Temperature as a Deep Determinant of Income Inequality in a Cross-Country Study

ERASMUS UNIVERSITY ROTTERDAM

Erasmus School of Economics

Master Thesis [programme International Economics]

Name student: Ruben Lehtonen

Student ID number: 423018rl

Supervisor: Dr. Anna Baiardi

Second assessor: Dr. Aart Gerritsen

Date final version: April 3, 2021

Abstract

We conduct an empirical analysis on the effect of temperature on net inequality in a long-term cross-country setting, asking why it is that modern welfare states have appeared in colder climates. We find that long-term average minimum temperature alone accounts for 32% of the cross-country variation in net Gini, and then controlling for potential confounders, find a 1°C increase in minimum temperature associated with a 0.50-point increase in the Gini coefficient on a scale from 0 to 100. A strong effect of temperature remains even after adding continent dummies, showing that the effect also holds in otherwise more similar regions. In exploration of the intermediate steps, we infer that redistribution and taxes are the primary channel through which temperature has its effect on net inequality. On the other hand, the effect does not appear to be through precipitation. The coefficients for temperature are statistically significant at the 1% level throughout the different model specifications. We start by explaining why it is important to study such deep determinants of inequality separately from those of income and growth in the development economics literature.

1. Introduction

What are the “deep determinants” of net income inequality, the factors that underlie cross-country differences in preferences for redistribution, government policies, and other more immediate determinants of the post-tax and transfer income distribution? This paper suggests one such determinant: temperature, or more specifically the average minimum daily temperature in the winter months. As cold is “the great executioner of nature” (Masters & McMillan, 2001), it is presumed that the cold has played a role in the necessity of building various social support systems, in the form of modern welfare states and other government and private institutions. Could it be that for instance the Nordic model is Nordic due to the nature of the region being cold?

Our research question is formulated as follows:

Does colder temperature lead to lower net inequality?

We choose net inequality (meaning inequality after taxes and transfers) because it captures both differences in redistribution and government policies that are not redistributive but nevertheless affect market inequality (inequality before taxes and transfers). This relationship between the variables is shown in Figure 2, which is borrowed from Berg et al. (2018) who underscore the need to clearly define these variables in such a manner. This will be explained in more detail in the next section.

It is important to note that inequality is not bad in itself, and we are not making a normative judgment on any particular levels of inequality being better or worse than others. In fact, laboratory experiments have found that the relevant variable for normative judgments is *fairness*: people prefer a fair distribution to an equal distribution (equal meaning everyone would receive the same amount), where fairness is evaluated according to things like levels of effort, ability, and disability (Starmans et al., 2017). There is an important subjective component to fairness: what one person considers fair may not seem fair to another. Isaksson & Lindskog (2009) find that, across countries, there is large variation in *beliefs of what determines income,* and that these different beliefs are important in explaining preferences for redistribution. In other words, different levels of inequality are seen as optimal depending on who you ask, based on their beliefs upon the fairness of that inequality, based in turn on what they believe causes that inequality – e.g. if they believe it is due to effort, inheritance, intelligence, etc.

This is to say, when studying the determinants of inequality, we are at least in part studying the determinants of fairness evaluations: To the extent that, over time, the income distribution of a country approaches the level that is seen to be fair in that country, we are proxying for differences in ideas of fairness by looking at inequality. As we cannot make normative judgments of normative judgments, an element of our question is whether temperature is an important determinant of these fairness beliefs, the judgments which then manifest as part of the differences we see across countries in their income distributions in the usual measures such as the Gini coefficient.

Figure 1 illustrates these notions with an example of two countries with identical realized income distributions, yet holding different normative implications. The Lorenz curve describes the cumulative share of the total income in the country that falls to the fraction of the population measured on the x-axis. This means that a 45-degree line would represent complete income equality, where each person has the same income, and conversely a line that is closer to the intersection of the two axes represents progressively more inequality. In this example, inequality is identical in both countries, yet it is *desired* to be higher in country 1, but lower in country 2.

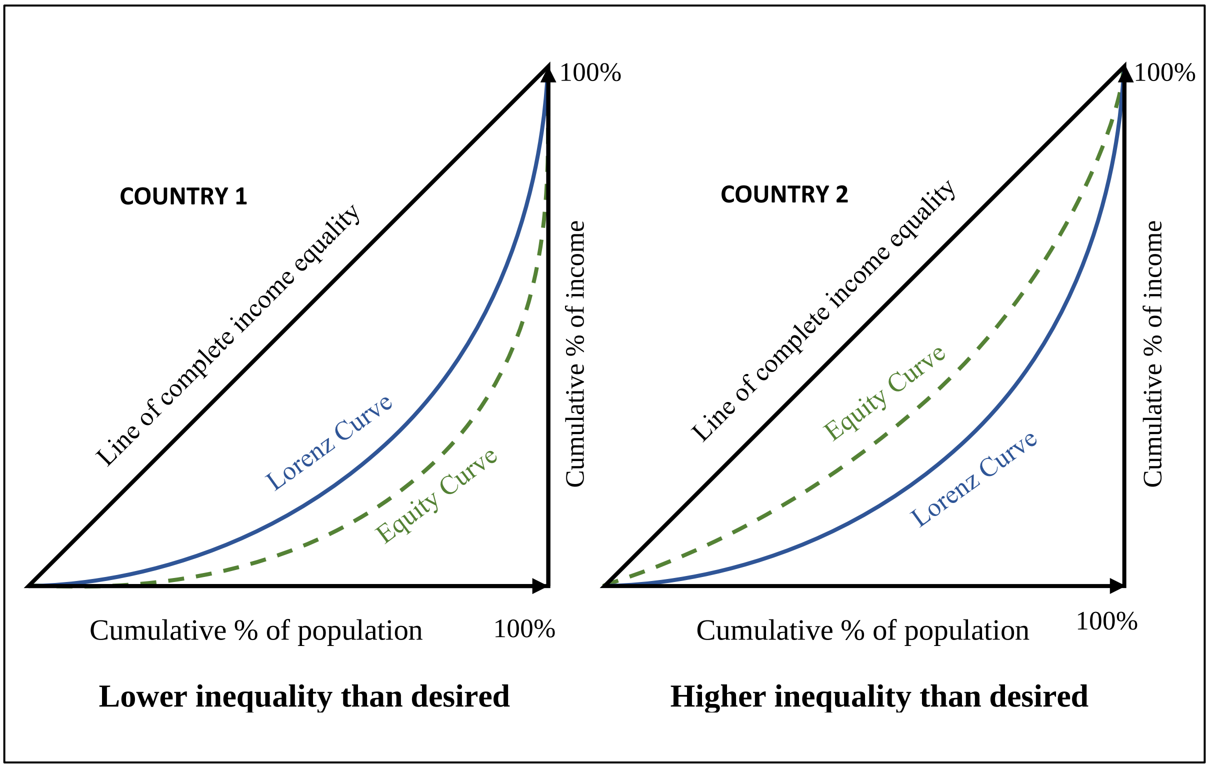

Figure 1 Illustration of the relationship between the actual income distribution (Lorenz curve) and one that would be considered equitable (fair) by the average person in the population of the country (equity/fairness curve).

One may ask why not study fairness directly, to see if temperature shifts the equity curve in this figure, to which we answer that we do not assume that fairness is the *only* thing that causes inequality, and we also do not assume ex ante that any effect of temperature on inequality would be fully due to fairness perceptions. There is a long list of potential determinants of redistributive preferences: fairness, the structure of the family, individual characteristics, perceptions of inequality, social mobility, religion, cognitive biases, etc. (Cruces et al. (2013), Alesina & Giuliano, (2011)). So being, there may be channels that are unrelated to fairness that cause the Lorenz curve to be distanced from the fairness curve. We do not want to assume that an effect we find of temperature is all through fairness, although it does appear possible that these other factors are components of different fairness perceptions.

This discussion has been to show that motivational structures may exist that maintain different inequality equilibria across different countries, in order to convey that our goal cannot be to prescribe any low levels of inequality to all countries, but to explore what causes countries to have different inequality outcomes in the first place. Such inequality equilibria have been proposed by Alesina & Angeletos (2005), (Bénabou & Tirole, 2006), and others. In summary, we are interested in the causes behind different countries having these different income distributions, be that through actual differences in the desires for such distributions but also for reasons unrelated to fairness ideals.

We approach our research question mainly in the context of the *total effect*, and we control for potential confounders that might harm this effect, namely altitude, a dummy for whether or not the country is landlocked, and dummy variables for which continent the country is located in. To address concerns of endogeneity, we argue that confounding of the effect should not be an issue after adding these controls, due to temperature being largely exogenous by nature and hence only a few such biasing "backdoor paths" exist (as per the backdoor criterion from Pearl (2009)). On the other hand, the intermediate structure between temperature and net inequality contains much more uncertainty as is reflected in prior research, but we do attempt to block some likely mechanisms to learn about the strengths of these mechanisms, such as the importance of redistribution in mediating the effect. Thus, the primary focus and contribution is the identification of a total effect of temperature on net inequality, largely abstracting away from how much of the effect is direct/indirect or historical/modern, meaning how much of it is through human capital, culture, institutions, and/or other unknown intermediate steps. To our knowledge no previous paper has connected temperature to cross-county income distributions, so even showing that a total effect exists may inspire future research to decouple more intermediate mechanisms in a stepwise approach.

There is a large amount of research on the "deep determinants" of income levels and growth, measured usually by GDP. Why is it important then to study inequality separately from income? If inequality would be perfectly predictable from income or growth, there would be no need to approach the deep determinants of inequality separately. However, and importantly, attempts to find robust theoretical and empirical connections between inequality and income have been inconclusive. The famous Kuznets curve, a bell-shaped hypothesized relationship between income and inequality, with inequality initially rising with income and then falling, became viewed as somewhat of an accepted law in economics. Later empirical tests show mixed results, with Barro (2000) finding that the curve does indeed emerge in a cross-section of countries albeit not fitting the data very well, but it cannot predict variation in a panel study approach across time: for instance, the East Asian Miracle showed countries growing rapidly while actually decreasing inequality at the same time with indications that less inequality actually promoted the growth (Stiglitz, 1996), and the same is true for many other Asian and sub-Saharan countries (Palma, 2011). Li et al. (1998) find that inequality is remarkably stable over time, that it is determined by factors that change only slowly over time but which have large differences across countries, and argue strongly against the Kuznet's curve of simply matching inequality to income.

More closely related to the current study, papers by Engerman and Sokoloff (e.g. Sokoloff & Engerman, 2000) hypothesized that geographic preconditions were an important determinant of slave use in colonies, this slave use then raised economic inequality, which in turn lowered later economic development in the income of these regions. This hypothesis has found support in empirical studies, except for the aspect of inequality: empirical evaluations by Nunn (2008) and (Summerhill, 2010) find that indeed regions with slavery show lower subsequent development, but this did not appear to happen through higher inequalities. Nunn also notes that income is unrelated to inequality both in current times and historically. These are further examples of the difficulty of connecting or predicting inequality from just income, serving as motivation for why determinants of inequality must be studied separately.

The question and the approach have important academic and policy relevance. Economic inequality has been suggested as a variable linked to a multitude of important outcomes, ranging from health outcomes and social cohesion to GDP growth and distributive efficiency in society. Despite this, the search for fundamental "deeper determinants" of inequality is very scarce, nearly non-existent. Such knowledge of underlying largely exogenous forces determining inequality may guide future research towards deliberate new empirical questions to uncover intermediate steps and relevant actionable variables, build structural-based models, and guide policy towards better understanding of perhaps disproportionate or different types of efforts required in some areas of the world compared to others in pursuit of set objectives. As is attributed to Winston Churchill, the following insight should capture the spirit of motivation for this research:

"The longer you can look back, the farther you can look forward."

We find that temperature, altitude, and being landlocked or not explains 52% of the variation in net Gini (this percentage being the adjusted R-squared measure), with a 1°C increase in our average minimum temperature measure associated with a 0.50-point increase in the Gini on a scale from 0 to 100, and significant at the 1%-level. This result is obtained with observations for 165 countries, which are listed in the Appendix (Table 5). When eight continent region dummies are added to identify the strength of the effect within continents instead of across the whole world, the coefficient drops to 0.24 while still maintaining the same significance level, and the adjusted R-squared is 60%. A log-linear regression or adding a quadratic term does not change the fit of the model much.

It is remarkable to us that temperature alone explains a third of the variation in inequalities between countries, and a few additional controls over half of the variation. The fact that adding the continent dummies maintains the significance level of the total effect is important, as it shows that temperature variation across countries with otherwise more similar environments and histories, and all that comes with being located on the same continent, still shows a clear effect of temperature on net inequality. Being landlocked appears as an insignificant coefficient even at the 10%-level, but altitude is positive and significant even at the 1%-level. The coefficient of altitude does not appear to have a useful interpretation by itself however, as it would mean that holding temperature and being landlocked constant, while raising altitude, would raise the net Gini coefficient. Since altitude can only be a determinant of temperature but not vice versa, such a situation where temperature is consistently the same at significantly different altitudes is likely to be rare or at least not of practical use, and it makes more sense to look at the effect of altitude through its effect on temperature, as we do by controlling for it.

As additional analysis, we control for precipitation, which appears as an insignificant coefficient with the value zero, so precipitation does not appear to generate the effect we are finding. We also control for preferences for redistribution, albeit with a dataset that severely limits the number of countries to only 28. The coefficient for preferences for redistribution is 0.32, but not significant even at the 10% level. We address the data limitation issue by controlling for market Gini instead, for which we have 165 observations, and note that the coefficient for temperature does not change much. This means that the effect of temperature must be through redistribution and taxes, by definition of net Gini being market Gini minus redistribution (Figure 2).

In the next section, we introduce the relevant literature towards understanding what has been done earlier regarding temperature and economics, as well as in some closely related fields relevant for guiding our analysis. Section 3 describes the data we use and how it is modified, Section 4 describes the choice of methods and reasoning behind them, results are presented and interpreted in Section 5, and Section 6 concludes.

2. Theoretical and Empirical Framework

Let us briefly outline the “deep determinants” literature in the field of economic development, as our goal is closely related in having just a different focus, inequality instead of income. This literature is a search for underlying explanations behind the more traditionally known immediate causes of country income differences, such as factor endowments, factor productivities, and innovation (Spolaore & Wacziarg, 2013). In other words, the question is what causes countries to have differences in these more proximate factors in the first place. In this field, multiple classes of historical effects have been identified, ranging from colonial rule, the French Revolution, and African slave trade all the way to the Neolithic Revolution, the transition from hunting and gathering to agriculture (Nunn, 2014). It has also long been established that hot countries tend to be poor: Dell et al. (2009) find a negative relationship between temperature and income, both when using a cross-sectional approach and panel data, finding that temperature alone explains 23% of differences in income per capita across countries.

There is a long line of such studies documenting a correlation between geographic and biological factors and income per capita, and the correlation is well established. In particular, favorable biological and geographical preconditions have been identified as fundamental determinants of development through explaining the earlier timing of the Neolithic Revolution in certain regions of the world compared to others due to the availability of domesticable plants and animals, as well as the spreading of the ensuing technologies to other regions more readily on an East-West oriented continent (Eurasia) than on North-South oriented continents. These advantages include metallurgy and writing, but also germ immunity due to close living conditions with livestock. These hypotheses were famously put forward by Jared Diamond (1997), and have thereafter found support in data by Olsson & Hibbs (2005), (Ashraf & Michalopoulos, 2011) and others.

The Diamond (1997) hypothesis can be roughly summarized as follows:

1) An East-West orientation of continents was beneficial for the spreading of domesticated plants, animals, and innovation from their original locations to others. Countries on North-South oriented continents did not receive as much influence from elsewhere due to the difficulty of climate differences on such an orientation.
2) The initial endowment of domesticable plants and animals was significantly different around the world, being most favorable combination of both in the Fertile Crescent.

According to Diamond, these deep determinants further generated the dominance of Eurasia over the rest of the world, via an earlier shift to agriculture and the *Guns, Germs, and Steel* that this eventually brought along, as is the name of Diamond's popular book. Spolaore & Wacziarg (2013) both summarize some of this literature and use a unified dataset finding that a handful of geographic variables (latitude, share of land area in the tropics, dummies for being landlocked or an island) together account for 44% of the variation in the cross-country per capita income of 2005. Of these variables, the most important is latitude, and omitting it reduces the adjusted R squared to 0.29.

However, the correct causal interpretation of such deep effects is not clear, namely the question of which intermediate mechanisms carry this effect from geography to current outcomes, and/or if there is an ongoing effect of geography. In other words, the consensus is that there is at least an indirect effect of historical preconditions on current outcomes, but competing studies claim various strengths of intermediate variables such as institution, human capital, trade openness, an ongoing direct effect of geography, etc. (Spolaore & Wacziarg, 2013). In this area, the "geography vs. institutions" debate is particularly noteworthy, as it has centered around the notions that the full effect is through either institutions, geography, or some other geography-oriented variables besides institutions such as differences in culture.

For current purposes, the important takeaway from this literature is the insight that there exist deep determinants for current outcomes in income/capita, the difficulty of untangling the intermediate steps, and the notion that geographical variables are one of the most important of these deep determinants. This informs the primary scope of our paper to the study of the *total effect*, as was mentioned earlier and will be further developed in the methodology section.

We will now explore literature on the determinants of inequality, which we also briefly mentioned in the Introduction. This literature is primarily centered around the motivation to study *growth*: inequality is usually introduced as a variable which through some means affects growth in GDP. This connection has very conflicting findings, with studies showing both positive and negative effects of inequality on growth. Nevertheless, a key contribution in the field by Berg et al. (2018) informs our study tremendously by consolidating this vast literature into a cohesive understanding of the relationships between market inequality, redistribution, net inequality, and growth in a logical framework, as well as the importance of always specifying exactly which one of these variables is being studied and arguing for why the chosen data adequately captures this variable. This framework from their paper is shown in Figure 2 below.

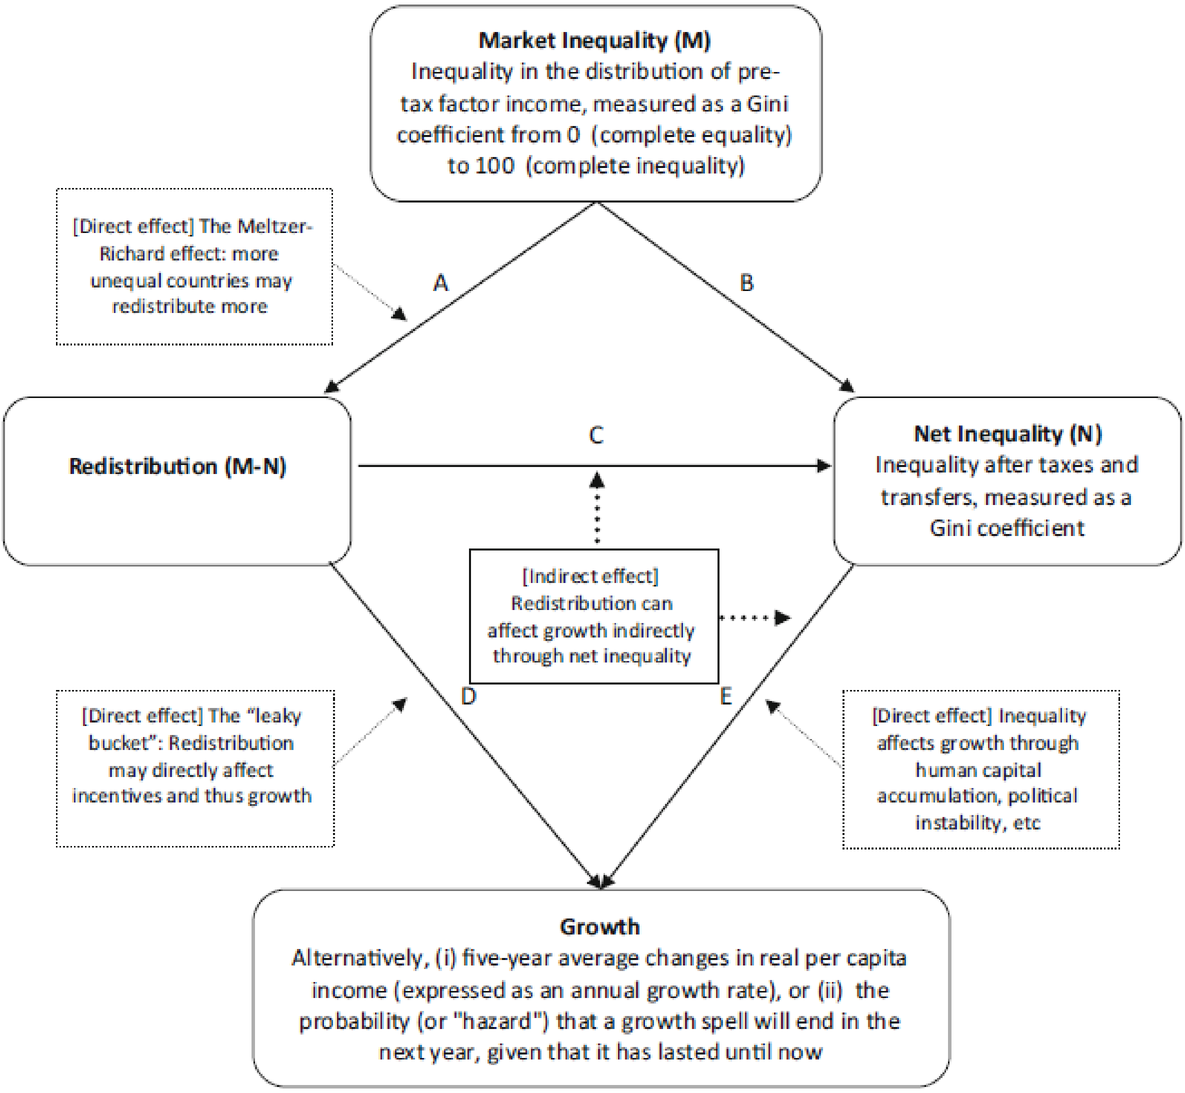

Figure 2 Copied from Berg et al. (2018), effectively showing the relationships between market inequality (M), net inequality (N), redistribution (M-N), and growth which are generally left vague and disconnected in prior research (Berg et al., 2018).

As we can see, redistribution is defined as the difference between market and net inequality. However, in addition to welfare state redistribution, the government also affects market inequality with measures such as health spending and education (Berg et al., 2018). This is a very important insight for our study, as it means we do not want to choose redistribution as our dependent variable, but instead net inequality, as we want to capture both the effects of redistributive and non-redistributive policies. Net inequality captures both welfare state redistribution (and social insurance) and non-fiscal measures that alter inequality before taxes and transfers, as Figure 1 illustrates. Things such as minimum wages, health care and education policies, tariffs and subsidies have an effect on market inequality (Berg et al., 2018), and we are also interested in capturing these. This is contained in the notion that the welfare state consists of two aspects: redistribution and social insurance. Some policies are mostly redistribution, such as the progressive income tax, whereas others mostly insure against adverse life events, such as unemployment insurance (Alesina & Giuliano, 2011).

There is a sizeable literature on the determinants of redistribution, which is often measured as preferences for redistribution from various survey results. An example is Alesina & Giuliano (2011), in which they make explicit the assumption that preferences for redistribution are eventually aggregated via the political process into policy outcomes. Afonso et al. (2010), find that there is variation in the efficiency with which different countries' spending has the desired result of lowering net inequality, the efficiency for example being higher in countries with better educational attainments.

The only paper that has studied a deeper determinant of inequality that we know of is that by Putterman & Weil (2010), and even there it is a minor excursion from their main topic which is the long-run determinants of GDP. They find that countries with populations that have more heterogeneity in their ancestral development history of experience with agriculture and organization of states (due to population movements from more and less historically developed areas) show higher levels of inequality. This heterogeneity in ancestry is a better predictor of income inequality than current ethnic or linguistic diversity. As further support for this finding, they show that the ranking of incomes among ethnic or racial groups follows these groups' early histories. They use a Gini measure, but they do not mention whether it is a measure for net inequality or market inequality, so this paper is also a good example of the importance of the insights in Figure 2 and the importance of always making these distinctions clear. Coincidentally, this paper by Putterman & Weil (2010) makes a very important contribution to the deep determinants literature in its main focus, by highlighting the need to adjust for population movements, and creating a matrix that traces the ancestry of the current populations for all of countries of the world. They find that recalculating the effect of early agriculture and state history on current GDP while accounting for these population movements greatly improves the fit of those regressions.

The findings of this section are consolidated in Figure 3 as a representation of this prior research, and to act as a foundation for the subsequent methodology and empirical analysis sections. It is not intended that this is a complete and necessarily accurate representation of the true data generation processes, but it places the current study into context while displaying the assumptions that we use in the further analysis. For our empirical purposes, the model only needs to be correct in some assumptions relevant to eliminating confounding in the effect of temperature on inequality and does not need to correctly capture and place all the intermediate mechanisms. Fortunately, as will be elaborated in the Methodology section, these requirements for the current research question place only relatively lenient demands on the model's accuracy.

3. Methodology

Looking at Figure 3, we see that it is possible to determine the effect of temperature on net inequality with purely observational data. As there are very few backdoor paths from minimum temperature to our outcome variable (as there are only a few arrows leading to temperature which could cause confounding), we can control for these in our regressions. This is echoed by Dell et al. (2014) and Rodrik et al. (2004), namely that geographic long-term analysis is a straightforward task empirically.

There is a related literature that focuses on the effect of the climate on economic outcomes, that has shifted from a focus on long-term correlations to using panel data in estimating the effect of climate change (Dell et al., 2014). This research looks at how changes in temperature affect the economy, largely motivated by global warming. Dell et al. (2014) summarize key developments in the field, noting that the focus has shifted to using panel data due to the ability to then use exogenous variation in weather to decouple its effects from potential confounders. However, they also note that if the interest is a long-run historical effect of temperature, then it is better to use cross-sectional data without controlling for intermediate mechanisms such as institutions, to avoid blocking these intermediate mechanisms.

The above statement of course relies on the important assumptions in Figure 3 being true. Luckily, we only need a small part of the model to be true. Namely, if there are determinants of temperature, which are also correlated with the other variables or their in turn unknown or otherwise omitted determinants, then these would create biasing paths in measuring our total effect of interest.

Latitude is clearly a determinant of temperature, but since temperature largely is the channel through which latitude causes the different habitats of the world, latitude should not cause a biasing path. Instead, temperature could be thought of as an instrumental variable for latitude. However, altitude, being landlocked, and ocean currents are variables that affect temperature

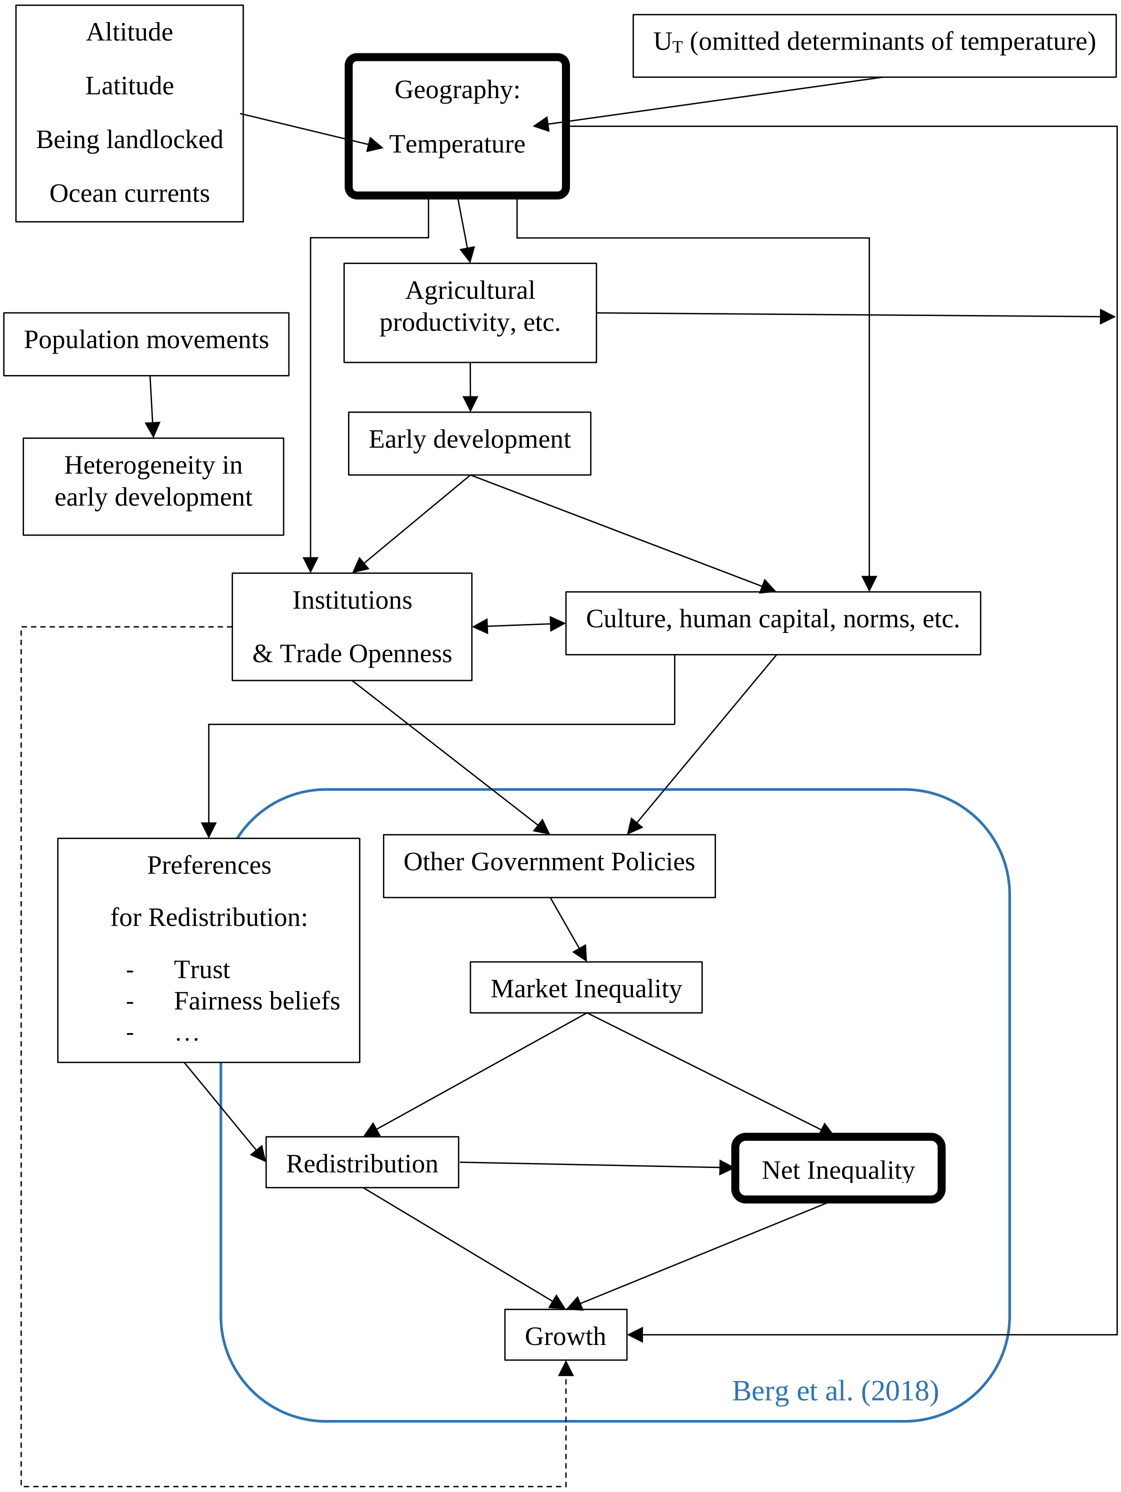

Figure 3 Graph of the prior literature made with simplifications, omitting arrows that are not important to the current study, such as suggested channels from institutions and trade openness to growth, in addition to unknown determinants of all the variables that are shown.

and may also have effects on other variables towards net inequality, so omitting them would bias any results we obtain. Dell et al. (2009) control for elevation, slope, and distance to sea, and we will do likewise, although we omit slope as their paper studies the level of municipalities and slope is not relevant at a country-level. In addition, The U_T in Figure 2 denotes further unknown determinants of temperature. If these other unknown factors only act through temperature, as drawn in the figure, then they follow the same empirical interpretation as latitude and should not cause an issue. However, if there is (in the true structural causal model) arrows from ocean currents and U_T to other variables or their unknown determinants, then these would disallow us from doing causal inference if we cannot close these backdoor channels. In other words, to the extent that we control for confounding paths with elevation and being landlocked, our regression should be robust to biasing paths. We cannot find country-level measures on the strength of ocean currents, so we assume that ocean currents have acted primarily through temperature, although they might have had a separate effect through things such as travelling by sea, or available seafood, etc. The other U_T we by definition do not know, and we assume that this is an empty group or else only acts through temperature. These are the assumptions we make for causal inference.

It is reminded that we are interested in the *total effect* of temperature on net inequality. Adding a multitude of controls to this main regression would close pathways from temperature to inequality, which would be detrimental in the goal of identifying the total effect. In addition, these controls would be practically arbitrary, as very little consensus exists on the intermediate steps from deep determinants to current outcomes, as for example geography vs. institutions-debate shows. However, we do pose sub-research questions and add two further controls with the goal of isolating particular channels within the total effect. Still, this is primarily left for future research, where perhaps machine learning algorithms can make solving such tasks easier, as conditional independencies of potentially thousands of variables might be possible to evaluate with data, in the process learning a model that matches all of the realities of that data.

The additional controls we add are precipitation, preferences for redistribution, and market Gini. By controlling for these variables incrementally, we can observe whether or not the coefficient for temperature changes and if so, into what direction, and from this infer whether they are important nodes in this mediating structure from temperature to inequalitys. Let us first examine our main empirical specifications, namely the one for answering our primary research question:

$$Net\ GINI = \beta_0 + \ \beta_1 * Temperature + \beta_2 * Elevation + \beta_3 * Landlocked + \ \beta_4 * Region\ Dummies + \varepsilon,$$

We estimate this with ordinary least squares regression, as is often done in the literature on inequality (e.g. (Afonso et al., 2010; Dell et al., 2009). We also add a quadratic term for temperature based on an initial observation of the scatterplot (Figure 4), because it appears the relationship could be quadratic. Dell et al. (2009) use logarithms for their dependent variable when studying the effect of temperature on income, and we will also do so and see if the fit is improved, as follows:

$$LOGNet\ GINI = \beta_0 + \ \beta_1 * Temperature + \beta_2 * Elevation + \beta_3 * Landlocked\ \beta_4 * Region\ Dummies + \varepsilon$$

Then, we add the intermediate controls for precipitation, redistribution preferences, and market Gini to the regression. We use heteroskedasticity-robust standard errors throughout, although the data does not seem to exhibit heteroskedasticity when looking at the scatterplots (Figure 4).

Next, we will describe how we obtain and modify the data used in order to implement these methods in answering the research question.

4. Data

We obtain net Gini and market Gini data from the Standardized World Income Inequality Database (SWIID) compiled by Solt (2020), designed for comparable cross-national research and incorporating data from the OECD, World Bank, Eurostat, the UN, and others. This data is made fully with actual measurements and not imputations. The data is available for 196 countries, with various years from 1960 to the present depending on the country. We create our variable “Net Gini” by taking an average for each country of the most recent 10 years, or less for those countries without ten years of observations. The motivation for this is twofold: inequality is remarkably stable over time (Li et al., 1998), and the focus of this of this study is on the long run, so taking an average of the years does not bias the GINI measures while being more robust to outlier years than picking a year at random.

Temperature data is from the Climate Research Unit (CRU) at the University of East Anglia. We use their “CRU CY 4.04” dataset, and the “tmn” variable within it, which stands for monthly average daily minimum temperature in Celsius. The data is provided for 290 countries over the years 1901-2019, and precalculated variables “DJF”, “MAM”, “JJA”, “SON” are included, the first of which measures the average daily minimum temperature of December, January, and February for that year, and the others similarly for the other months of the year. This suits our needs well, as taking an average over the coldest months of the year is more likely to capture climate (source). However, as the seasons of the year occur at different times e.g. in New Zealand compared to Finland, we need to use the June-July-August measure for New Zealand and the December-January-February one for Finland. For this reason, we choose the smallest of the four variables for each year, as we are interested in cold temperatures. For each country, an average is taken over 30 years, as this has been a standard in defining climate, although rather arbitrary and many different timeframes could be selected (Hsiang, 2016). Because the December-January-February observation is not available for 2019 as January and February would go into 2020, we remove the year 2019 from all countries' observations.

To control for confounding, we add data for altitude, being landlocked, and region dummies: Europe & Central Asia, Latin America & Caribbean, Middle East & North Africa, North America, South Asia, Sub-Saharan Africa, and East Asia & Pacific. Altitude data is obtained through the R "raster" package from the Shuttle Radar Topography Mission (SRTM) 90 m resolution dataset, referenced as (Jarvis et al., 2008). This altitude data is available for 171 of our countries. The data used to create a dummy variable for being landlocked is from CEPII geo_cepii dataset which is available in R through the "cepiigeodist" library and is referenced as Mayer & Zignago (2011).

The above variables are the ones of primary importance, as they relate to our main research question. However, we also attempt to examine the strength of some intermediate causal channels suggested in the literature. One of these controls is preferences for redistribution, for which Alesina & Giuliano (2011) use the following question from the World Value Surveys (WVS), on which people answer a number in the range between 1 and 10: 'People should take more responsibility to provide for themselves' (1), and 'The government should take more responsibility to ensure that everyone is provided for' (10). This question is a poor measure of redistributive preferences in a sense that would be useful for our current study. The downfall is the use of the word *more*, which means that the answer will be endogenous to the already existing social support systems, rather than being a logical precursor to redistribution as the authors of that paper argue. Instead, we use The International Social Survey Programme (ISSP), such as Scheve & Stasavage (2005). We use the 2016 Role of Government V edition, where the relevant question is as follows: "It is the responsibility of the government to reduce the differences in income between people with high incomes and those with low incomes." Values for this variable range from 1 to 4: (1) Definitely should be, (2) Probably should be, (3) Probably should not be, and (4) Definitely should not be. This data is only available for 35 countries, which further drops to 28 when it is matched with our other data, so regressions that include it are more subject to error due to the limited number of observations in this factor.

The other intermediate variable we control for is precipitation, which we obtain also through the R raster package, this time from WorldClim, which is a modified use of the “CRU-TS-4.03” from the Climatic Research Unit, with a spatial resolution of 5 minutes and referenced as (Fick & Hijmans, 2017). The precipitation data is the average of the years 1970-2000, and we choose the total precipitation for the coldest quarter of the year.

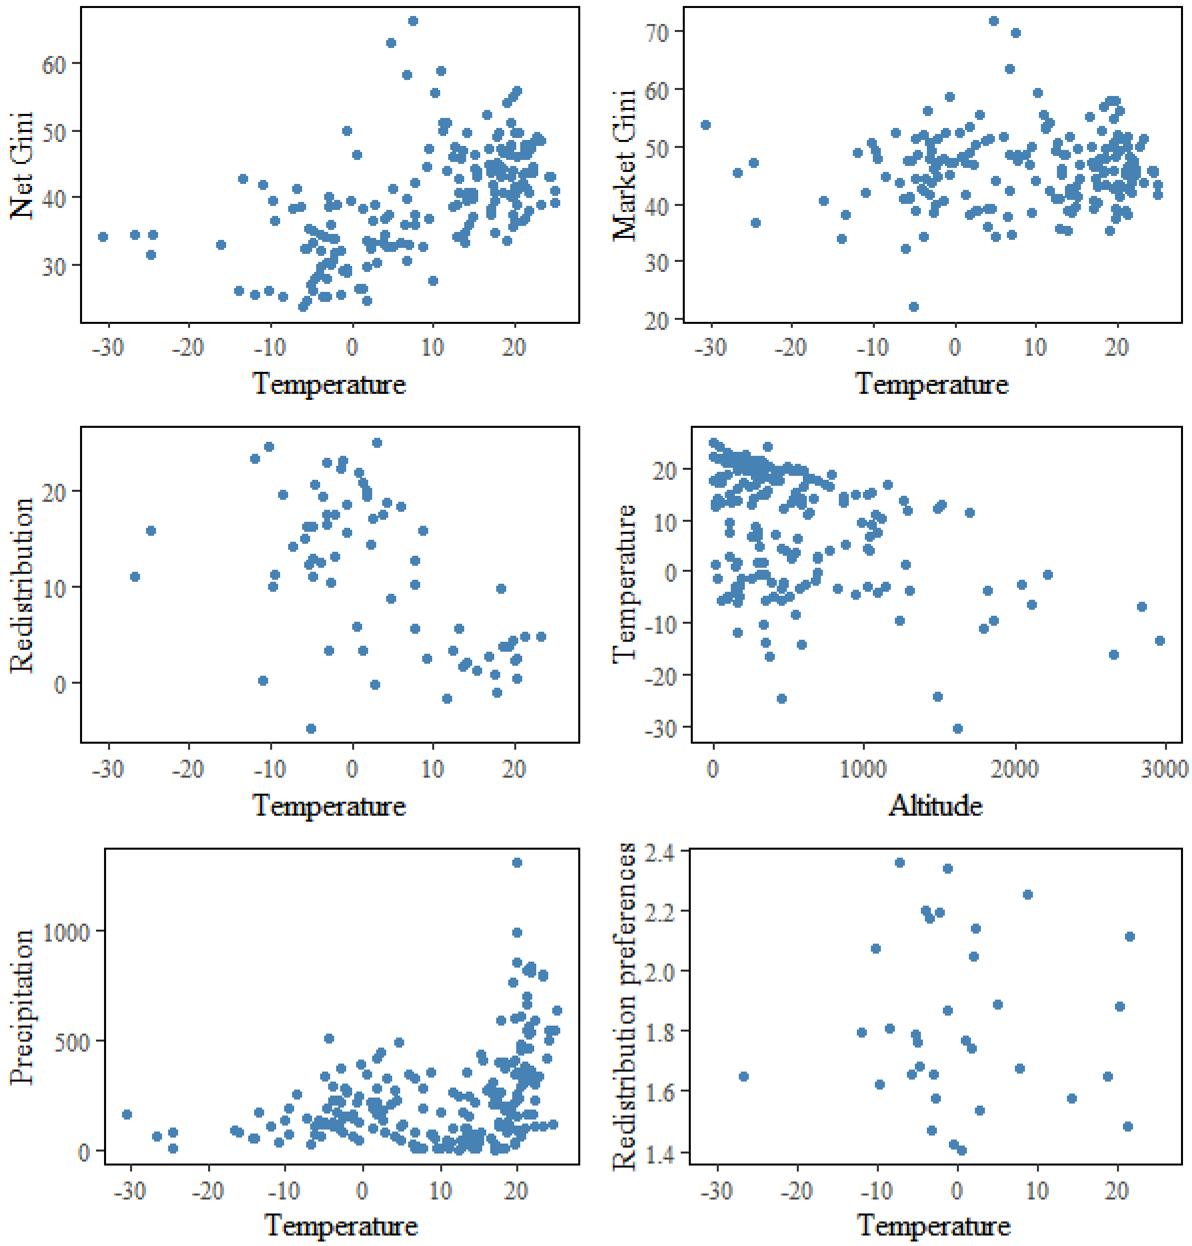

Figure 4 Scatterplots of our variables, the most important being the first on the top-left with Net Gini and Temperature, where a positive correlation is visible. This is not the case for Market Gini in the second graph.

Figure 2 shows scatterplots of the variables for preliminary analysis, the most important being the one with Net Gini on Temperature. We can see there is a clear positive correlation between temperature and Net Gini, whereas no such relationship exists for Market Gini. Temperature Table 1 further presents correlation coefficients for these variables, and Table 2 summary statistics. A list of the 165 countries with available data for all of the main regressions, as well as the 28 the countries for the redistribution preferences is available in the Appendix.

Table 1: Correlation Matrix

	Net Gini	Market Gini	Temperature	Landlocked	Altitude	Precipitation	Redistribution preferences
Net Gini	1	0.575	0.563	-0.037	0.072	0.157	-0.185
Market Gini	0.575	1	0.062	-0.023	-0.035	0.077	0.084
Temperature	0.563	0.062	1	-0.286	-0.419	0.380	-0.029
Landlocked	-0.037	-0.023	-0.286	1	0.488	-0.262	0.136
Altitude	0.072	-0.035	-0.419	0.488	1	-0.208	-0.153
Precipitation	0.157	0.077	0.380	-0.262	-0.208	1	-0.024
Redistribution preferences	-0.185	0.084	-0.029	0.136	-0.153	-0.024	1

Table 2: Summary Statistics

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max
Net Gini	196	39.505	7.909	23.480	33.690	44.508	66.300
Market Gini	196	45.792	6.445	22.000	41.465	49.475	71.840
Temperature	217	9.405	11.693	-30.587	-0.613	19.513	25.160
Landlocked	211	0.171	0.377	0.000	0.000	0.000	1.000
Altitude	196	533.281	528.831	4.756	179.567	672.493	2,954.920
Precipitation	220	234.905	213.622	0.031	79.195	334.375	1,306.660
Redistribution preferences	35	1.822	0.269	1.405	1.647	2.062	2.356

5. Results

Let us now examine our results, which are presented in Table 3. We find that temperature alone accounts for around 32% of the variation in Net Gini, and when adding altitude the two variables together explain around 52%. The statistical significances for the coefficients on

Table 3: Main Results and Intermediate Mechanisms

	Dependent variable: Net Gini							
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**	**(7)**	**(8)**
Temperature	0.336***	0.488***	0.358***	0.495***	0.227***	0.235***	0.317**	0.186***
	(0.037)	(0.044)	(0.039)	(0.045)	(0.054)	(0.059)	(0.156)	(0.058)
Temperature^2	0.005**	0.006**	0.006**	0.006**	0.003	0.004	0.012	0.005*
	(0.002)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.009)	(0.003)
Altitude		0.006***		0.006***	0.004***	0.004***	0.007**	0.004***
		(0.001)		(0.001)	(0.001)	(0.001)	(0.003)	(0.001)
Landlocked			2.881*	0.144	-0.779	-0.822	-4.354**	-0.080
			(1.476)	(1.481)	(1.335)	(1.368)	(1.889)	(0.908)
Europe & Central Asia					-4.207**	-4.207**	-4.093**	-1.166
					(1.711)	(1.729)	(2.953)	(1.611)
Latin America & Caribbean					4.725***	4.725***	4.697***	4.425
					(1.095)	(1.115)	(4.728)	(0.840)
Middle East & North Africa					1.155	1.155	1.083	1.131
					(1.619)	(1.647)	(2.381)	(1.191)
North America					-0.492	-0.429		-5.034*
					(2.043)	(2.044)		(2.651)
South Asia					-0.128	-0.161	7.488***	0.915
					(2.712)	(2.727)	(2.698)	(1.669)
Sub-Saharan Africa					5.390***	5.310***		2.342***
					(1.384)	(1.412)		(0.898)
Precipitation						-0.001	-0.011	-0.001
						(0.002)	(0.009)	(0.001)
Redistribution preferences							-2.637	
							(3.178)	
Market Gini								0.630***
								(0.079)
Constant	35.312***	30.533***	34.400***	30.389***	33.436***	33.442***	36.051***	6.734*
	(0.808)	(0.707)	(0.913)	(0.743)	(1.765)	(1.773)	(7.492)	(3.578)
Observations	187	171	179	165	165	165	28	165
R^2	0.329	0.525	0.345	0.531	0.622	0.622	0.765	0.833
Adjusted R^2	0.322	0.517	0.333	0.519	0.597	0.595	0.626	0.820
Residual Std. Error	6.558 (df = 184)	5.525 (df = 167)	6.553 (df = 175)	5.561 (df = 160)	5.090 (df = 154)	5.105 (df = 153)	4.314 (df = 17)	3.403 (df = 152)
F Statistic	45.089*** (df = 2; 184)	61.592*** (df = 3; 167)	30.661*** (df = 3; 175)	45.232*** (df = 4; 160)	25.298*** (df = 10; 154)	22.870*** (df = 11; 153)	5.523*** (df = 10; 17)	63.203*** (df = 12; 152)

Notes: ***Significant at the 1 percent level.; **Significant at the 5 percent level.; *Significant at the 10 percent level. Region dummy variable coefficients are in relation to East Asia & Pacific, which is therefore dropped. In Model (4), coefficients for North America and Sub-Saharan Africa are missing due to the limited number of observations (28 observations for Model 7).

temperature and altitude are high throughout, being significant at the 1% level in all models except Model 7, and even then significant at the 5% level. On the other hand, controlling for the country being landlocked presents as an insignificant coefficient which does not change the coefficient of temperature or affect the fit of the model much. We can interpret the results as follows, starting with Model 4 which includes controls for altitude and being landlocked: a 1°C increase in minimum temperature leads to around a 0.5-point increase in the Gini coefficient when it is measured from 1 to 100. When region dummies are added as per Model 5, the fit of the regression improves to around 60%, and the coefficient on temperature drops to 0.23. This means that even in otherwise more similar regions, we still find a significant effect of temperature on net inequality.

When using the log of Net Gini as dependent variable, the adjusted R-squared values for these results improves by around 0.05. However, we find it more intuitive to interpret the coefficients that are in Table 3, as they inform us of how the Gini would change on the scale from 1 to 100 instead of as percentage changes. As the improved fit of the log model is small, we relegate them to the Appendix (Table 4). Adding a quadratic term to temperature also only slightly improves the fit of the model, but we have included it in our main results of Table 3 as it can be interpreted in the same context as the linear term as just an interaction of temperature with itself. Figure 5 is a graphical representation of the base regression, showing how the linear and quadratic equations fit the observations. As also seen from this figure, there is no sizeable difference in the fit of the two.

The final three models (Models 6-8) show results with controls for our intermediate channels. Precipitation presents as an effect very close to zero and statistically insignificant, so the mechanism of temperature does not appear to be through countries having different precipitation levels. Model 7 adds preferences for redistribution, which presents as a negative but insignificant coefficient. This model has the limitation of only containing 28 observations, so the results in this column are tenuous. The coefficient on temperature is increased by adding these preferences, which we feel certain is just due to the data limitation as it would mean that preferences for redistribution have no effect on net Gini, and a correlation is only opened by controlling for them.

Model 8 adds market Gini, which presents as a significant and positive coefficient, which also decreases the coefficient of temperature from 0.24 to 0.19. This decrease is in line with what is expected based on Figure 2, as market Gini is an important determinant of net Gini. It is noteworthy that the decrease is only small, and a significant effect of temperature remains. This suggests that the channel whereby temperature affects net Gini is not primarily through market Gini, but must be through redistribution which we have failed to capture in Model 7 due to the data limitation. As net Gini is by definition the difference between market Gini and redistribution and closing off the effect of market Gini does not change the results much, it must be the case that the primary effect occurs through redistribution and taxes. In this manner, we are able to sidestep the data limitation issue for redistribution by inferring its effect through controlling for market Gini. This means that redistribution is a more important channel that temperature flows through than other government measures that change the market income distribution such as education policies.

Figure 6 shows that there is nearly no relationship between temperature and market Gini, which also points towards the effect of temperature on net Gini as being mediated through redistribution and taxes, instead of other government policies which affect the pre-tax and transfer wage distribution.

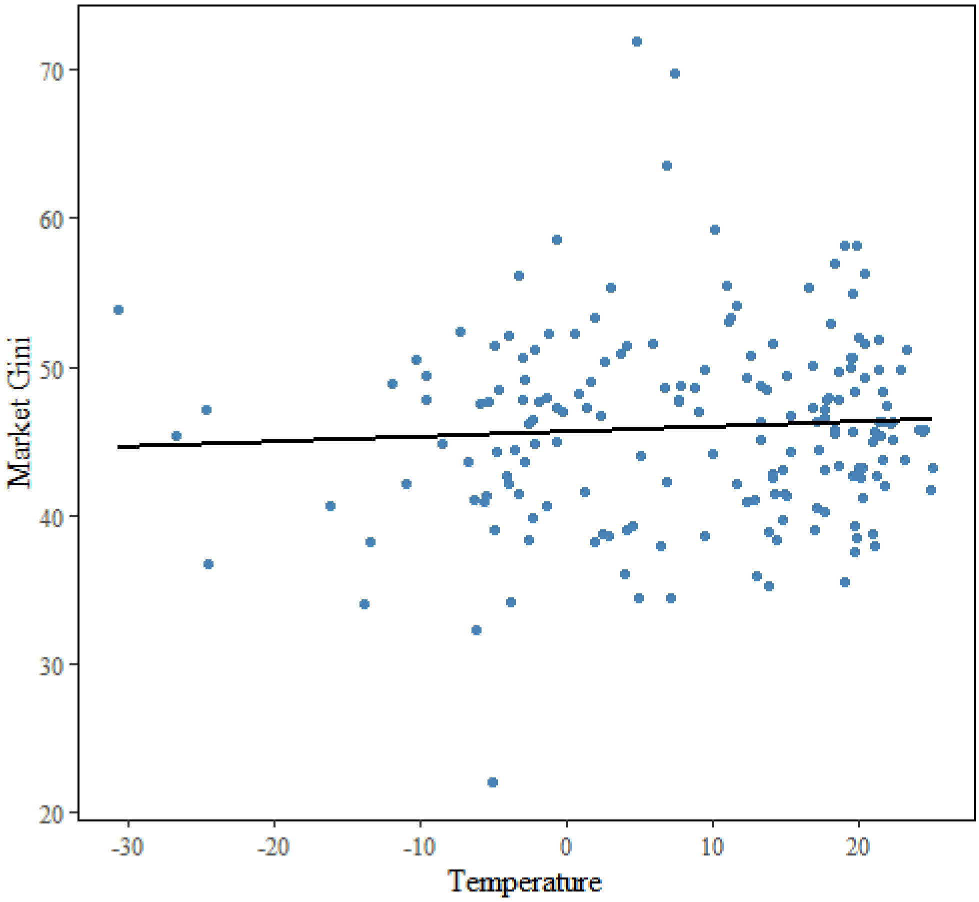

Figure 6 Fitted line of Market Gini on Temperature. The difference compared to Net Gini of Figure 5 is apparent: the income distribution before taxes and transfers is nearly unrelated to temperature.

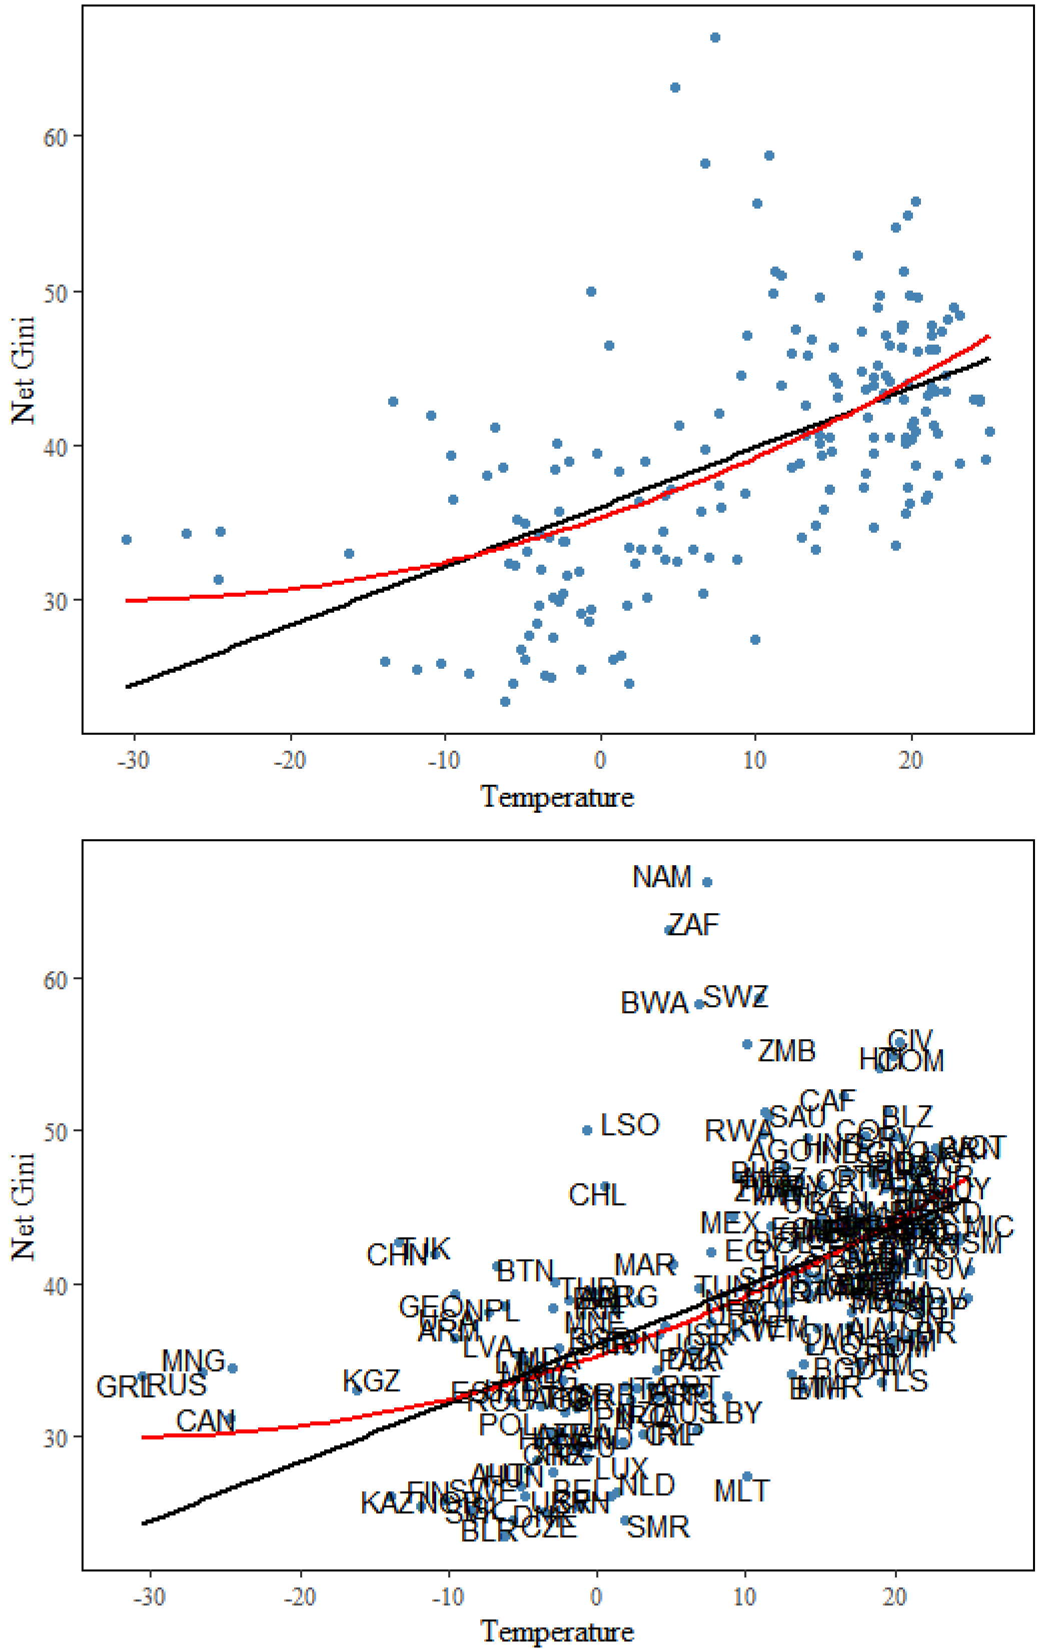

Figure 5 Fitted basic OLS regression lines with linear and quadratic forms, with the lower graph showing country abbreviations.

6. Discussion and Conclusion

We have introduced a deep determinant of income inequality, which to our knowledge has not been previously studied. We study the effect of temperature on net inequality in a cross-country setting, finding that a 1°C increase in temperature is associated with around a 0.5-point increase in net Gini when controlling for the country being landlocked and elevation. With continent region dummies, this effect of temperature becomes around 0.23, which shows that even in otherwise similar regions, there is still a significant effect. We find that a log-linear model fits the data slightly better, but have chosen to interpret the linear model as the improvement is not very large, with the log model available in the Appendix. These finding place only lenient demands the extent of the assumptions we must make, as there are relatively few things that can affect temperature and thus confound the effect.

It is remarkable that minimum temperature alone explains over 32% of net inequality across countries, and that a significant effect remains when controlling for the confounders elevation and being landlocked, and adding continent dummies. Further controls indicate that the effect of temperature happens primarily through taxes and redistribution, which we infer by controlling for market Gini and observing that the coefficient for temperature changes only slightly.

Taking the next steps is significantly more difficult, as the geography vs. institutions-debate indicates: more intricate knowledge as to how the effect takes place, and the relative importances of various causal channels is a much more complex task, and perhaps best left to machine learning algorithms that can test conditional independencies on large amounts of potential variables. However, understanding the large scale of the problem in this study may help humans better direct this search.

The social and policy relevance of the results would be to encourage more redistribution if the desire is to lower net inequality, to the extent that redistribution can be affected. However, as was discussed in the introduction, factors like fairness perceptions which underlie the chosen level of redistribution may be tied to the historical effects such as temperature which cannot be changed rapidly. This would mean that it is not possible to drastically change the extent of redistribution, as the preferences for redistribution themselves are different. Then, the policy relevance should be the understanding that different countries exist in different states of preference, whereby it is simply not desirable to copy redistribution levels from another country.

References

Afonso, A., Schuknecht, L., & Tanzi, V. (2010). Income distribution determinants and public spending efficiency. *The Journal of Economic Inequality*, *8*(3), 367–389. https://doi.org/10.1007/s10888-010-9138-z

Alesina, A., & Angeletos, G.-M. (2005). Fairness and Redistribution. *American Economic Review*, *95*(4), 960–980. https://doi.org/10.1257/0002828054825655

Alesina, A., & Giuliano, P. (2011). Chapter 4—Preferences for Redistribution. In J. Benhabib, A. Bisin, & M. O. Jackson (Eds.), *Handbook of Social Economics* (Vol. 1, pp. 93–131). North-Holland. https://doi.org/10.1016/B978-0-444-53187-2.00004-8

Ashraf, Q. H., & Michalopoulos, S. (2011). *The Climatic Origins of the Neolithic Revolution: Theory and Evidence* (SSRN Scholarly Paper ID 903847). Social Science Research Network. https://doi.org/10.2139/ssrn.903847

Barro, R. J. (2000). Inequality and Growth in a Panel of Countries. *Journal of Economic Growth*, *5*(1), 5–32. https://doi.org/10.1023/A:1009850119329

Bénabou, R., & Tirole, J. (2006). Belief in a Just World and Redistributive Politics*. *The Quarterly Journal of Economics*, *121*(2), 699–746. https://doi.org/10.1162/qjec.2006.121.2.699

Berg, A., Ostry, J. D., Tsangarides, C. G., & Yakhshilikov, Y. (2018). Redistribution, inequality, and growth: New evidence. *Journal of Economic Growth*, *23*(3), 259–305. https://doi.org/10.1007/s10887-017-9150-2

Cruces, G., Perez-Truglia, R., & Tetaz, M. (2013). Biased perceptions of income distribution and preferences for redistribution: Evidence from a survey experiment. *Journal of Public Economics*, *98*, 100–112. https://doi.org/10.1016/j.jpubeco.2012.10.009

Dell, M., Jones, B. F., & Olken, B. A. (2009). Temperature and Income: Reconciling New Cross-Sectional and Panel Estimates. *American Economic Review*, *99*(2), 198–204. https://doi.org/10.1257/aer.99.2.198

Dell, M., Jones, B. F., & Olken, B. A. (2014). What Do We Learn from the Weather? The New Climate-Economy Literature. *Journal of Economic Literature*, *52*(3), 740–798. https://doi.org/10.1257/jel.52.3.740

Diamond, J. (1997). *Guns, Germs, and Steel: The Fates of Human Societies*. W. W. Norton & Company.

Fick, S. E., & Hijmans, R. J. (2017). WorldClim 2: New 1-km spatial resolution climate surfaces for global land areas. *International Journal of Climatology*, *37*(12), 4302–4315. https://doi.org/10.1002/joc.5086

Hsiang, S. (2016). Climate Econometrics. *Annual Review of Resource Economics*, *8*(1), 43–75. https://doi.org/10.1146/annurev-resource-100815-095343

Isaksson, A.-S., & Lindskog, A. (2009). Preferences for redistribution—A country comparison of fairness judgements. *Journal of Economic Behavior & Organization*, *72*(3), 884–902. https://doi.org/10.1016/j.jebo.2009.08.006

Jarvis, A., Reuter, H. I., Nelson, A., Guevara, E., & others. (2008). Hole-filled SRTM for the globe Version 4. *Available from the CGIAR-CSI SRTM 90m Database (Http://Srtm. Csi. Cgiar. Org)*, *15*, 25–54.

Li, H., Squire, L., & Zou, H. (1998). Explaining International and Intertemporal Variations in Income Inequality. *The Economic Journal*, *108*(446), 26–43. https://doi.org/10.1111/1468-0297.00271

Masters, W. A., & McMillan, M. S. (2001). Climate and Scale in Economic Growth. *Journal of Economic Growth*, *6*(3), 167–186. https://doi.org/10.1023/A:1011398431524

Mayer, T., & Zignago, S. (2011). *Notes on CEPII's Distances Measures: The GeoDist Database* (SSRN Scholarly Paper ID 1994531). Social Science Research Network. https://doi.org/10.2139/ssrn.1994531

Nunn, N. (2008). *Slavery, Inequality, and Economic Development in the Americas: An Examination of the Engerman-Sokoloff Hypothesis*. 41.

Nunn, N. (2014). Chapter 7—Historical Development. In P. Aghion & S. N. Durlauf (Eds.), *Handbook of Economic Growth* (Vol. 2, pp. 347–402). Elsevier. https://doi.org/10.1016/B978-0-444-53538-2.00007-1

Olsson, O., & Hibbs, D. A. (2005). Biogeography and long-run economic development. *European Economic Review*, *49*(4), 909–938. https://doi.org/10.1016/j.euroecorev.2003.08.010

Palma, J. G. (2011). Homogeneous Middles vs. Heterogeneous Tails, and the End of the 'Inverted-U': It's All About the Share of the Rich. *Development and Change*, *42*(1), 87–153. https://doi.org/10.1111/j.1467-7660.2011.01694.x

Pearl, J. (2009). *Causality*. Cambridge University Press.

Putterman, L., & Weil, D. N. (2010). Post-1500 Population Flows and The Long-Run Determinants of Economic Growth and Inequality*. *The Quarterly Journal of Economics*, *125*(4), 1627–1682. https://doi.org/10.1162/qjec.2010.125.4.1627

Rodrik, D., Subramanian, A., & Trebbi, F. (2004). Institutions Rule: The Primacy of Institutions Over Geography and Integration in Economic Development. *Journal of Economic Growth*, *9*(2), 131–165. https://doi.org/10.1023/B:JOEG.0000031425.72248.85

Scheve, K., & Stasavage, D. (2005). Religion and Preferences for Social Insurance. *London School of Economics. Manuscript*.

Sokoloff, K. L., & Engerman, S. L. (2000). Institutions, Factor Endowments, and Paths of Development in the New World. *Journal of Economic Perspectives*, *14*(3), 217–232. https://doi.org/10.1257/jep.14.3.217

Solt, F. (2020). Measuring Income Inequality Across Countries and Over Time: The Standardized World Income Inequality Database. *Social Science Quarterly*, *101*(3), 1183–1199. https://doi.org/10.1111/ssqu.12795

Spolaore, E., & Wacziarg, R. (2013). How Deep Are the Roots of Economic Development? *Journal of Economic Literature*, *51*(2), 325–369. https://doi.org/10.1257/jel.51.2.325

Starmans, C., Sheskin, M., & Bloom, P. (2017). Why people prefer unequal societies. *Nature Human Behaviour*, *1*(4), 1–7. https://doi.org/10.1038/s41562-017-0082

Stiglitz, J. E. (1996). SOME LESSONS FROM THE EAST ASIAN MIRACLE. *The World Bank Research Observer*, *11*(2), 151–177. https://doi.org/10.1093/wbro/11.2.151

Summerhill, W. (2010). *Colonial Institutions, Slavery, Inequality, and Development: Evidence from São Paulo, Brazil* (SSRN Scholarly Paper ID 1591449). Social Science Research Network. https://doi.org/10.2139/ssrn.1591449

Appendix

Table 4: Log Model Tests

	Dependent variable: log(Net Gini)					
	(1)	(2)	(3)	(4)	(5)	(6)
Temperature	0.009***	0.013***	0.006***	0.007***	0.010**	0.005***
	(0.001)	(0.001)	(0.001)	(0.002)	(0.005)	(0.001)
Temperature^2	0.0002***	0.0002***	0.0001*	0.0001*	0.0003	0.0001**
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0003)	(0.0001)
Altitude		0.0002***	0.0001***	0.0001***	0.0002**	0.0001***
		(0.00003)	(0.00002)	(0.00002)	(0.0001)	(0.00002)
Landlocked		-0.006	-0.024	-0.025	-0.143**	-0.008
		(0.036)	(0.032)	(0.033)	(0.063)	(0.023)
Europe & Central Asia			-0.123***	-0.123***	-0.119***	-0.043
			(0.042)	(0.042)	(0.094)	(0.040)
Latin America & Caribbean			0.113***	0.113***	0.112***	0.120
			(0.026)	(0.026)	(0.129)	(0.020)
Middle East & North Africa			0.038	0.038	0.036	0.035
			(0.039)	(0.039)	(0.075)	(0.029)
North America			-0.014	-0.012		-0.120*
			(0.056)	(0.056)		(0.064)
South Asia			-0.006	-0.007	0.143*	0.018
			(0.064)	(0.065)	(0.083)	(0.040)
Sub-Saharan Africa			0.123***	0.120***		0.051**
			(0.032)	(0.032)		(0.021)
Precipitation				-0.00002	-0.0003	-0.00003
				(0.0001)	(0.0003)	(0.00002)
Redistribution preferences					-0.068	
					(0.104)	
Market Gini						0.015***
						(0.002)
Constant	3.539***	3.407***	3.494***	3.494***	3.563***	2.868***
	(0.021)	(0.022)	(0.043)	(0.043)	(0.244)	(0.077)
N	187	165	165	165	28	165
R^2	0.368	0.579	0.665	0.666	0.728	0.836
Adjusted R^2	0.361	0.568	0.644	0.642	0.568	0.823
Residual Std. Error	0.165 (df = 184)	0.137 (df = 160)	0.125 (df = 154)	0.125 (df = 153)	0.136 (df = 17)	0.088 (df = 152)
F Statistic	53.536*** (df = 2; 184)	54.999*** (df = 4; 160)	30.615*** (df = 10; 154)	27.687*** (df = 11; 153)	4.556*** (df = 10; 17)	64.605*** (df = 12; 152)

Notes: ***Significant at the 1 percent level.; **Significant at the 5 percent level.; *Significant at the 10 percent level. Region dummy variable coefficients are in relation to East Asia & Pacific, which is therefore dropped. In Model (4), coefficients for North America and Sub-Saharan Africa are missing due to the limited number of observations (28 observations for Model 7).

Table 5: Lists of Countries Used in the Main Regressions

165 Countries	Afghanistan, Angola, Albania, Andorra, United Arab Emirates, Argentina, Armenia, Antigua & Barbuda, Australia, Austria, Azerbaijan, Burundi, Belgium, Benin, Burkina Faso, Bangladesh, Bulgaria, Bahrain, Bahamas, Bosnia & Herzegovina, Belarus, Belize, Bolivia, Brazil, Barbados, Brunei, Bhutan, Botswana, Central African Republic, Canada, Switzerland, China, Cote d'Ivoire, Cameroon, Congo - Brazzaville, Colombia, Comoros, Cape Verde, Costa Rica, Cyprus, Czechia, Germany, Djibouti, Dominica, Denmark, Dominican Republic, Algeria, Ecuador, Egypt, Spain, Estonia, Ethiopia, Finland, France, Gabon, United Kingdom, Georgia, Ghana, Gambia, Guinea-Bissau, Equatorial Guinea, Greece, Grenada, Greenland, Guatemala, Guyana, Hong Kong SAR China, Honduras, Croatia, Haiti, Hungary, Indonesia, India, Ireland, Iran, Iraq, Iceland, Israel, Italy, Jamaica, Jordan, Japan, Kazakhstan, Kenya, Kyrgyzstan, Cambodia, St. Kitts & Nevis, Kuwait, Laos, Lebanon, Liberia, Libya, Sri Lanka, Lesotho, Lithuania, Luxembourg, Latvia, Morocco, Moldova, Madagascar, Mexico, North Macedonia, Mali, Myanmar (Burma), Mongolia, Mozambique, Mauritania, Mauritius, Malawi, Malaysia, Namibia, Niger, Nigeria, Nicaragua, Netherlands, Norway, Nepal, Oman, Pakistan, Panama, Peru, Philippines, Papua New Guinea, Poland, Paraguay, Qatar, Rwanda, Saudi Arabia, Sudan, Senegal, Solomon Islands, Sierra Leone, El Salvador, San Marino, Somalia, Sao Tome & Principe, Suriname, Slovakia, Slovenia, Sweden, Eswatini, Syria, Chad, Togo, Thailand, Tajikistan, Turkmenistan, Tonga, Trinidad & Tobago, Tunisia, Turkey, Tuvalu, Tanzania, Uganda, Ukraine, Uruguay, Uzbekistan, St. Vincent & Grenadines, Venezuela, Vietnam, Vanuatu, Samoa, Yemen, Zambia, Zimbabwe
28 Countries (regression including redistribution preferences)	Australia, Belgium, Switzerland, Czechia, Germany, Denmark, Spain, Finland, France, United Kingdom, Georgia, Croatia, Hungary, India, Iceland, Israel, Japan, Lithuania, Latvia, Norway, Philippines, Suriname, Slovakia, Slovenia, Sweden, Thailand, Turkey, Venezuela