy (BEIS) (2018a) *Annual fuel poverty statistics report: 2018.* London: BEIS.

Department of Business, Energy and Industrial Strategy (BEIS) (2018b) *Digest of UK Energy Statistics (DUKES) 2018.* London: BEIS.

Dragulescu, A. and Yakovenko, V. (2000) Statistical mechanics of money. *The European Physical Journal B* 17, 723–729.

Dragulescu, A. and Yakovenko, V. (2001a) Evidence for the exponential distribution of income in the USA. *The European Physical Journal B* 20, 585-589.

Dragulescu, A. and Yakovenko, V. (2001b) Exponential and power-law probability distributions of wealth and income in the United Kingdom and the United States. *Physica A* 299, 213-221.

Dragulescu, A. and Yakovenko, V. (2003) Statistical Mechanics of Money, Income, and Wealth: A Short Survey. In Garrido, P. and Marro, J. (eds.) *Modeling of Complex Systems: Seventh Granada Lectures, AIP Conference Proceedings* 661. New York, American Institute of Physics.

Ebell, M. (2017) *Regional Inequality in Productivity in the UK: A Closer Look.* London: National Institute of Economic and Social Research.

Florida, R. (2017) *The New Urban Crisis.* New York: Basic Books.

Gibrat, R. (1931) *Les inégalités économiques: applications: aux inégalitês des richesses, à la concentration des entreprises, aux populations des villes, aux statistiques des familles, etc: d'une loi nouvelle: la loi de l'effet proportionnel.* Paris: Librairie du Recueil Sirey.

Gini, C. (1912) Variabilita e Mutabilita. *Journal of the Royal Statistical Society* 76(3), 326-327.

Gini, C. (1936) On the Measure of Concentration with Special Reference to Income and Statistics. *Colorado College Publication, General Series* 208, 73–79.

Hasegawa, A., Mima, K. and Duong-van, M. (1985) Plasma distribution function in a superthermal radiation field. *Physical Review Letters* 54(24), 2608.

Hijmans, R., Kapoor, J., Wieczorek, J., Garcia, N., Maunahan, A., Rala, A., Mandel, A. (2018) *GADM v3.6.* `https://gadm.org` [5.12.18]

Hoggart, K. (1990) Let's do away with rural. *Journal of rural studies* 6(3), 245-257.

Hood, A., Water, T. (2017) *Living standards, poverty and inequality in the UK: 2016-2017 to 2021-2022.* London: Institute for Fiscal Studies.

Hui, S. (2001) Low energy building design in high density urban cities. *Renewable Energy* 24, 627-640.

International Energy Agency (IEA) (2013) *World Energy Outlook 2013.* Paris: International Energy Agency.

Jacobson, A., Milman, A., Kammen, D. (2005) Letting the (energy) Gini out of the bottle: Lorenz curves of cumulative electricity consumption and Gini coefficients as metrics of energy distribution and equity. *Energy Policy* 33, 1825-1832.

Jarvis, S. (2018) *Global Gridded Electricity (GGE) dataset.* Unpublished.

Kalecki, M. (1945) On the Gibrat Distribution. *Econometrica* 13(2), 161-170.

Koop, G., McIntyre, S., Mitchell, J., Poon, A. (2018) *Regional Output Growth in the United Kingdom: More Timely And Higher Frequency Estimates, 1970-2017.* London: Economic Statistics Centre of Excellence.

Krapivsky, P. and Redner, S. (1998) Wealth distributions in asset exchange models. *The European Physical Journal* B 2, 267–276.

Lawrence, S., Liu, Q., Yakovenko, V. (2013) Global Inequality in Energy Consumption from 1980 to 2010. *Entropy* 15, 5565-5579.

Levy, M. and Levy, H. (2003) Investment talent and the Pareto wealth distribution: Theoretical and experimental analysis. *Review of Economics and Statistics* 85(3), 709-725.

Lorenz, M. (1905) Methods of Measuring the Concentration of Wealth. *Publications of the American Statistical Association* 9(70), 209-219.

Mandelbrot, B. (1960) The Pareto-Levy law and the distribution of income. *International Economic Review* 1(2), 79-106.

Mimkes, J. (2000) Society as a many-particle system. *Journal of Thermal Analysis and Calorimetry* 60(3), 1055-1069.

Minx, J., Baiocchi, G., Wiedmann, T., Barrett, J., Creutzig, F., Feng, K., Förster, M., Pichler, P.P., Weisz, H. and Hubacek, K. (2013) Carbon footprints of cities and other human settlements in the UK. *Environmental Research Letters* 8(3), 035039.

Mohajeri, N., Gudmundssonb, A., Scartezzini, J. (2015) Statistical-thermodynamics modelling of the built environment in relation to urban ecology. *Ecological Modelling* 307, 32–47.

Norman, J., Maclean, H., Kennedy, C. (2006) Comparing High and Low Residential Density: Life-Cycle Analysis of Energy Use and Greenhouse Gas Emissions. *Journal of Urban Planning and Development* 132(1) 10-21.

O'Brien, W., Kennedy, C., Athienitis, A., Kesik, T. (2010) The Relationship between Net Energy Use and the Urban Density of Solar Buildings. *Environment and Planning B: Planning and Design* 37(6), 1002–1021.

Office for National Statistics (ONS) (2013) *2011 Census: Population Estimates by five-year age bands, and Household Estimates, for Local Authorities in the United*

Kingdom. https://www.ons.gov.uk/file?uri=/peoplepopulationandcommunity/populationandmigration/populationestimates/datasets/2011censuspopulationestimatesbyfiveyearagebandsandhouseholdestimatesforlocalauthoriti r12ukrttablep04ukv2_tcm77-304141.xls [20.12.18]

Office for National Statistics (ONS) (2014) *National Population Projections: 2014-based Statistical Bulletin.* https://www.ons.gov.uk/peoplepopulationandcommunity/populationandmigration/populationprojections/bulletins/nationalpopulationprojections/2015-10-29 [23.12.18]

Office for National Statistics (ONS) (2017a) *Built-up Areas (December 2011) Boundaries V2.* https://data.gov.uk/dataset/15e3be7f-66ed-416c-b0f2-241e87668642/built-up-areas-december-2011-boundaries-v2 [20.12.18]

Office for National Statistics (ONS) (2017b) *Rural Urban Classification (2011) of Wards in England and Wales.* http://geoportal.statistics.gov.uk/datasets/rural-urban-classification-2011-of-wards-in-england-and-wales [20.12.18]

Organisation for Economic Co-operation and Development (OECD) (2018) *Income Distribution and Poverty* https://stats.oecd.org/Index.aspx?DataSetCode=IDD [8.1.19]

Pareto, V. (1897) Cours d'Économie Politique. *The Annals of the American Academy of Political and Social Science* 9.3, 128-131.

Pateman, T. (2011) Rural and urban areas: comparing lives using rural/urban classifications. *Regional Trends* 43(1), 11-86.

Persky, J. (1992) Retrospectives: Pareto's Law. *Journal of Economic Perspectives* 6(2), 181-192.

Pigou, A. (1912) *Wealth and Welfare.* London: Macmillan.

Prendergast, L. and Staudte, R. (2016) Quantile versions of the Lorenz curve. *Electronic Journal of Statistics* 10(2), 1896-1926.

Raikes, L., Millward, L., Longlands, S. (2018) *State of the North 2018: Reprioritising the Northern Powerhouse.* Manchester: IPPR North.

Rodd, J. (1995) Pareto's Law of Income Distribution, or the 80/20 rule. Journal of Nonprofit and Voluntary Sector Marketing 1(1), 77-89.

Rode, P., Keim, C., Robazza, G., Viejo, P., Schofield, J. (2014) Cities and Energy: Urban Morphology and Residential Heat-Energy Demand. *Environment and Planning B: Planning and Design* 41(1), 138–162.

Silva, A. and Yakovenko, V. (2005) Temporal evolution of the 'thermal' and 'superthermal' income classes in the USA during 1983-2001. *Europhysics Letters* 69, 304-310.

Stanley, H., Afanasyev, V., Amaral, L., Buldyrev, S., Goldberger, A., Havlin, S., Leschhorn, H., Maass, P., Mantegna, R., Peng, C. and Prince, P. (1996) Anomalous fluctuations in the dynamics of complex systems: from DNA and physiology to econophysics. *Physica A: Statistical Mechanics and its Applications* 224(1-2), 302-321.

Wiedenhofer, D., Lenzen, M., Steinberger, J. (2013) Energy requirements of consumption: Urban form, climatic and socio-economic factors, rebounds and their policy implications. *Energy Policy* 63, 696–707.

Wilkinson, R., Pickett, K. (2009) *The Spirit Level: Why More Equal Societies Almost Always Do Better.* London: Allen Lane.

World Energy Council (2013) *World Energy Resources: 2013 Survey.* London: World Energy Council.

Wu, Q., Maslyuk, S., Clulow, V. (2012) Energy consumption inequality and human development. In Morvaj, Z (ed.) *Energy Efficiency-A Bridge to Low Carbon Economy.* London: InTech.

Yakovenko, V. and Rosser, J. (2009) *Colloquium: Statistical mechanics of money, wealth, and income.* Reviews of Modern Physics 81(4), 1703-1725.

Yakovenko, V. and Rosser, J. (2009) Colloquium: Statistical mechanics of money, wealth and income. *Reviews of Modern Physics* 81(4), 1703.

Zoega, G. (2016) *On the causes of Brexit: Regional differences in economic prosperity and voting behaviour.* `https://voxeu.org/article/brexit-economic-prosperity-and-voting-behaviour` [24.12.18]