2023-01
Working paper. Economics
ISSN 2340-5031

HETEROGENEOUS PREDICTIVE ASSOCIATION OF CO_2 WITH GLOBAL WARMING

LIANG CHEN, JUAN J. DOLADO, JESÚS GONZALO and ANDREY RAMOS

uc3m | Universidad Carlos III de Madrid

Serie disponible en http://hdl.handle.net/10016/11

Web: http://economia.uc3m.es/

Correo electrónico: departamento.economia@eco.uc3m.es

Heterogeneous Predictive Association of CO_2 with Global Warming*

Liang Chen[1], Juan J. Dolado[2], Jesús Gonzalo[3], and Andrey Ramos[3]

[1] *Peking University HSBC Business School, chenliang@phbs.pku.edu.cn*
[2] *Department of Economics, Universidad Carlos III de Madrid, dolado@eco.uc3m.es*
[3] *Department of Economics, Universidad Carlos III de Madrid, jgonzalo@est-eco.uc3m.es*
[3] *Department of Economics, Universidad Carlos III de Madrid, anramosr@eco.uc3m.es*

January 25, 2023

Abstract

Global warming is a non-uniform process across space and time. This opens the door to a heterogeneous relationship between CO_2 and temperature that needs to be analyzed going beyond the standard analysis based on mean temperature found in the literature. We revisit this topic through the lenses of a new class of factor models for high-dimensional panel data, labeled Quantile Factor Models (QFM). This technique extracts quantile-dependent factors from the distributions of temperature across a wide range of stable weather stations in the Northern and Southern Hemispheres over 1959-2018. In particular, we test whether the (detrended) growth rate of CO_2 concentrations help predict the underlying factors of the different quantiles of the distribution of (detrended) temperature in the time dimension. We document that predictive association is greater at the lower and medium quantiles than at the upper quantiles and provide some conjectures about what could be behind this non-uniformity. These findings complement recent results in the literature documenting steeper trends in lower temperature levels than in other parts of the spatial distribution.

Keywords: Global warming, CO2 concentrations, Quantile factor models, Predictive association.
JEL codes: C31, C33, Q54.

*Prepared for the *Economica 100 challenge*. We are indebted to an Editor and three referees for very helpful comments, as well as to participants at the XII Time Series Workshop in Zaragoza for useful suggestions. Financial support from the National Natural Science Foundation of China (Grant No.71703089), the Spanish Ministerio de Economía y Competitividad (grants PID2019-104960GB-I00, PID2020-118659RB-I00 and TED2021-129784B-I00), and MadEco-CM (grant S205/HUM-3444) is gratefully acknowledged. The usual disclaimer applies.

1 Introduction

As stressed by world leaders at the 2021 UN Climate Change Conference of the Parties (COP26), one of the most pressing issues in the international policy agenda is the fight against the rise of global surface temperatures. Since this phenomenon is mainly due to the increasing concentrations of greenhouse gases in the atmosphere, a proper design of climate policy requires a deep understanding of the relationship of global warming (GW, henceforth) with carbon dioxide concentrations (CO_2).[1]

From an econometrics viewpoint, several studies have used time-series techniques to explore this topic empirically (see e.g. Kaufmann et al. (2006), Stips et al. (2016), Castle and Hendry (2020), Montamat and Stock (2020), Phillips et al. (2020), Pretis (2020), or Chen et al. (2022), among many others). The standard practice is to use a time series of average global temperature across a large number of stations to quantify the so-called Equilibrium Climate Sensitivity (ECS), defined as the temperature response to a doubling in the CO_2 concentrations, and Transient Climate Response (TCR), which measures the strength and speed at which climate responds to greenhouse gas forcing. The reliability of this evidence depends on the statistical properties assumed for the time series of interest, such as their the order of integration or the type of trends they present.

This standard analysis of average temperature has to be complemented by a broader one that takes into account the well-known fact that GW is a spatially and temporally non-uniform process (Chapman et al., 2013; Shindell, 2014; Ji et al., 2014; Previdi et al., 2021; Gadea and Gonzalo, 2023). Our contribution to this literature goes in this direction by proposing a novel econometric methodology aimed at establishing predictive-association between temperature and CO_2 concentrations allowing for heterogeneity in this relationship. It relies on Quantile Factor Models (QFM) (Chen et al., 2021), whose use in this context can be motivated as follows.

Let a a panel of temperature $\{X_{it}\}$ be available for $i = 1, ..., N$ stations and $t = 1, ..., T$ periods together with time observations on CO_2, denoted as $\{Z_t\}$, $t = 1, ..., T$, which are uniform across all stations. Different approaches are available to analyze the association between these two variables. For example, researchers can adopt a conditional regression approach for each individual station, $\mathbb{E}[X_{it}|Z_t] = \gamma_{0i} + \gamma_{1i} Z_t$, $i = 1, ..., N$, and then compute the distribution of the parameter of interest γ_1 using station-level estimates of $\hat{\gamma_{1i}}$. However, when N is large and T

[1] When sunlight reaches Earth, its surface absorbs some of the light's energy and re-radiates it as infrared waves. These waves travel up into the atmosphere and will escape back into space if unimpeded. For example, while oxygen and nitrogen do not interfere with infrared waves in the atmosphere because molecules are picky about the range of wavelengths they interact with, CO_2 and other greenhouse gases absorb energy at a variety of wavelengths whose ranges do overlap with that of infrared energy. As CO_2 soaks up this infrared energy, it vibrates and re-emits the infrared energy back in all directions. About half of that energy goes out into space, while the other half returns to Earth as heat, contributing to GW through the so-called "greenhouse effect", first discovered by Fourier (1824), experimentally verified by Foote (1856) and Tyndall (1863), and quantified by Arrhenius (1896).

is smaller, not only this procedure can be computationally burdensome but also the individual estimates from each separate regression would lack precision. As a result, a standard practice in this literature is to aggregate temperatures over the cross-sectional dimension and estimate a single conditional regression of the form $\mathbb{E}\left[\bar{X}_t|Z_t\right] = \bar{\gamma_0} + \bar{\gamma_1}Z_t$, where $\bar{X}_t$ is the series of mean temperature, and $\bar{\gamma_1}$ is the mean of γ_1. Likewise, by assuming a single common factor structure, $X_{it} = \lambda_i f_t + \epsilon_{it}$, an alternative procedure would be to extract the common factor $\hat{f}_t$ through Principal Components Analysis (PCA), and estimate a single conditional regression of the form $\mathbb{E}\left[\hat{f}_t|Z_t\right] = \delta_0 + \delta_1 Z_t$. Under the standard factor model conditions (see Bai (2003)), $\bar{\gamma_1}$ and δ_1 are expected to be proportional.[2]

However, a serious limitation of these approaches is that they do not allow for the presence of heterogeneous patterns in the association of CO_2 with temperature. Heterogeneity can be accounted for by adopting quantile approaches. For example, consider a conditional quantile regression (QR) model for each individual station in the form of $Q_\tau[X_{it}|Z_t] = \beta_{0i}(\tau) + \beta_{1i}(\tau)Z_t$, where $0 < \tau < 1$ denotes the quantile level, and $\beta_{1i}(\tau)$ is the object of interest. In this scenario, two different sources of heterogeneity can be considered. First, for a given station i, $\beta_{1i}(\tau)$ can vary across τ, capturing heterogeneity along the temperature distribution, e.g. by finding that CO_2 affects more strongly the lowest than the highest temperatures in a given location. Second, for a given τ, $\beta_{1i}(\tau)$ can vary across i, implying the existence of spatial heterogeneity e.g. when CO_2 affects more the lowest temperature in the Antarctica than in the Sahara.

Considering the difficulty embodied in incorporating jointly both types of heterogeneity, in this paper we explore methods to characterize it along the temperature distribution. A natural approach in that direction would be to estimate a standard QR for the mean temperature, $Q_\tau[\bar{X}_t|Z_t] = \bar{\beta}_0(\tau) + \bar{\beta}_1(\tau)Z_t$, $0 < \tau < 1$ which, however, still disregards all the information on the distribution of temperatures across different stations.[3] Alternatively, one could run QR for each station, and then average the estimates $\hat{\beta}_{1i}(\tau)$ across units for each τ. Yet, this would be a doubtful statistical procedure because the average of quantiles differs from the quantile of averages.

For this reason, our proposal relies on Chen et al. (2021, CDG hereafter) where a quantile factor structure is assumed for a panel data on temperatures. Accordingly, $X_{it} = \lambda_i'(\tau)f_t(\tau) + u_{it}(\tau)$, where $f_t(\tau)$ and $\lambda_i(\tau)$ are a $r(\tau) \times 1$ vectors of factors and loadings, respectively, which may differ at each τ. Once consistent estimates of the quantile-dependent objects are obtained, a natural approach is to relate the estimated common factors at each quantile $\hat{f}_t(\tau)$ with Z_t through time-series methods. Note that the quantile-dependent common factors in this setup are interpreted as aggregators for the objects of interest, in the same way as the PCA factors

[2]It is possible to show that, under certain assumptions on the individual loadings and the idiosyncratic error in the single common factor structure, the aggregated common factor over the cross-sectional dimension is proportional to the mean. Therefore, PCA common factors are useful aggregators for the mean.

[3]Throughout the paper we use $\mathsf{Q}_X[\tau|Z]$ to denote the conditional quantile of X given Z.

are aggregators for the mean.

In line with the generalization of linear regression to QR models, QFM can be understood as an extension of approximate factor models (AFM) to allow for hidden factors shifting specific characteristics (moments or quantiles) of the distribution of temperature.[4] As a simple illustration of the advantages of QFM over AFM, consider the factor structure in a standard *location-scale shift* model with the following Data Generating Process (DGP): $X_{it} = \alpha_i f_{1t} + \eta_i f_{2t}\epsilon_{it}$, with $f_{1t} \neq f_{2t}$ (both are scalars), $\eta_i, f_{2t} > 0$ and $\mathbb{E}(\epsilon_{it}) = 0$. The first factor ($f_{1t}$) shifts location, whereas the second (f_{2t}) shifts the scale and therefore governs the volatility of shocks to X_{it}.[5] Such a DGP can be rewritten in QR format as $X_{it} = \lambda_i'(\tau) f_t + u_{it}(\tau)$, with $0 < \tau < 1$, $\lambda_i(\tau) = [\alpha_i, \eta_i \mathsf{Q}_\epsilon(\tau)]'$, where $\mathsf{Q}_\epsilon(\tau)$ represents the quantile function of ϵ_{it}, $f_t = [f_{1t}, f_{2t}]'$, $u_{it}(\tau) = \eta_i f_{2t}[\epsilon_{it} - \mathsf{Q}_\epsilon(\tau)]$, and the conditional quantile $\mathsf{Q}_{u_{it}(\tau)}[\tau|f_t] = 0$.It can be checked that PCA will only extract the location-shifting factor f_{1t} in this model, but it will fail in capturing the scale-shifting factor f_{2t}. By contrast, the estimation of QFM by the so-called *Quantile Factor Analysis* (QFA) will be able to retrieve the space spanned by both factors in this DGP. [6]

As explained further below, the QFA estimation procedure relies on the minimization of the standard *check* function in QR (instead of the quadratic loss function used in AFM) to estimate jointly the common factors and their loadings at a given quantile, once the number of factors has been selected by a consistent criterion. Lastly, it is noteworthy that, given that QFA captures all quantile-shifting factors (including those affecting the means of observed variables), it provides a natural way to differentiate AFM from QFM. This is specially relevant in the presence of outliers where QFA will render valid estimation and inference while AFM may not work well[7] Indeed, since outliers are present in the panel data of temperature, we will illustrate the advantages of using QFA in such a case through several Monte-Carlo simulations.

In the empirical analysis, we make use of a balanced panel of 441 station-level annual mean temperature series over the period 1959-2018. A key requirement for extracting the QFA (and PCA) factors is that the individual time-series processes do not have stochastic or deterministic trends. Therefore, prior to implementing the QFA, we establish the statistical properties of these series and apply the corresponding filtering to achieve that condition. Note that the identification of the type of process followed by the temperature series has been subject to an intense debate in the climate econometrics literature. On the one hand, authors as Kaufmann et al. (2006), Chang et al. (2020), Phillips et al. (2020), or Pretis (2020), among others, argue

[4] Ando and Bai (2020) use a similar setup with an unobservable factor structure which is also allowed to be quantile dependent; yet, their assumptions are more restrictive since all the moments of the idiosyncratic errors are required to exist.

[5] Note that the simplifying assumption of a known number of factors in this specific example is later relaxed.

[6] Note that, since f_{1t} can be consistently estimated by PCA in this specific DGP, it is also feasible to estimate f_{2t} by applying PCA to the squared residuals obtained from subtracting the factor structure at the mean from the original variables. However, in practice, the DGP is unknown and therefore QFA is needed

[7] The insight for this relative performance is similar to the one underlying the use of robust least median regression when outliers abound, as in Huber (1981)

that temperature has stochastic trends, as it is also the case for (CO_2), and therefore defend the use of cointegration methods if appropriate. On the other hand, authors as Gao and Hawthorne (2006), Gay-Garcia et al. (2009), or Estrada et al. (2013), claim that global temperature follows a trend-stationary process with nonlinear deterministic trends. Our findings are more in line with the latter view: we cannot reject that the temperature series are trend-stationary, though trends appear to be linear.[8] Hence, we linearly detrend X_{it} before QFA is implemented.

Once the quantile-dependent common factors have been obtained, the next step in our proposed methodology is to determine the existence of predictive-association between them and CO_2 concentrations. In line with most of the literature, the statistical properties of the latter variable points to the existence of a unit root in levels, while its first difference is trend-stationary, again with a linear trend. Therefore, we consider linearly-detrended changes in atmospheric CO_2 concentrations as a predictor for the quantile-dependent common factor series of (detrended) temperature. The existence of predictive-association is determined through an F-test for the joint significance of the coefficients associated to the lags of the CO_2 concentrations-related variable.

Two important features of our proposed methodology need to be clarified from the outset. First, our proposed procedure is akin to a Granger-causality analysis for a given information set in a restrictive sense, namely, by assuming CO_2 concentrations (together with past temperatures) are the whole information set available by the econometrician. Since this is not realistic, we only claim the existence of predictive-power or predictive-association. Second, the object of interest in our analysis differs from the complementary studies relating temperature and CO_2 concentrations. Here, the QFA allows us to extract the common factors that drive variations of temperature around a linear trend in all the stations. For instance, at a low quantile of reference, we are extracting the common factors that drive large-negative fluctuations in detrended local temperatures in all available units, to later examine if such fluctuations can be predicted by past changes in CO_2 concentrations around its trend. This is a different association indicator from the TCS or the TCR analyzed in the standard literature which seek to establish a causal relationship through a more structural model.

In summary, starting from the well-known fact that GW is a non-uniform process (spatially and temporarily), this paper provides a novel quantitative methodology which helps analyze the heterogeneous predictive association between CO_2 and the GW process. We find that this association, at the temporal level station by station, is statistically significant at the lower part of the temperature distribution and non-significant at the upper one. This feature can have more serious consequences than an increase in the middle part of the distribution. In this sense, our results complement the available ones in the literature on climate sensitivity (Sherwood et al., 2020) that mostly focus on the mean temperature. They also point out that future

[8] A complete theoretical as well as empirical trend analysis can be found in Gadea and Gonzalo (2020).

climate agreements should go beyond the mean-temperature target and instead consider the whole temperature distribution and a CO_2 concentrations objective.

The rest of the paper is organized as follows. Section 2 defines QFM. In Section 3, we introduce the QFA estimator and its computational algorithm, a consistent selection criterion to choose the number of factors at each quantile, and finally run a Monte Carlo simulation to highlight the advantages of using QFA instead of PCA in finite samples with big outliers. Section 4 considers the empirical application about the predictive association between CO_2 concentrations and GW using a large panel dataset on the annual distributions of temperatures. Section 5 discusses our main findings. Finally, Section 6 concludes. An Appendix gathers supplementary material related to further robustness results and some auxiliary procedures referred to in the main text.

2 Quantile Factor Model

To motivate our empirical analysis, this section reviews the basic concepts and tools underlying CDG's (2021) QFM approach.

Let $\{X_{it}\}$ be a panel of N observed variables (units), each with T observations. Then, X_{it}, with $i = 1, 2, .., N$ and $t = 1, 2, ..., T$, has the following QFM structure at some $\tau \in (0, 1)$:

$$\mathsf{Q}_{X_{it}}[\tau|f_t(\tau)] = \lambda_i'(\tau)f_t(\tau),$$

where the common factors $f_t(\tau)$ are gathered in a $r(\tau)\times 1$ vector of unobserved random variables, $\lambda_i(\tau)$ is a $r(\tau) \times 1$ vector of non-random factor loadings with $r(\tau) \ll N$. Note that in the QFM defined above, the factors, the loadings, and the number of factors are all allowed to be quantile-dependent.

Alternatively, the above equation implies that

$$X_{it} = \lambda_i'(\tau)f_t(\tau) + u_{it}(\tau), \tag{1}$$

where the quantile-dependent idiosyncratic error $u_{it}(\tau)$ is assumed to satisfy the quantile restriction $P[u_{it}(\tau) \leq 0|f_t(\tau)] = \tau$.

As mentioned in the Introduction, location-scale shift models provide nice illustrations of potential DGPs with the above QFM representation. In particular, recall the example given above, i.e. $X_{it} = \alpha_i'f_{1t} + (\eta_i'f_{2t})\epsilon_{it}$, where $\{\epsilon_{it}\}$ are zero-mean i.i.d errors independent of $\{f_{1t}\}$ and $\{f_{2t}\}$, with cumulative distribution function (CDF) F_ϵ such that the median of ϵ_{it} is 0, i.e., $\mathsf{Q}_\epsilon(0.5) = 0$, $\alpha_i, f_{1t} \in \mathbb{R}^{r_1}$, η_i, $f_{2t} \in \mathbb{R}^{r_2}$, and $\eta_i'f_{2t} > 0$. Then, when f_{1t} and f_{2t} do not share common elements, this model has a QFM representation as in (1) with $\lambda_i(\tau) = [\alpha_i', \eta_i'\mathsf{Q}_\epsilon(\tau)]'$,

$f_t(\tau) = [f_{1t}', f_{2t}']$ for $\tau \neq 0.5$, and $\lambda_i(\tau) = \alpha_i$, $f_t(\tau) = f_{1t}$ for $\tau = 0.5$. Note that, for this DGP, the loadings are quantile-dependent objects while the factors are not. An alternative example where factors do depend on quantiles is provided by a similar DGP where now different (positive) factors affect the first three moments of the data, i.e. $X_{it} = \alpha_i f_{1t} + f_{2t}\epsilon_{it} + f_{3t}\epsilon_{it}^3$, where ϵ_{it} follows a standard normal random variable with CDF $\Phi(\cdot)$. Then X_{it} has an equivalent representation in form of (1), with $\lambda_i(\tau) = [\alpha_i, \Phi^{-1}(\tau), c_i\Phi^{-1}(\tau)^3]'$, $f_t(\tau) = (f_{1t}, f_{2t}, f_{3t})'$ for $\tau \neq 0.5$, and $\lambda_i(\tau) = \alpha_i$, $f_t(\tau) = f_{1t}$ for $\tau = 0.5$. In particular, since the mapping $\tau \mapsto \Phi^{-1}(\tau)^3$ is strictly increasing, then there exists a QFM representation as in (1) with $\lambda_i(\tau) = [\alpha_i, \Phi^{-1}(\tau)]'$ and $f_t(\tau) = [f_{1t}, f_{2t} + f_{3t}\Phi^{-1}(\tau)^2]'$ for $\tau \neq 0.5$, so that the second factor in $f_t(\tau)$ is quantile dependent even for $\tau \neq 0.5$.

Finally, recall that applying PCA to the data in the two previous DGPs will fail to capture the extra factors shifting quantiles, other than the means. Hence the need to use QFA to estimate all quantile-dependent objects in the QFM.

3 QFA Estimators

To simplify the notations, we suppress hereafter the dependence of $f_t(\tau), \lambda_i(\tau), r(\tau)$ and $u_{it}(\tau)$ on τ, so that the QFM in (1) is rewritten as:

$$X_{it} = \lambda_i' f_t + u_{it}, \quad P[u_{it} \leq 0 | f_t] = \tau, \tag{2}$$

where $\lambda_i, f_t \in \mathbb{R}^r$. Let $\{f_{0t}\}$ and $\{\lambda_{0i}\}$ be the true values of $\{f_t\}$ and $\{\lambda_i\}$, respectively. A fixed-effects approach is taken by treating $\{\lambda_{0i}\}$ and $\{f_{0t}\}$ as parameters to be estimated, so that the asymptotic analysis is conditional on $\{f_{0t}\}$. In subsection 3.1, we consider the estimation of $\{\lambda_{0i}\}$ and $\{f_{0t}\}$ while r is assumed to be known, while the estimation of r at each quantile is discussed later in subsection 3.3.

3.1 Estimating Quantile Factors and Loadings

It is well known in the literature on factor models that $\{\lambda_{0i}\}$ and $\{f_{0t}\}$ cannot be separately identified without imposing normalizations (see Bai and Ng (2002)). Without loss of generality, the following normalizations are imposed:

$$\frac{1}{T}\sum_{t=1}^{T} f_t f_t' = \mathbb{I}_r, \quad \frac{1}{N}\sum_{i=1}^{N} \lambda_i \lambda_i' \text{ is diagonal with non-increasing diagonal elements.} \tag{3}$$

Let $M = (N+T)r$, $\theta = (\lambda_1', \ldots, \lambda_N', f_1', \ldots, f_T')'$, and $\theta_0 = (\lambda_{01}', \ldots, \lambda_{0N}', f_{01}', \ldots, f_{0T}')'$ denotes the vector of true parameters, where the dependence of θ and θ_0 on M is also suppressed

to save notation. Let $\mathcal{A}, \mathcal{F} \subset \mathbb{R}^r$ and define:

$$\Theta^r = \left\{\theta \in \mathbb{R}^M : \lambda_i \in \mathcal{A}, f_t \in \mathcal{F} \text{ for all } i, t, \{\lambda_i\} \text{ and } \{f_t\} \text{ satisfy the normalizations in (3)}\right\}.$$

Further, define:

$$\mathbb{M}_{NT}(\theta) = \frac{1}{NT}\sum_{i=1}^{N}\sum_{t=1}^{T}\rho_\tau(X_{it} - \lambda_i' f_t),$$

where $\rho_\tau(u) = (\tau - \mathbf{1}\{u \leq 0\})u$ is the check function. The QFA estimator of θ_0 is defined as:

$$\hat{\theta} = (\hat{\lambda}_1', \ldots, \hat{\lambda}_N', \hat{f}_1', \ldots, \hat{f}_T')' = \arg\min_{\theta \in \Theta^r} \mathbb{M}_{NT}(\theta).$$

This estimator extends the PCA estimator studied by Bai and Ng (2002) and Bai (2003) in the same way as QR is related to standard least-squares regressions. However, unlike these PCA estimators, $\hat{\theta}$ does not yield an analytical closed form. Thus the need of CDG's (2021) computational algorithm, labeled *Iterative Quantile Regression* (IQR), that can effectively find the stationary points of the object function and can be described as follows.

Let $\Lambda = (\lambda_1, \ldots, \lambda_N)'$, $F = (f_1, \ldots, f_T)'$, and define the following averages:

$$\mathbb{M}_{i,T}(\lambda, F) = \frac{1}{T}\sum_{t=1}^{T}\rho_\tau(X_{it} - \lambda' f_t) \quad \text{and} \quad \mathbb{M}_{t,N}(\Lambda, f) = \frac{1}{N}\sum_{i=1}^{N}\rho_\tau(X_{it} - \lambda_i' f).$$

Note that we have $\mathbb{M}_{NT}(\theta) = N^{-1}\sum_{i=1}^{N}\mathbb{M}_{i,T}(\lambda_i, F) = T^{-1}\sum_{t=1}^{T}\mathbb{M}_{t,N}(\Lambda, f_t)$. The main difficulty in finding the global minimum of $\mathbb{M}_{NT}$ is that this object function is not convex in θ. However, for given F, $\mathbb{M}_{i,T}(\lambda, F)$ happens to be convex in λ for each i and likewise, for given Λ, $\mathbb{M}_{t,N}(\Lambda, f)$ is also convex in f for each t. Thus, both optimization problems can be efficiently solved by various linear programming methods (see Chapter 6 of Koenker (2005)). Based on this observation, the following iterative procedure is proposed:

Iterative quantile regression (IQR):
Step 1: Choose random starting parameters: $F^{(0)}$.
Step 2: Given $F^{(l-1)}$, solve $\lambda_i^{(l-1)} = \arg\min_\lambda \mathbb{M}_{i,T}(\lambda, F^{(l-1)})$ for $i = 1, \ldots, N$; given $\Lambda^{(l-1)}$, solve $f_t^{(l)} = \arg\min_f \mathbb{M}_{t,N}(\Lambda^{(l-1)}, f)$ for $t = 1, \ldots, T$.
Step 3: For $l = 1, \ldots, L$, iterate the second step until $\mathbb{M}_{NT}(\theta^{(L)})$ is close to $\mathbb{M}_{NT}(\theta^{(L-1)})$, where $\theta^{(l)} = (\text{vech}(\Lambda^{(l)})', \text{vech}(F^{(l)})')'$.
Step 4: Normalize $\Lambda^{(L)}$ and $F^{(L)}$ so that they satisfy the normalizations in (3).

In the general case where $r \geq 1$, replacing the check function in the IQR algorithm by the least-squares loss function and normalizing $F^{(l-1)}, \Lambda^{(l-1)}$ to satisfy (3) at step 2, IQR is equivalent to the method of *orthogonal iterations* proposed by Golub and Van Loan (2013) to compute the eigenvectors associated with the r largest eigenvalues of XX'.

The asymptotic properties of the QFA estimators are not presented here for brevity, but can be reviewed directly in CDG (2021). In any case, it should be noted that they achieve asymptotic normality with the same convergence rates as PCA and, foremost, that these properties hold even when the distribution of the idiosyncratic errors has no moments.

3.2 Selecting the Number of Factors at Quantiles

To allow for an unknown number of quantile-dependent factors CDG (2021) propose a rank-minimization criterion to select the correct number at each τ with probability approaching one. Suppressing once again the dependence of $r(\tau)$ on τ to ease notation, the criterion works as follows.

Let k be a positive integer larger than r, and $\mathcal{A}^k$ and $\mathcal{F}^k$ be compact subsets of $\mathbb{R}^k$. In particular, let us assume that $[\lambda_{0i}' \quad \mathbf{0}_{1\times(k-r)}]' \in \mathcal{A}^k$ for all i.

Let $\lambda_i^k, f_t^k \in \mathbb{R}^k$ for all i, t and write $\theta^k = (\lambda_1^{k\prime}, \ldots, \lambda_N^{k\prime}, f_1^{k\prime}, \ldots, f_T^{k\prime})'$, $\Lambda^k = (\lambda_1^k, \ldots, \lambda_N^k)'$, $F^k = (f_1^k, \ldots, f_T^k)'$. Consider the normalizations for factors and loadings discussed above, define $\hat{\Lambda}^k = (\hat{\lambda}_1^k, \ldots, \hat{\lambda}_N^k)'$ and write

$$(\hat{\Lambda}^k)'\hat{\Lambda}^k/N = \text{diag}\left(\hat{\sigma}_{N,1}^k, \ldots, \hat{\sigma}_{N,k}^k\right).$$

The rank minimization criterion to estimate the number of factors r is defined as:

$$\hat{r}_{\text{rank}} = \sum_{j=1}^{k} \mathbf{1}\{\hat{\sigma}_{N,j}^k > P_{NT}\},$$

where P_{NT} is a sequence that goes to 0 as $N, T \to \infty$. In other words, $\hat{r}_{\text{rank}}$ can be interpreted as a rank estimator of $(\hat{\Lambda}^k)'\hat{\Lambda}^k/N$ since this average converges to a matrix with rank r, where P_{NT} can be viewed as a cutoff value determining that asymptotic rank. In particular, CDG (2021) find that the following choice of P_{NT} works well in practice,

$$P_{NT} = \hat{\sigma}_{N,1}^k \cdot \left(\frac{1}{L_{NT}^2}\right)^{1/3}$$

3.3 Relative Performance of PCA and QFA in a DGP with Outliers

One of the main characteristics of climate change is the existence of extreme events and the presence of outliers. In particular, the proportion of records that can be considered as outliers in the temperature data that we use in the empirical analysis is around 4% of the sample per year

(see Figure 1).[9] Therefore, this subsection studies how robust is our novel proposed methodology to outliers.

As mentioned earlier, at $\tau = 0.5$, QFA can be viewed as a robust QR alternative to the PCA estimators. By the same token, the QFA estimator of the number of factors should also be more robust to outliers and heavy tails than the IC-based method of Bai and Ng (2002). In what follows, we confirm these two claims by means of a few Monte Carlo simulations.

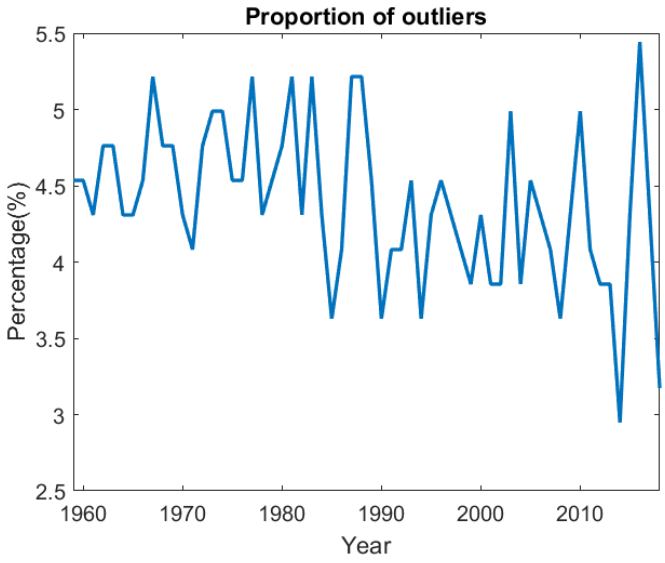

Figure 1: Proportion of outliers in temperature data

In particular, the following DGP is considered:

$$X_{it} = \sum_{j=1}^{3} \lambda_{ji} f_{jt} + u_{it},$$

where $f_{1t} = 0.8 f_{1,t-1} + \epsilon_{1t}$, $f_{2t} = 0.5 f_{2,t-1} + \epsilon_{2t}$, $f_{3t} = 0.2 f_{3,t-1} + \epsilon_{3t}$, $\lambda_{ji}, \epsilon_{jt}$ are all independent draws from $\mathcal{N}(0,1)$, and $u_{it} \sim i.i.d\ B_{it} \cdot \mathcal{N}(0,1) + (1 - B_{it}) \cdot \text{Cauchy}(0,1)$, where B_{it} are i.i.d Bernoulli random variables with means equal to 0.98 and Cauchy$(0,1)$ is the standard Cauchy distribution. Thus, approximately 2% of the idiosyncratic errors are generated as outliers.

We consider four estimators of the number of factors r: two estimators based on PC_{p1}, IC_{p1} of Bai and Ng (2002), the Eigenvalue Ratio (ER) estimator proposed by Ahn and Horenstein (2013) and CDG's (2021) rank-minimization estimator with P_{NT} chosen as in section 3.3. We set $k = 8$ for all four estimators, and consider $N, T \in \{50, 100, 200, 500\}$.

[9] The proportion of outliers in a given year is determined by the number of stations whose temperature record is above two standard deviations around the mean global temperature in that year.

Table 1 reports the following fractions for each estimator having run 1000 replications:

$$[\text{proportion of } \hat{r} < 3 \ , \ \text{proportion of } \hat{r} = 3 \ , \ \text{proportion of } \hat{r} > 3 \ ]$$

It becomes evident that PC_{p1} and IC_{p1} almost always overestimate the number factors, while the ER estimator tends to underestimate them, though to a lesser extent than what PC_{p1} and IC_{p1} overestimate them. By contrast, the rank-minimization estimator selects more accurately the right number of factors.

Table 1: AFM with Outliers in the Idiosyncratic Errors: Estimating the Number of Factors

N	T	PC_{p1} of BN			IC_{p1} of BN			Eigenvalue Ratio			Rank Estimator		
50	50	[0.00	0.04	0.96]	[0.00	0.14	0.86]	[0.26	0.30	0.44]	[0.47	0.53	0.00]
50	100	[0.00	0.02	0.98]	[0.00	0.05	0.95]	[0.33	0.19	0.48]	[0.40	0.60	0.00]
50	200	[0.00	0.00	1.00]	[0.00	0.01	0.99]	[0.41	0.12	0.47]	[0.33	0.67	0.00]
50	500	[0.00	0.00	1.00]	[0.00	0.00	1.00]	[0.56	0.07	0.37]	[0.29	0.71	0.00]
100	50	[0.00	0.02	0.98]	[0.00	0.05	0.95]	[0.34	0.18	0.48]	[0.39	0.61	0.00]
100	100	[0.00	0.00	1.00]	[0.00	0.01	0.99]	[0.41	0.13	0.46]	[0.10	0.90	0.00]
100	200	[0.00	0.00	1.00]	[0.00	0.00	1.00]	[0.48	0.07	0.45]	[0.06	0.94	0.00]
100	500	[0.00	0.00	1.00]	[0.00	0.00	1.00]	[0.65	0.05	0.30]	[0.02	0.98	0.00]
200	50	[0.00	0.00	1.00]	[0.00	0.01	0.99]	[0.45	0.10	0.45]	[0.37	0.63	0.00]
200	100	[0.00	0.00	1.00]	[0.00	0.00	1.00]	[0.48	0.08	0.44]	[0.10	0.90	0.00]
200	200	[0.00	0.00	1.00]	[0.00	0.00	1.00]	[0.63	0.06	0.31]	[0.00	1.00	0.00]
200	500	[0.00	0.00	1.00]	[0.00	0.00	1.00]	[0.76	0.08	0.16]	[0.00	1.00	0.00]
500	50	[0.00	0.00	1.00]	[0.00	0.00	1.00]	[0.57	0.08	0.35]	[0.36	0.64	0.00]
500	100	[0.00	0.00	1.00]	[0.00	0.00	1.00]	[0.68	0.06	0.26]	[0.05	0.95	0.00]
500	200	[0.00	0.00	1.00]	[0.00	0.00	1.00]	[0.76	0.08	0.16]	[0.00	1.00	0.00]
500	500	[0.00	0.00	1.00]	[0.00	0.00	1.00]	[0.80	0.10	0.10]	[0.00	1.00	0.00]

Note: The DGP considered in this Table is: $X_{it} = \sum_{j=1}^{3} \lambda_{ji} f_{jt} + u_{it}$, where $f_{1t} = 0.8 f_{1,t\ 1} + \epsilon_{1t}$, $f_{2t} = 0.5 f_{2,t\ 1} + \epsilon_{2t}$, $f_{3t} = 0.2 f_{3,t\ 1} + \epsilon_{3t}$, $\lambda_{ji}, \epsilon_{jt} \sim i.i.d\ \mathcal{N}(0,1)$, $u_{it} \sim i.i.d\ B_{it} \cdot \mathcal{N}(0,1) + (1 - B_{it}) \cdot \text{Cauchy}(0,1)$ where $B_{it} \sim i.i.d$ Bernoulli(0.98). For each estimation method, the [proportion of $\hat{r} < 3$, proportion of $\hat{r} = 3$, proportion of $\hat{r} > 3$] is reported from 1000 replications.

Next, to compare the PCA and QFA estimators of the common factors in the previous DGP, let us assume that $r = 3$ is known. We first get the PCA estimator (denoted $\hat{F}_{PCA}$), and then obtain the QFA estimator at $\tau = 0.5$ (denoted $\hat{F}^{0.5}_{QFA}$) using the IQR algorithm. Next, each of the true factors is regressed on $\hat{F}_{PCA}$ and $\hat{F}^{0.5}_{QFA}$ separately, and their average (adjusted) R^2 from 1000 replications are reported in Table 2 as an indicator of how well the space of the true factors

is spanned by the estimated factors. As can be inspected, the PCA estimators are not very successful in capturing the true common factors, while the QFA estimators approximate them very satisfactorily, even when N, T are not too large. Thus, this simulation exercise provides strong evidence in favour of using QFA instead of PCA in those cases where the idiosyncratic error terms in AFM exhibit heavy tails and outliers.

Table 2: AFM with Outliers in the Idiosyncratic Errors: Estimation of the Factors

		Regress F on $\hat{F}_{PCA}$			Regress F on $\hat{F}^{0.5}_{QFA}$		
N	T	f_1	f_2	f_3	f_1	f_2	f_3
50	50	0.939	0.810	0.686	0.987	0.975	0.968
50	100	0.931	0.718	0.578	0.987	0.975	0.968
50	200	0.890	0.589	0.412	0.987	0.975	0.968
50	500	0.807	0.405	0.252	0.988	0.975	0.968
100	50	0.928	0.738	0.595	0.993	0.986	0.984
100	100	0.921	0.630	0.441	0.994	0.988	0.984
100	200	0.857	0.479	0.285	0.994	0.988	0.985
100	500	0.713	0.294	0.138	0.994	0.988	0.984
200	50	0.890	0.657	0.513	0.997	0.994	0.992
200	100	0.858	0.514	0.333	0.997	0.994	0.993
200	200	0.779	0.358	0.178	0.997	0.994	0.992
200	500	0.530	0.131	0.051	0.997	0.994	0.992
500	50	0.819	0.501	0.371	0.998	0.997	0.996
500	100	0.725	0.327	0.196	0.999	0.998	0.997
500	200	0.546	0.165	0.062	0.999	0.998	0.997
500	500	0.273	0.036	0.018	0.999	0.998	0.997

Note: The DGP considered in this Table is: $X_{it} = \sum_{j=1}^{3} \lambda_{ji} f_{jt} + u_{it}$, where $f_{1t} = 0.8 f_{1,t\ \ 1} + \epsilon_{1t}$, $f_{2t} = 0.5 f_{2,t\ \ 1} + \epsilon_{2t}$, $f_{3t} = 0.2 f_{3,t\ \ 1} + \epsilon_{3t}$, $\lambda_{ji}, \epsilon_{jt} \sim i.i.d\ \mathcal{N}(0,1)$, $u_{it} \sim i.i.d\ B_{it} \cdot \mathcal{N}(0,1) + (1 - B_{it}) \cdot \text{Cauchy}(0,1)$ where $B_{it} \sim i.i.d$ Bernoulli(0.98). For each estimation method, we report the average R^2 in the regression of (each of) the true factors on the estimated factors by PCA and QFA (assuming the number of factors to be known).

4 Climate Change and CO_2 Concentrations

4.1 Data description

For the empirical analysis, we use data from the Climatic Research Unit (CRU) at University of East Anglia. In principle, CRU provides monthly and annual data of land and sea temperatures in North and South hemispheres from 1850 to the present, collected at different stations around the globe. However, a limitation of this dataset is that the number of stations fluctuates each year and its geographic distribution of stations is far from being homogeneous. In effect, a higher concentration of stations is reported for the U.S., Southern Canada, Europe, and Japan, while lesser coverage is reported in South America, Africa, and Antarctica. Thus, to guarantee some stability in the distribution of temperatures, we restrict the sample to 1959-2018 when data for those stations is available each year. Applying this procedure, we construct a balanced panel of local mean annual temperatures for 441 stations (N) observed over 60 periods (T). Figure 2 plots the temperature time series for five selected stations in our dataset.

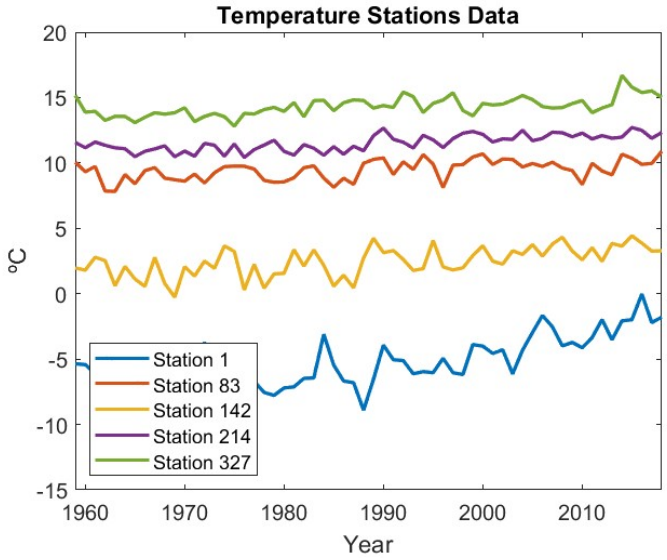

Figure 2: Temperature series in selected stations

Data on CO_2 is obtained from the Global Carbon Budget (GCB) series, compiled by Friedlingstein et al. (2021) and available at https://www.icos-cp.eu/science-and-impact/global-carbon-budget/2021. The specific information on atmospheric CO_2 concentrations that we use is drawn from Dlugokencky and Tans (2020) for the period 1959-2018 and measured in gigatons of carbon (GtC) per year.[10] Figure 3 displays the above series, both in levels (Panel [a]) and in first dif-

[10]The same series has been used recently by Bennedsen et al. (2020) in the estimation of a multivariate dynamic model involving the main variables included in the Global Carbon Budget (GCB).

ferences (Panel [b]). For the early period 1959-1980, estimations are based on Mauna Loa and South Pole stations as observed by the CO_2 Program at Scripps Institution of Oceanography; from 1980 onward it corresponds to global averages estimated from multiple stations run by the National Oceanic and Atmospheric Administration (NOAA) and Earth System Research Laboratory (ESRL). The robustness of our findings to the data source and transformations is analyzed in Appendix A.1 where we consider the CO_2 concentrations from Mauna Loa only and the Radiative forcing of CO_2 estimated by Hansen et al. (2011).

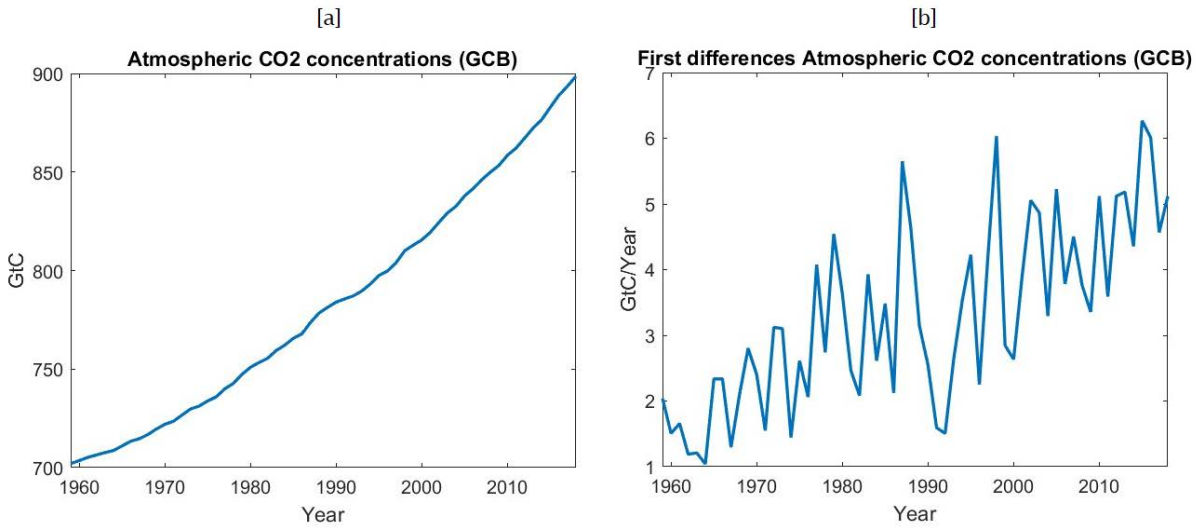

Figure 3: Atmospheric CO2 concentrations data (GCB)

In addition, our choice of sample period is also determined by the characteristics of the data on CO_2. In fact, before 1958, CO_2 concentrations were inferred from ice drilling, and it is only from 1959 that they started to be measured with instruments. As Pretis and Hendry (2013) discuss, pooling over different measurement regimes hinders the statistical analysis of the series and, in particular, the identification of the order of integration. Hence, these arguments also support restricting attention to the above-mentioned sample period. [11]

4.2 Testing for trends in temperature data

Gadea and Gonzalo (2020) provide a methodology to test for the existence of trends in the unconditional distributional characteristics (moments, quantiles, etc.) of global temperatures. Treating temperatures as a functional stochastic process, their distributional characteristics can be thought of as time-series objects to which one could apply standard testing procedures. For

[11] Other papers analyzing the relationship between temperature and CO_2 over the same time window include Pretis (2020) and Bennedsen et al. (2022)

example, the proposed robust linear trend-test is based in the statistical significance of the β coefficient in the following least-squares regression:

$$C_t = \alpha + \beta t + v_t, t = 1, ...T, \tag{4}$$

where C_t denotes a particular distributional characteristic of interest (e.g. a given quantile). The asymptotic properties of the OLS estimator in (4) depend on the summability order of the unknown trend component, $k \geq 0$, defined as follows. Let $C_t = h(t) + v_t$, where v_t is an $I(0)$ process and $h(t)$ is the unknown trend polynomial process of order k with coefficients β_k. Then its summability order becomes $S_T = \frac{1}{T^{1+k}} \sum_{t=1}^{T} h(t)$. The OLS estimated coefficient in (4) becomes:

$$\hat{\beta} = \frac{\sum tC_t - T\bar{t}\bar{C}}{\sum t^2 - T\bar{t}^2}, \tag{5}$$

where $\sum tC_t = T^{2+k}\frac{1}{T}\sum(\frac{t}{T})\frac{C_t}{T^k})$ and $\sum t^2 = T^3\frac{1}{T}\sum(\frac{t}{T})^2$, so that $T^{3/2}(\hat{\beta} - T^{k-1}\beta_k) = O_p(1)$, implying consistency if $k = 0, 1$. In such cases, it can be verified that $t_{\beta=0} \to N(0, 1)$, implying that the linear test based on Equation (4) can detect any type of trend even if it is non-linear.

This test is implemented as a preliminary inspection of the statistical properties of temperatures in our dataset. Using their cross-sectional distribution, a set of representative distributional characteristics (mean, sd, quantiles, etc.) are estimated for the sample period 1959-2018. Figure 4 presents the plots of the mean and quantiles q10, q25, q50, q75, and q90. According to the evidence shown in Table 3, a linear trend component is detected in most of these characteristics (except in the inter-quantile rage, iqr). Moreover, the GW phenomenon is clearly heterogeneous along the temperature distribution since the slope of the trend coefficients in the lower quantiles is steeper than those in the mean, median, and upper quantiles. Interestingly, these results are qualitatively similar to those reported by Gadea and Gonzalo (2020) using a different dataset over the longer period 1880-2015.

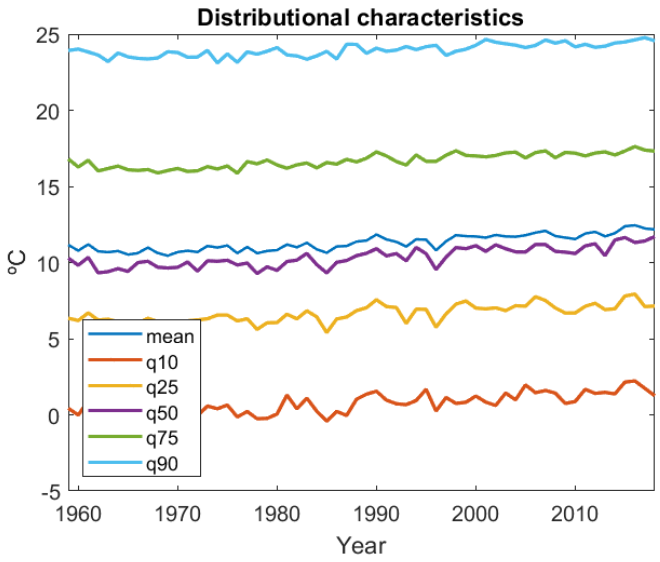

Figure 4: Estimated distributional characteristics

Table 3: Gadea-Gonzalo Trend test (1959-2018).

Characteristic	Test-statistic	p-value
mean	0.0264	0.0000
sd	-0.0050	0.0000
min	0.0605	0.0000
max	0.0407	0.0000
iqr	-0.0013	0.2461
kur	-0.0018	0.0054
skw	0.0007	0.0114
q01	0.0468	0.0000
q05	0.0385	0.0000
q10	0.0302	0.0000
q25	0.0240	0.0000
q50	0.0301	0.0000
q75	0.0227	0.0000
q90	0.0181	0.0000
q95	0.0161	0.0000
q99	0.0215	0.0000

Note: Annual distributional characteristics are estimated using the cross-sectional distribution at each year (1959-2018). OLS estimates and HAC $t_{\beta=0}$ p-values from regression (4) are reported.

4.3 Quantile Factor Analysis

The previous evidence opens the door to analyze heterogeneous association patterns between local (station-level) temperatures and atmospheric CO_2 concentrations using our proposed QFA procedure. To do so, we first estimate the quantile-dependent common factors of the temperature panel, and then use them as dependent variables in predictive-association regressions with CO_2 concentrations.

As mentioned earlier, a key requirement for extracting the QFA (and the PCA) factors is the absence of stochastic and deterministic trends in the individual temperature processes. Appendix A.2 reports the results of applying standard Augmented Dickey-Fuller (ADF) tests in regressions with a linear trend component. As can be observed, the null of unit root is rejected in almost all the stations (except in five cases), as well as in the distributional characteristics of interest. Based on this evidence, we conclude that the individual temperature series are trend-stationary. Thus, linear detrending is implemented to achieve the above requirements.[12]

Accordingly, QFA is applied on the linearly detrended panel of station-level temperatures in its standardized format, so that the lower (higher) quantiles capture large negative (positive) variations of temperature around a linear trend. The number of factors are selected according to the rank-minimization criterion discussed in section 3.3 for a fine grid of quantile levels, τ, ranging from 0.01 to 0.99. As pointed out before, the number of factors varies across quantiles, declining as we move away from the median. In particular, these numbers are: 1 (at $\tau = 0.01$, 0.05, 0.10, 0.95, 0.90, and 0.99), 3 (at $\tau = 0.25$ and 0.75), and 4 (at $\tau = 0.5$). For illustrative purposes, Figure 5 shows that the estimated factors for the quantiles 0.01, 0.50, and 0.99 are fairly different. In addition, PCA is used to estimate the factors at the mean, where the number of factors being chosen according to the PC_{p1} criterion of Bai and Ng (2002) which selects 8. This is the maximum number imposed in the IQR computational algorithm.

To compare the QFA factors (denoted as $\hat{F}^{\tau}_{QFA}$) with the PCA factors (denoted as $\hat{F}_{PCA}$), we regress each element of $\hat{F}^{\tau}_{QFA}$ on the 8 $\hat{F}_{PCA}$ and compute their corresponding R^2 as a measure of correlation.[13] The results are shown in Table 4. It becomes clear that, for the quantiles at the center of the distribution ($\tau = 0.25$, 0.5, and 0.75), the estimated factors are highly correlated with the PCA factors, with all the R^2s exceeding 0.90, especially in the case of the factors for the median (above 0.98). By contrast, the QFA factors at the upper and lower quantiles (e.g. $\tau = 0.01$, 0.05, 0.95, and 0.99) exhibit much lower correlations with the PCA factors, with R^2s fluctuating between 0.6 and 0.7. Thus, there seems to be room for using QFA in this application since the factors at the extreme quantiles help identify different features of the temperature distribution which the factors at the medium quantiles are unable to capture.

[12] Note that these properties of the dependent variables precludes the use of cointegration in a bivariate setup.

[13] Recall that the PC_{p1} is chosen to select the number of PCA factors estimated in these regressions to play conservative.

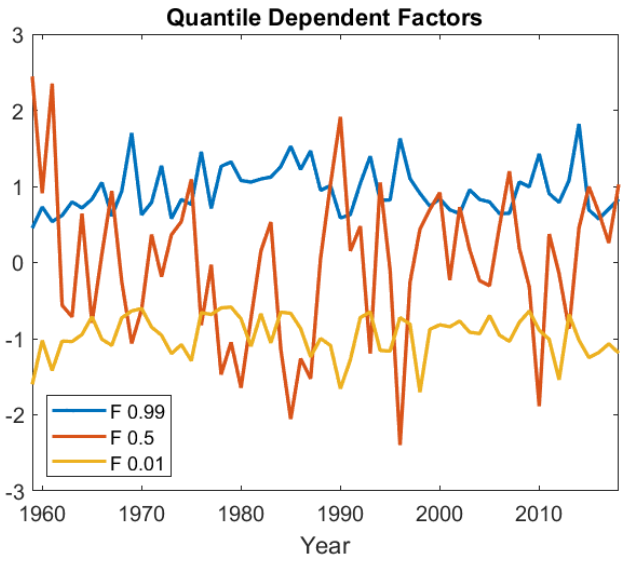

Figure 5: Estimated QFA factors

Table 4: Comparison of $\hat{F}_{QFA}$ and $\hat{F}_{PCA}$

	Elements of $\hat{F}^{\tau}_{QFA}$			
τ	1	2	3	4
0.01	0.6846			
0.05	0.7179			
0.10	0.7852			
0.25	0.9560	0.9624	0.9571	
0.50	0.9967	0.9968	0.9839	0.9806
0.75	0.9368	0.9481	0.9372	
0.90	0.7290			
0.95	0.6624			
0.99	0.6043			

Note: This Table reports the R^2 of regressing each element of $\hat{F}_{QFA}$ on $\hat{F}_{PCA}$. For $\hat{F}_{QFA}$, the numbers of estimated factors is obtained using the rank-minimization criterion, while for $\hat{F}_{PCA}$ the numbers of estimated factors are 8 for all datasets.

4.4 Temperature Factors and CO_2 Concentrations

Bivariate tests characterizing the predictive association between the estimated QFA temperature factors and a suitable transformation of the atmospheric CO_2 concentrations are implemented next. The statistical properties of the latter determine how it should be transformed. Panel [a] in Figure 3, makes it is clear that the CO_2 concentrations series from the GCB is upward trending and exhibits the typical dynamics of a unit root process. Moreover, the first-differenced series presented in Panel [b] point to an acceleration over the sample period following a linear trend. These features are confirmed by an ADF test reported in Appendix A.2, which suggest that the level of CO_2 concentrations has a unit root, while the growth rate is trend-stationary. As discussed in Bennedsen et al. (2020), these properties are consistent with a dynamic statistical model relating CO_2 concentrations, antrhopogenic CO_2 emissions, and absorption of CO_2 by the terrestrial, ocean and marine biospheres. The same statistical properties hold for the (logged) CO_2 concentrations from Mauna Loa and the CO_2 radiative forcing, series that we use in the robustness exercises in Appendix A.1.

Building on the previous observations, the tests to be implemented rely on linear regressions of the estimated QFA factors at each relevant quantile ($\hat{F}^{\tau}_{QFA,t}$) on p own lags and q lags of the linearly detrended changes in atmospheric CO_2 concentrations ($\hat{\Delta}CO_2$). In other words, this approach sheds light on whether past fluctuations in CO_2 concentrations changes around a linear trend have a predictive power on the common quantile-dependent factors of the stations-level temperature fluctuations, again around a linear trend. The specific regressions are given by:

$$\hat{F}^{\tau}_{QFA,t} = \alpha(\tau) + \sum_{i=1}^{p} \beta_i \times \hat{F}^{\tau}_{QFA,t-i} + \sum_{j=1}^{q} \gamma_j \times \hat{\Delta}CO_{2,t-j} + u_t, \tag{6}$$

where the lag lengths p and q are selected through a general-to-specific approach.

Given that QFA and PCA estimated factors share the same rates of convergence, a similar condition to that used in Bai and Ng (2006) is required to replace the true quantile factors by the QFA estimated ones in Equation (6), namely, $\sqrt{T}/N \to 0$. This condition is easily verified in our finite sample since $\sqrt{60}/441 = 0.017$. Thus, the proposed test looks at the joint significance of the γ_j coefficients, $j = 1, ..., q$, by means of an F-statistic, and can be interpreted as a predictive-association test.[14]

Table 5 reports the p-values of the proposed tests. At a 10% significance level, we find that past values of the linearly detrended changes in atmospheric CO_2 concentrations have a predictive power on the current values of QFA factors of the temperature at the lower and middle

[14]The specification in first differences of (logged) CO_2 is further corroborated when run predictive regressions using $q + 1$ lags of that variable in levels since the sum of their estimated coefficients is not significantly different from zero, pointing to the use of q lags of the first-differenced series as the correct choice .

quantiles (τ from 0.01 to 0.75). Yet, this is not the case for the QFA factors at the extreme upper quantiles ($\tau = 0.90, 0.95$ and 0.99) where the null hypothesis of the test is not rejected. The detected heterogeneous predictive-association pattern is the main finding of our paper which, as Appendix A.1 shows, happens to be robust to alternative sources, units of measurement, and transformations of the CO_2 concentrations series.

Table 5: p-values of the predictive association tests for the QFA factors

		Elements of $\hat{F}^{\tau}_{QFA}$			
Regressor	τ	1	2	3	4
	0.01	**0.0713**			
	0.05	**0.0317**			
	0.10	**0.0546**			
	0.25	**0.0140**	0.5626	0.7714	
$\hat{\Delta}CO_2$	0.50	**0.0057**	0.2568	0.1899	**0.0835**
	0.75	0.9149	**0.0878**	0.2499	
	0.90	0.2747			
	0.95	0.6395			
	0.99	0.8308			

Note: This Table reports the p-values of the proposed F test for the joint significance of the coefficients γ_j, $j = 1, ...q$ in Equation (6). Lag lengths are chosen following a general to specific approach. p-values smaller than 0.1 are in bold.

As a complementary exercise, we study the predictive association between the 8 PCA mean factors selected with the PC_{p1} criterion and the suitable transformation of the CO_2 concentrations. Notice that under the standard factor model conditions, the average temperature across stations is equivalent to a linear combination of these 8 PCA factors. Table 6 shows the p-values for the proposed tests. In agreement with the results for the median factors, a significant predictive association is detected for some of the PCA factors.

Table 6: p-values of the predictive association tests for the PCA factors

	Elements of $\hat{F}_{PCA}$							
Regressor	1	2	3	4	5	6	7	8
$\hat{\Delta}CO_2$	**0.0138**	0.3556	0.2285	0.2202	**0.0046**	0.6469	0.4759	**0.0135**

Note: This Table reports the p-values of the proposed F test for the joint significance of the coefficients γ_j, $j = 1, ...q$ in the corresponding version of Equation (6). Lag lengths are chosen following a general to specific approach. p-values smaller than 0.1 are shown in bold.

5 Discussion

The link between temperature and CO_2 concentrations is a long standing issue in the climate science literature. Our analysis of climate sensitivity departs from the standard approach that is typically quantified as warming per doubling of CO_2 (see Sherwood et al. (2020) for a recent overview of this line of research). In particular, we use a different metric, namely, one that focuses on the predicting power of CO_2 on temperature, where the former is measured as levels of global concentrations. In this fashion, our approach complements other research on heterogeneous climate sensitivity like Shindell and Faluvegi (2009) where the sensitivity of regional climate to changes in CO_2 (spatial heterogeneity) is investigated. Moreover, this non-uniformity seems to be a finding associated only to CO_2. In Appendix A.3 we consider the radiative forcing of other greenhouse gases as methane, ozone, or nitrous oxide, as well as natural forcings such as solar irradiance or volcanic activity, finding that predictive association holds for all quantile factors.

Admittedly, we still do not have a clear physical reason explaining our findings. Yet, a possible conjecture would go in parallel with the causes that explain the well-known diurnal asymmetry: the night-time temperatures have increased more rapidly than day-time temperatures (see Davy et al. (2017)). The proposed factors behind this asymmetry could rely on changes in cloud covering, precipitation, soil moisture, the planetary boundary layer, etc.

Finally, in spite of addressing a different research question related to GW, our results seem to be in line with the evidence reported in Gadea and Gonzalo (2020, 2023), where GW is also found to be non-uniform: lower temperatures increase much more than the medium and higher ones. The lower unconditional quantiles in their study correspond to the Arctic region. However, an increase of CO_2 concentrations will have unforeseen consequences (that is, whatever happens in the Poles does not remain there): ice melting, sea level increases, floods, migrations, extreme events, etc. All these events are further aggravated by their own feedback effects due to the reduction in the surface albedo (less solar energy is reflected out to space) and by the release of more greenhouse gasses (CO_2, and Methane) from the permafrost melting. In this respect, we highlight that non-uniform climate sensitivity is not regionally concentrated but rather affects

all the regions around the Globe. In particular, the growth rate of CO_2 emissions predicts (positively) the periods where the temperature decreases or does not increase much.

In view of this evidence, further research should aim at jointly analyzing the heterogeneous predictive power of (CO_2 for GW at both the spatial (across stations) and temporal dimensions. Preliminary results in a ongoing project on this issue point in the same direction as the ones reported in this paper. Hopefully, this quantitative analysis will help in the design of more efficient mitigation climate policies.

6 Conclusions

In this paper we test for predictive association between (detrended) CO_2 concentrations changes and temperatures from 441 weather stations in the Northern and Southern hemispheres over the period 1959-2018. Using the QFA methodology proposed by CDG (2021), we retrieve the quantile-dependent common factors and their number at different quantiles. We apply predictive association tests of different CO_2 concentration measures on those factors. The specification of the corresponding dynamic predictive equations is helped by the methodology proposed by Gadea and Gonzalo (2020) to detect deterministic and stochastic trends in different moments/quantiles of the distribution of temperature and by ADF tests for unit roots. As a by-product of the analysis, it is shown that QFA is a much more robust estimation method that standard PCA in the presence of outliers, as is the case in climate data.

Our main finding is that CO_2 concentrations changes have stronger predictive power for factors at the lower quantiles of the temperature fluctuations around a linear trend than at the middle and upper quantiles. We stress once again that this result is not picked up by the use the PCA mean factors since they capture common features of all temperatures whereas QFA factors capture common features at each quantile. Thus, as discussed earlier, we interpret our results as complementary to the available on climate sensitivity (see e.g. (Sherwood et al., 2020)).

A Appendix

A.1 Robustness

The main finding of the paper is the heterogeneous predictive-association between a suitable transformation of the atmospheric CO_2 concentrations series from the GCB and the QFA factors of the linearly detrended panel of local temperatures. To investigate the robustness of our findings to the data source, units of measurement, and transformation, the analysis is repeated using two additional series of atmospheric CO_2 concentrations as regressors. The first is the mean annual CO_2 concentrations series measured in parts per million by volume (ppmv) as obtained from direct measurements at Mauna Loa ($CO_{2,MLO}$).[15] The second corresponds to the Effective Radiative Forcing (ERF) from the CO_2 series ($ERF_t^{CO_2}$)[16] measured in Watts per meter squared (Wm^{-2}) and obtained from Hansen et al. (2011).[17] From Figures A.1 and A.2 it is clear that the dynamics of both series are similar to the dynamics of the CO_2 concentrations from the GCB.

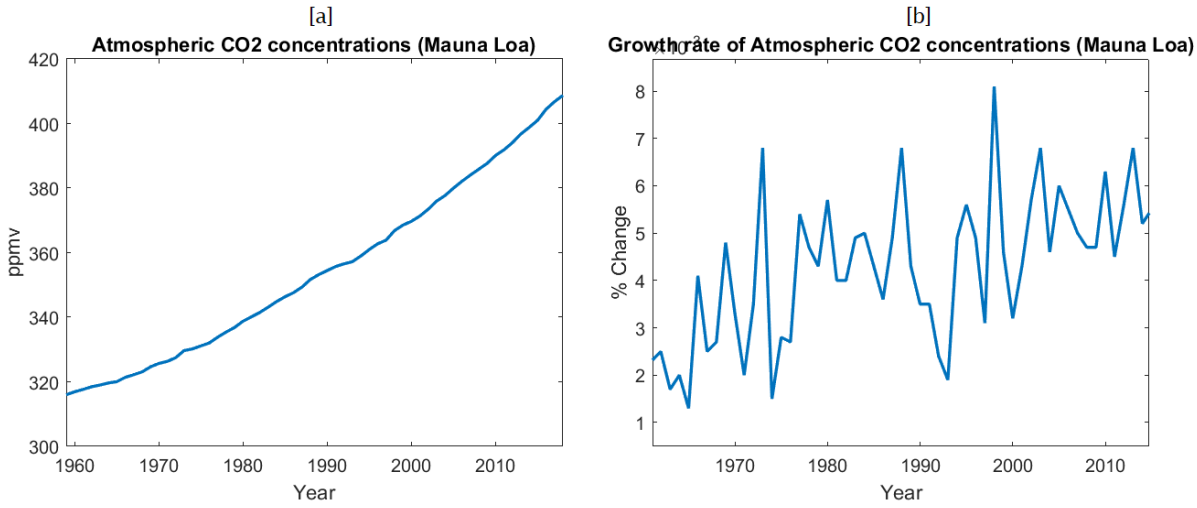

Figure A.1: Atmospheric CO2 concentrations (MLO)

[15] Available at https://gml.noaa.gov/ccgg/trends/data.html.

[16] By definition, $ERF_t^{CO_2} = 5.35 \times ln(CO2_t/CO2_{base})$, where $CO2_t$ are the CO_2 concentrations at a given year t and $CO2_{base}$ are the CO_2 concentrations at a given base year.

[17] Available at http://www.columbia.edu/ mhs119/Forcings/.

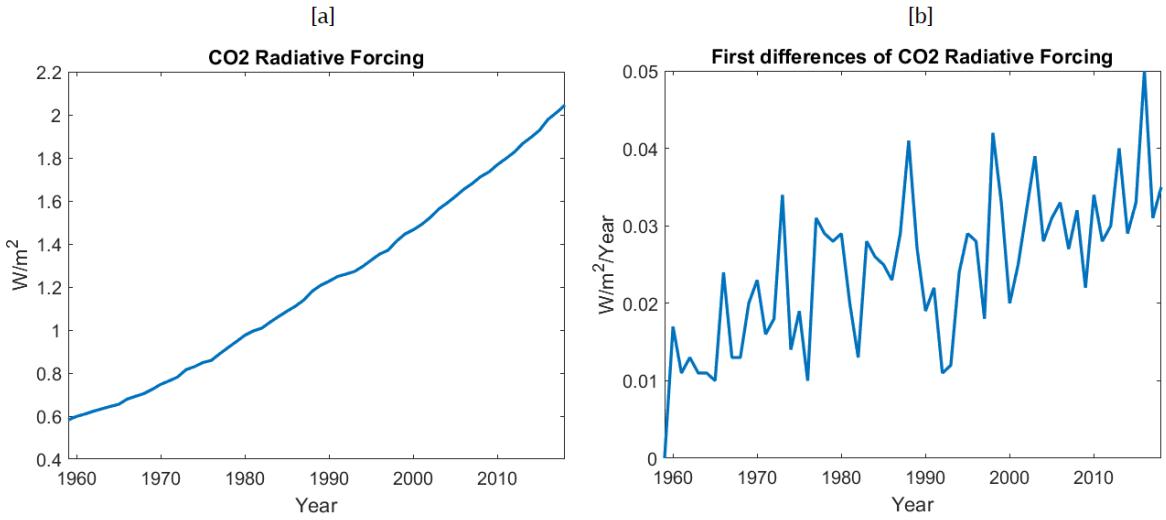

Figure A.2: Effective Radiative Forcing from CO2

Provided that the series in Panel [b] of the previous Figures are trend-stationary according to the ADF tests for unit roots described in Appendix A.2, in Equation (6) we consider as regressors the linearly detrended growth rate of CO_2 concentrations from Mauna Loa ($\hat{\Delta} log(\mathrm{CO}_{2,MLO})$) and the linearly detrended first differences the ERF from CO_2 ($\hat{\Delta} ERF_t^{CO_2}$). Results reported in Table A.1 indicate that the main finding of the paper holds for the two additional series. In both cases we reject the null of the implemented test for QFA factors at lower and middle quantiles (τ from 0.01 to 0.75), while the null of the test is not rejected at upper quantiles ($\tau = 0.90, 0.95$ and 0.99).

Table A.1: p-values of the predictive association tests for the QFA factors

		Elements of $\hat{F}^{\tau}_{QFA}$			
Regressor	τ	1	2	3	4
$\hat{\Delta} log(\mathrm{CO}_{2,MLO})$	0.01	**0.0754**			
	0.05	0.1959			
	0.10	**0.0143**			
	0.25	**0.0533**	**0.0131**	0.3997	
	0.50	**0.0057**	**0.0154**	0.5053	**0.0583**
	0.75	**0.0357**	0.1356	0.6663	
	0.90	0.2092			
	0.95	0.2763			
	0.99	0.2076			
$\hat{\Delta} ERF_t^{CO_2}$	0.01	**0.0088**			
	0.05	**0.0970**			
	0.10	**0.0187**			
	0.25	**0.0259**	**0.0041**	0.3057	
	0.50	**0.0236**	**0.0002**	0.1732	**0.0706**
	0.75	**0.0365**	**0.0582**	0.4800	
	0.90	0.1213			
	0.95	0.1027			
	0.99	0.3969			

Note: This Table reports the p-values of the proposed F test for the joint significance of the coefficients γ_j, $j = 1, ...q$ in the corresponding version of Equation (6) for $\hat{\Delta} log(\mathrm{CO}_{2,MLO})$ and $\hat{\Delta} ERF_t^{CO_2}$. Lag lengths are chosen following a general to specific approach. p-values smaller than 0.1 are in bold.

A.2 Unit root tests

Augmented Dickey-Fuller (ADF) tests for unit roots are implemented to guide the choice of the suitable transformations of the data coherent with the methodological devices at hand. For the case of the temperature series, individual unit root tests are implemented on the full set of stations. Additionally, such tests are implemented on the distributional characteristics of interest (moments, quantiles, etc.) estimated using the cross sectional distribution of temperatures. In the test specification we include intercept and a linear trend. As reported in Table A.2, the null hypothesis of unit root is rejected in 98.66% of the stations; it is not rejected only in 5 stations out for 441. In a similar direction, the ADF test suggests no unit roots on the different distributional characteristics. This piece of evidence suggest that temperature series do not contain unit roots, but rather follow trend-stationary processes.

Table A.2: ADF unit root tests.

ADF test by stations		
Percentage of rejections	98.66 %	
Number of nonrejections	5	
ADF test by characteristics		
Characteristics	Test-statistic	p-value
mean	-5.6261	0.0001
sd	-6.5609	0.0000
min	-6.2759	0.0000
max	-8.1441	0.0000
iqr	-6.1115	0.0000
kur	-6.5601	0.0000
skw	-7.6596	0.0000
q01	-8.7850	0.0000
q05	-8.0555	0.0000
q10	-5.0763	0.0006
q25	-5.4339	0.0002
q50	-6.1891	0.0000
q75	-6.7940	0.0000
q90	-6.3746	0.0000
q95	-6.6529	0.0000
q99	-6.2450	0.0000

Note: Annual distributional characteristics estimated using the cross-sectional distribution at each year (1959-2018). Significance level of 5% is considered in the individual tests. ADF-test equations include intercept and trend. Lag-selection conducted using SBIC criterion.

In Table A.3 the p-values of the ADF test implemented on the series of CO_2 and ERF are reported. The three series related to CO_2 ($CO_{2,GCB}$, $log(CO_{2,MLO})$, and ERF^{CO_2}) contain a

unit root in its levels, while the first differences are trend-stationary. Regarding the ERF series, the test results indicates that the non-CO_2 ERF is stationary in its levels while the total ERF is $I(1)$. The unit root analysis for this set of series is consistent with Pretis (2020) and Bennedsen et al. (2022).

Table A.3: ADF unit root tests

	Level series		**First differences**	
Variable	Constant	Constant and trend	Constant	Constant and trend
$CO_{2,GCB}$	7.7983	-0.4758	-0.4491	**-6.9110**
$log(CO_{2,MLO})$	4.3538	-1.3849	-1.3903	**-6.4887**
ERF^{CO_2}	6.4539	-0.9759	**-4.5028**	**-7.0202**
ERF^{nonCO_2}	**-4.3825**	**-5.7877**	**-7.6238**	**-7.5589**
ERF^{Tot}	-1.372	**-5.8092**	**-7.6105**	**-7.5558**

Note: This Table reports test-statistic of the ADF tests on the corresponding variable when the test includes only a constant or a constant and an intercept. Lags are selected using the BIC criterion. The values in bold indicate that the null of the test is rejected.

A.3 Other warming sources

Even though the interest of the paper is on the bivariate association between CO_2 and station-level temperatures, in this section we briefly examine the association with other greenhouse gases as methane, ozone, or nitrous oxide, as well as natural forcings such as solar irradiance or volcanic activity. In Figure A.3 we present the total ERF (ERF^{Tot}) (Panel [a]) and the ERF from other sources different than CO_2 (ERF^{nonCO_2}) (Panel [b]) as taken from Hansen et al. (2011). The ADF tests for unit root described in Appendix A.2 indicates that ERF^{Tot} is I(1) and ERF^{nonCO_2} is I(0).

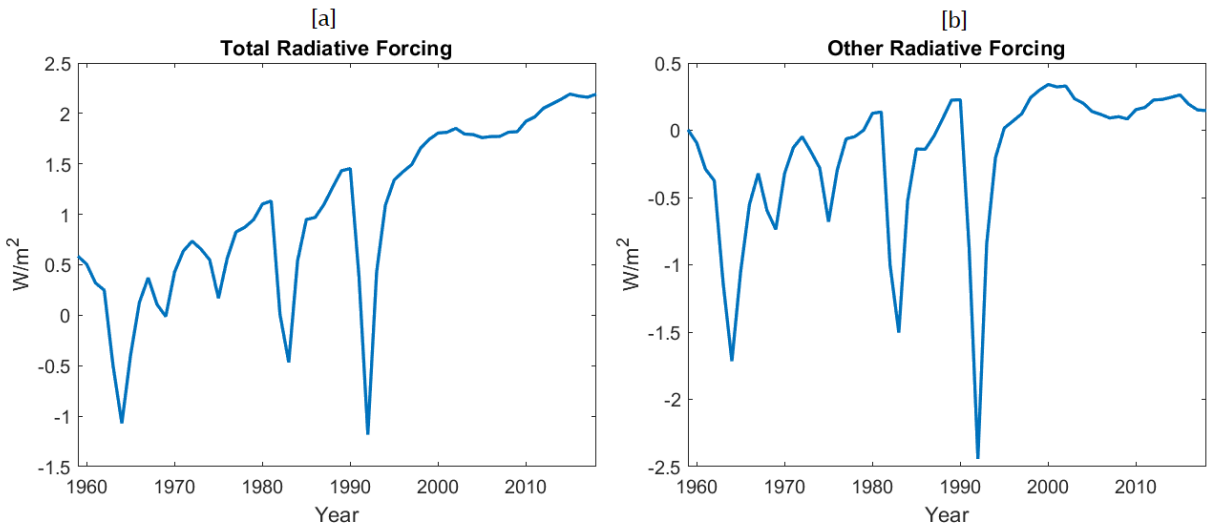

Figure A.3: Effective Radiative Forcing from all sources

The predictive-association analysis is conducted considering as regressors the first differences of ERF^{Tot} (ΔERF^{Tot}) or the levels of ERF^{nonCO_2}. As observed in Table A.4, when including other sources of warming (natural or anthropogenic) different from CO_2, a more homogeneous pattern of predictive association is obtained. In fact, when considering ERF^{nonCO_2}, it seems that the predictive association is stronger for the QFA factors of the medium and upper quantiles. A deeper examination of the patterns of association between local temperatures and different warming is beyond the scope of this study but constitutes an interesting avenue for future research.

Table A.4: p-values of the predictive association tests for the QFA factors

		Elements of $\hat{F}^{\tau}_{QFA}$			
Regressor	τ	1	2	3	4
ΔERF^{Tot}	0.01	**0.0000**			
	0.05	**0.0592**			
	0.10	**0.0077**			
	0.25	**0.0001**	0.6193	**0.0000**	
	0.50	**0.0008**	0.3757	**0.0057**	**0.0782**
	0.75	**0.0014**	0.7196	**0.0007**	
	0.90	**0.0065**			
	0.95	**0.0040**			
	0.99	**0.0434**			
ERF^{nonCO_2}	0.01	**0.0572**			
	0.05	**0.0554**			
	0.10	**0.0180**			
	0.25	**0.0001**	0.5056	**0.0000**	
	0.50	**0.0012**	0.5520	0.6629	0.1269
	0.75	**0.0155**	0.8938	**0.0011**	
	0.90	**0.0070**			
	0.95	**0.0147**			
	0.99	**0.0055**			

Note: This Table reports the p-values of the proposed F test for the joint significance of the coefficients γ_j, $j = 1, ...q$ in the corresponding version of Equation (6) for ΔERF^{Tot} and ERF^{nonCO_2}. Lag lengths are chosen following a general to specific approach. p-values smaller than 0.1 are in bold.

References

Ahn, S. C. and A. R. Horenstein (2013). Eigenvalue ratio test for the number of factors. *Econometrica 81*(3), 1203–1227.

Ando, T. and J. Bai (2020). Quantile co-movement in financial markets: A panel quantile model with unobserved heterogeneity. *Journal of the American Statistical Association 115*(529), 266–279.

Arrhenius, S. (1896). On the influence of carbonic acid in the air upon the temperature of the ground. *Philosophical Magazine and Journal of Science 41*(5), 237–276.

Bai, J. (2003). Inferential theory for factor models of large dimensions. *Econometrica 71*(1), 135–171.

Bai, J. and S. Ng (2002). Determining the number of factors in approximate factor models. *Econometrica 70*(1), 191–221.

Bai, J. and S. Ng (2006). Confidence intervals for diffusion index forecasts and inference for factor-augmented regressions. *Econometrica 74*(4), 1133–1150.

Bennedsen, M., E. Hillebrand., and S. J. Koopman (2020). A statistical model of the global carbon budget. *CREATES Research Papers 2020-18, Department of Economics and Business Economics, Aarhus University*.

Bennedsen, M., E. Hillebrand, and J. Zhou-Lykke (2022). Global temperature projections from a statistical energy balance model using multiple sources of historical data. *arXiv:2205.10269*.

Castle, J. and D. Hendry (2020). Climate econometrics: An overview. *Foundations and Trends in Econometrics 10*(3-4).

Chang, Y., R. Kaufmann, C. Sik-Kim, I. Miller, Y. Park, and S. Park (2020). Evaluating trends in time series of distributions: A spatial fingerprint of human effects on climate. *Journal of Econometrics 214*(1), 274–294.

Chapman, S., D. Stainforth, and N. Watkins (2013). On estimating local long-term climate trends. *Phil Trans R Soc A 371:20120287*.

Chen, L., J. Dolado, and J. Gonzalo (2021). Quantile factor models. *Econometrica 89*(2), 875–910.

Chen, L., J. Gao, and F. Vahid (2022). Global temperatures and greenhouse gases: A common features approach. *Journal of Econometrics 230*, 240–254.

Davy, T., I. Esau, A. Chernokulsky, S. Outten, and S. Zilitinkevich (2017). Diurnal asymmetry to the observed global warming. *International Journal of Climatology 37*(1).

Dlugokencky, E. and P. Tans (2020). Trends in atmospheric carbon dioxide. technical report, national oceanic and atmospheric administration, earth system research laboratory.

Estrada, F., P. Perron, and B. Martínez-López (2013). Statistically derived contributions of diverse human influences to twentieth-century temperature changes. *Nature Geoscience 6*(12), 1050–1055.

Foote, E. (1856). Circumstances affecting the heat of the sun's rays. *The American Journal of Science and Arts 22*(2), 382–383.

Fourier, J. (1824). On the temperatures of the terrestrial sphere and interplanetary space. *Chimie et de Physique 27*, 136–167.

Friedlingstein, P., M. Jones, and et al (2021). Global carbon budget 2021. *Earth System Science Data Discussions*, 1–191.

Gadea, M. and J. Gonzalo (2020). Trends in distributional characteristics: Existence of global warming. *Journal of Econometrics 214*(1), 153–174.

Gadea, M. and J. Gonzalo (2023). Climate change heterogeneity: A new quantitative approach. *arXiv:2301.02648*.

Gao, J. and K. Hawthorne (2006). Semiparametric estimation and testing of the trend of temperature series. *Journal of Econometrics 9*(2), 332–355.

Gay-Garcia, C., F. Estrada, and A. Sánchez (2009). Global and hemispheric temperatures revisited. *Climatic Change 94*(3), 333–349.

Golub, G. H. and C. F. Van Loan (2013). *Matrix Computations*, Volume 3. JHU Press.

Hansen, J., M. Sato, P. Kharecha, and K. von Schuckmann (2011). Earth's energy imbalance and implications. *Atmospheric Chemistry and Physics 11*, 13421–13449.

Huber, P. (1981). *Robust Statistics*. John Wiley and Sons.

Ji, F., Z. Wu, J. Huang, and E. Chassignet (2014). Evolution of land surface air temperature trend. *Nature Climate Change 4*, 462–466.

Kaufmann, R., K. Kauppi, and J. Stock (2006). Emissions, concentrations, and temperature: A time series analysis. *Climatic Change 77*, 249–278.

Koenker, R. (2005). *Quantile Regression*. Number 38. Cambridge University Press.

Montamat, G. and J. Stock (2020). Quasi-experimental estimates of the transient climate response using observational data. *Climatic Change 160*, 361–371.

Phillips, P., T. Leirvik, and T. Storelvmo (2020). Econometric estimates of earth's transient climate sensitivity. *Journal of Econometrics 214*, 6–32.

Pretis, F. (2020). Econometric modelling of climate systems: The equivalence of energy balance models and cointegrated vector autoregressions. *Journal of Econometrics 214*(1), 256–273.

Pretis, F. and D. Hendry (2013). Comment on "polynomial cointegration tests of anthropogenic impact on global warming" by beenstock et al. (2012) – some hazards in econometric modelling of climate change. *Earth System Dynamics 4*, 375–384.

Previdi, M., K. Smith, and L. Polvani (2021). Arctic amplification of climate change: a review of underlying mechanisms. *Environmental Research Letters 16*(9), 1–26.

Sherwood, S., M. Webb, J. Annan, K. Armour, P. Forster, J. Hargreaves, G. Hegerl, K. Klein, S. Marvel, E. Rohling, M. Watanabe, T. Andrews, P. Braconnot, C. Bretherton, G. Foster, Z. Hausfather, A. von der Heydt, R. Knutti, T. Mauritsen, J. Norris, C. Proistosescu, M. Rugenstein, G. Schmidt, K. Tokarska, and M. Zelinka (2020). An assessment of earth's climate sensitivity usingmultiple lines of evidence. *Reviews of Geophysics, 58*(4).

Shindell, D. (2014). Inhomogeneous forcing and transient climate sensitivity. *Nature Climate Change 4*, 274–277.

Shindell, D. and G. Faluvegi (2009). Climate response to regional radiative forcing during the twentieth century. *Nature Climate Change 2*, 294–300.

Stips, A., D. Macias, C. Coughlan, E. Garcia-Gorriz, and X. Liang (2016). On the causal structure between co2 and global temperature. *Scientific Reports 6*(21691).

Tyndall, J. (1863). On radiation through the earth's atmosphere. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science 25*(167), 200–206.