

Heterogeneous predictive association of CO₂ with global warming

Liang Chen¹ | Juan J. Dolado² | Jesús Gonzalo² | Andrey Ramos²

¹Peking University HSBC Business School, China

²Universidad Carlos III de Madrid, Spain

Correspondence

Juan J Dolado, Department of Economics, Universidad Carlos III de Madrid, Madrid, Spain

Email: dolado@eco.uc3m.es

Abstract

Global warming is a non-uniform process across space and time. This opens the door to a heterogeneous relationship between CO₂ and temperature that needs to be explored going beyond the standard analysis based on mean temperature. We revisit this topic through the lens of a new class of factor models for high-dimensional panel data, called quantile factor models. 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 CO₂ concentrations helps to predict the underlying factors of the different quantiles of the distribution of (detrended) temperatures in the time dimension. We document that predictive association is greater at the lower and medium quantiles than at the upper quantiles of temperature in all stations, 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.

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 due mainly 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) with carbon dioxide (CO₂) concentrations.¹

From an econometrics viewpoint, several studies have used time series techniques to explore this topic empirically (see, for example, Kaufmann *et al.* 2006; Stips *et al.* 2016; Castle and

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2023 The Authors. *Economica* published by John Wiley & Sons Ltd on behalf of London School of Economics and Political Science.

Hendry 2020; Montamat and Stock 2020; Phillips *et al.* 2020; Pretis 2020; Chen *et al.* 2022). 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 CO₂ 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, like the order of integration or the type of deterministic 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. In particular, we propose a novel econometric methodology aimed at establishing predictive association between temperature and CO₂ concentrations, allowing for heterogeneity in this relationship. It relies on quantile factor models (QFMs) (Chen *et al.* 2021), whose use in this context can be motivated as follows.

Let a panel of temperature $\{X_{it}\}$ be available for $i = 1, \dots, N$ stations and $t = 1, \dots, T$ periods, together with time observations on CO₂, denoted as $\{Z_t\}$, $t = 1, \dots, T$, which are uniform across all stations. Different approaches are available to analyse the association between these two variables (or stationary transformations of them). For example, researchers can adopt a conditional regression approach for each individual station, $\mathbb{E}[X_{it}|Z_t] = \gamma_{0i} + \gamma_{1i}Z_t$, $i = 1, \dots, N$, 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 is smaller, not only can this procedure 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 average temperatures across weather stations for each period, and then estimate a single conditional regression of the form $\mathbb{E}[\bar{X}_t|Z_t] = \bar{\gamma}_0 + \bar{\gamma}_1 Z_t$, where \bar{X}_t is the series of mean temperatures, and $\bar{\gamma}_1$ is the mean of γ_1 . Likewise, by assuming a single common factor structure $X_{it} = \lambda_i f_t + \varepsilon_{it}$, 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 $\mathbb{E}[\hat{f}_t|Z_t] = \delta_0 + \delta_1 Z_t$. 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 the previous approaches is that they do not allow for the presence of heterogeneous patterns in the predictive association of CO₂ with temperature. Heterogeneity can be accounted for by adopting different quantile-estimation approaches depending on the dimension on which one wishes to focus: either heterogeneity of temperature distribution across space (i.e. weather stations) or across time. For example, to consider *spatial* heterogeneity—that is, whether CO₂ has stronger effects on temperatures at certain locations (e.g. when the ECS is higher for northern latitudes than in the tropics)—one could run OLS regressions of the unconditional quantile curves of the temperature distribution at time t (across N stations) on Z_t .³ Noticing that the unconditional temperature quantiles correspond to different latitudes, this would be equivalent to, for example, analysing whether CO₂ has a higher ECS on northern latitudes than in the tropics, an approach that has been adopted by in the climate literature by Leduc *et al.* (2016) and Miller (2023), among others. Yet a much less explored topic in that literature is the so-called heterogeneity of the temperature distribution over *time*: for example, whether CO₂ affects more strongly the lowest rather than the highest temperature in a given location, and if this effect remains uniform across all locations. This is the research question on which we focus in this paper. To address this issue, consider the conditional quantiles for temperature over time at each station i in the form of the quantile regression (QR) $Q_\tau[X_{it}|Z_t] = \beta_{0i}(\tau) + \beta_{1i}(\tau)Z_t$, where $0 < \tau < 1$ denotes the quantile level, Q_τ is the conditional quantile, and $\beta_{1i}(\tau)$ is the object of interest.⁴ One possibility is to run station-by-station QR of the temperature distribution on CO₂. However, given that N is large, this procedure is likely to

face the same problems as using the conditional regression approach for each individual station. Alternatively, one could run QR for each station, and then average the estimates $\hat{\beta}_{1i}(\tau)$ across units for each τ . However, this would yield 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), where a quantile factor structure is assumed for panel data on temperature by weather station. Accordingly, $X_{it} = \lambda'_i(\tau) f_i(\tau) + u_{it}(\tau)$, where $f_i(\tau)$ and $\lambda_i(\tau)$ are $r(\tau) \times 1$ 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 $\hat{f}_i(\tau)$ 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 (temperature quantiles), in the same way as the PCA factors are aggregators for the mean. In other words, our aim is to estimate common factors (to all stations) driving temperature at different quantiles, namely, those forces that drive lower, medium and higher temperatures in both the Antarctic and the Sahara (irrespective of whether the lowest temperature in the latter is much higher than the hottest temperature in the former), to then study the predictive power of CO₂ on these quantile factors. Hence, insofar as the quantiles are computed separately for each station, we allow for some individual heterogeneity while keeping our main focus on heterogeneity along the temperature distribution over time.

In line with the generalization of linear regression to QR models, QFMs can be understood as an extension of approximate factor models (AFMs) to allow for hidden factors shifting specific characteristics (moments or quantiles) of the distribution of temperature.⁵ As a simple illustration of the advantages of QFMs over AFMs, consider the factor structure in a standard *location-scale shift* model with the following data generating process (DGP): $X_{it} = \alpha_i f_{1t} + \eta_i f_{2t} \varepsilon_{it}$, with $f_{1t} \neq f_{2t}$ (both are scalars), $\eta_i, f_{2t} > 0$ and $\mathbb{E}(\varepsilon_{it}) = 0$. 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} .⁶ Such a DGP can be rewritten in QR format as $X_{it} = \lambda'_i(\tau) f_i + u_{it}(\tau)$, with $0 < \tau < 1$, $\lambda_i(\tau) = [\alpha_i, \eta_i Q_\varepsilon(\tau)]'$, where $Q_\varepsilon(\tau)$ represents the quantile function of ε_{it} , $f_i = [f_{1t}, f_{2t}]'$, $u_{it}(\tau) = \eta_i f_{2t} [\varepsilon_{it} - Q_\varepsilon(\tau)]$, and the conditional quantile is $Q_{u_{it}(\tau)}[\tau | f_i] = 0$. It can be checked that PCA will extract only the location-shifting factor f_{1t} in this model; it will fail to capture the scale-shifting factor f_{2t} . By contrast, the estimation of a QFM by the so-called *quantile factor analysis* (QFA) will be able to retrieve the space spanned by both factors in this DGP.⁷

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 an 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 an AFM from a QFM. This is especially relevant in the presence of outliers, where QFA will render valid estimation and inference, while an AFM may not work well.⁸ 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 such as Kaufmann

et al. (2006), Chang *et al.* (2020), Phillips *et al.* (2020) and Pretis (2020), among others, argue that temperature has stochastic trends, as is also the case for CO₂, and therefore defend the use of cointegration methods if appropriate. On the other hand, authors such as Gao and Hawthorne (2006), Gay-Garcia *et al.* (2009) and Estrada *et al.* (2013) claim that global temperature follows a trend-stationary process with non-linear 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. For example, Figure 1 plots temperature series (without detrending) for five randomly selected stations in our dataset, where it becomes clear that the station with lower temperature (station 1) exhibits a clear upward trend, whereas the others have smoother linear trends. Hence we linearly detrend X_{it} before QFA is implemented.⁹

Once the quantile-dependent common factors have been extracted, the next step in our proposed methodology is to determine the existence of predictive association between those factors and (a convenient transformation of) CO₂ 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 upward trend as shown in Figure 2 (see similar evidence in Bennedsen *et al.* 2023). Therefore we consider linearly detrended changes in atmospheric CO₂ concentrations as a predictor for the quantile-dependent common factor series of (detrended) temperature. In this respect, it noteworthy that the Arrhenius (1896) hypothesis stated the existence of a linear relation between the levels of temperature and the levels (or logs) of atmospheric CO₂. Yet there is strong evidence that during our sample period (1959–2018), temperature behaves as a trend-stationary variable whereas CO₂ levels are integrated of order 1 (denoted I(1)) around a quadratic trend. Hence, to run balanced regressions, one has to regress levels of temperature on first differences of CO₂ and a trend or, using the Frisch–Waugh–Lovell Theorem, detrended series of the dependent variable on detrended changes of the regressor plus lags. Taking this key point into consideration, the existence of predictive association is determined through an F -test for the joint significance of the coefficients associated to the lags of the CO₂ concentrations related variable in first differences.

Given the previous discussion, a brief summary of our empirical methodology follows to facilitate its understanding. For simplicity, let us focus on a given station i and a given

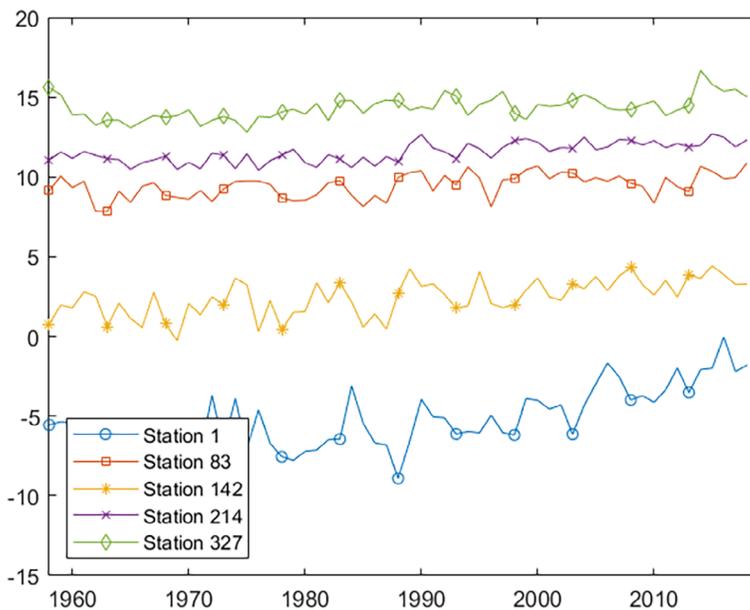


FIGURE 1 Temperature series in selected stations.

quantile τ of its temperature distribution over time. Denote temperature by X_{it} , and changes in CO_2 concentrations by Z_t . Then

$$X_{it} = p_x(t) + \tilde{X}_{it}, \tag{1}$$

$$\tilde{X}_{it}(\tau) = \lambda_i(\tau) f_i(\tau) + u_{it}(\tau), \tag{2}$$

$$Z_t = p_z(t) + \tilde{Z}_t, \tag{3}$$

$$f_i(\tau) = \beta_\tau \tilde{Z}_t + e_t, \tag{4}$$

where $p_x(t)$ and $p_z(t)$ are trend polynomials (selected to be linear) in temperature and other greenhouse gases (e.g. methane) different from CO_2 that also exhibit trends, \tilde{X}_{it} and \tilde{Z}_t are deviations of X_{it} and Z_t from their trends, $f_i(\tau)$ is the factor at quantile τ , and $\lambda_i(\tau)$ is the corresponding loading. Equation (4) is the predictive OLS regression run separately for each τ , and u_{it} and e_t are error terms. Substituting equations (3) and (4) into (2) yields

$$\tilde{X}_{it} = \lambda_i(\tau) \beta_\tau \tilde{Z}_t + (\lambda_i(\tau) e_t + u_{it}(\tau)) = \gamma_i(\tau) \tilde{Z}_t + v_{it}(\tau),$$

which implies the conditional quantile on which QFM is based, that is, $Q_\tau(\tilde{X}_{it}|\tilde{Z}_t) = \gamma_i(\tau) \tilde{Z}_t$.

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 that CO_2 concentrations (together with past temperatures) are the whole information set available for the econometrician. Since this is not realistic, we claim only 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 ECS or the TCR analysed in the standard literature, which seeks to establish a causal relationship through a more structural model.

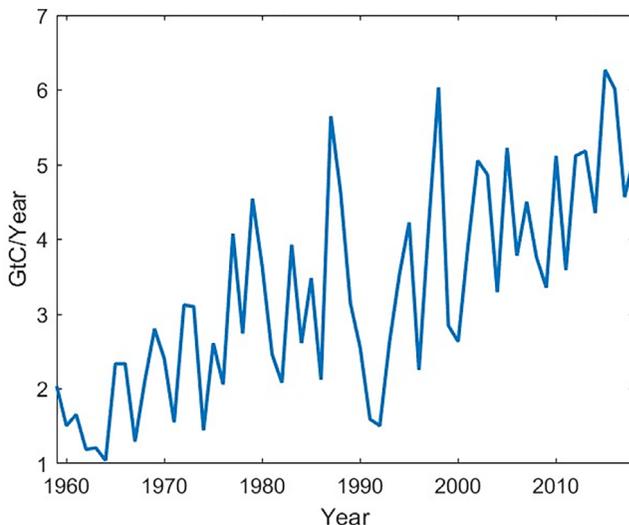


FIGURE 2 First differences of atmospheric CO_2 concentrations data (GCB).

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 that helps us to analyse the heterogeneous predictive association between CO₂ 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 focus mostly on the mean temperature. They also point out that future climate agreements should go beyond the mean temperature target and instead consider the whole temperature distribution and a CO₂ concentrations objective.

The rest of the paper is organized as follows. Section 2 defines the 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 CO₂ 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 the Chen *et al.* (2021) QFM approach.

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

$$Q_{X_{it}}[\tau|f_i(\tau)] = \lambda'_i(\tau) f_i(\tau),$$

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

Alternatively, the above equation implies that

$$X_{it} = \lambda'_i(\tau) f_i(\tau) + u_{it}(\tau), \quad (5)$$

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

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, namely $X_{it} = \alpha'_i f_{1t} + (\eta'_i f_{2t}) \varepsilon_{it}$, where $\{\varepsilon_{it}\}$ are zero-mean independent and identically distributed (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, that is, $Q_\varepsilon(0.5) = 0$, $\alpha_i, f_{1t} \in \mathbb{R}^{r_1}$, $\eta_i, f_{2t} \in \mathbb{R}^{r_2}$ 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 equation (5), with $\lambda_i(\tau) = [\alpha'_i, \eta'_i Q_\varepsilon(\tau)]'$, $f_i(\tau) = [f'_{1t}, f'_{2t}]$ for $\tau \neq 0.5$ and $\lambda_i(\tau) = \alpha_i$, $f_i(\tau) = f_{1t}$ 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, that is, $X_{it} = \alpha_i f_{1t} + f_{2t} \varepsilon_{it} + f_{3t} \varepsilon_{it}^3$, where ε_{it} follows a standard normal random variable with CDF $\Phi(\cdot)$. Then X_{it} has an equivalent representation in the form of equation (5), with $\lambda_i(\tau) =$

$[\alpha_i, \Phi^{-1}(\tau), c_i \Phi^{-1}(\tau)^3]'$, $f_i(\tau) = [f_{1t}, f_{2t}, f_{3t}]'$ for $\tau \neq 0.5$ and $\lambda_i(\tau) = \alpha_i$, $f_i(\tau) = f_{1t}$ for $\tau = 0.5$. In particular, since the mapping $\tau \mapsto \Phi^{-1}(\tau)^3$ is strictly increasing, there exists a QFM representation as in equation (5) with $\lambda_i(\tau) = [\alpha_i, \Phi^{-1}(\tau)]'$ and $f_i(\tau) = [f_{1t}, f_{2t} + f_{3t} \Phi^{-1}(\tau)^2]'$ for $\tau \neq 0.5$, so that the second factor in $f_i(\tau)$ 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 notation, we suppress hereafter the dependence of $f_i(\tau)$, $\lambda_i(\tau)$, $r(\tau)$ and $u_{it}(\tau)$ on τ , so that the QFM in equation (5) is rewritten as

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

where $\lambda_i, f_i \in \mathbb{R}^r$. Let $\{f_{0t}\}$ and $\{\lambda_{0i}\}$ be the true values of $\{f_i\}$ 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 the first subsection, we consider the estimation of $\{\lambda_{0i}\}$ and $\{f_{0t}\}$ while r is assumed to be known; the estimation of r at each quantile is discussed in the second subsection.

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 identified separately without imposing normalizations (see Bai and Ng 2002). Without loss of generality, the following normalizations are imposed:

$$\frac{1}{T} \sum_{t=1}^T f_t f_t' = \mathbb{I}_r, \quad \frac{1}{N} \sum_{i=1}^N \lambda_i \lambda_i' \text{ is diagonal with non-increasing diagonal elements.} \quad (6)$$

Let $M = (N + T)r$ and $\theta = [\lambda_1', \dots, \lambda_N', f_1', \dots, f_T']'$, and let $\theta_0 = [\lambda_{01}', \dots, \lambda_{0N}', f_{01}', \dots, f_{0T}]'$ denote 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

$$\Theta^r = \{ \theta \in \mathbb{R}^M : \lambda_i \in \mathcal{A}, f_t \in \mathcal{F} \text{ for all } i, t, \text{ and } \{\lambda_i\}, \{f_t\} \text{ satisfy the normalizations in (6)} \}.$$

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

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

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 for the Chen *et al.* (2021)

computational algorithm, labelled *iterative quantile regression* (IQR), that can effectively find the stationary points of the object function, described in the first subsection of the Appendix.

Likewise, the asymptotic properties of the QFA estimators are not presented here, for brevity, but can be reviewed directly in Chen *et al.* (2021). In any case, it should be highlighted 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, Chen *et al.* (2021) propose a rank-minimization criterion to select the correct number at each τ with probability approaching 1. 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 let \mathcal{A}^k and \mathcal{F}^k be compact subsets of \mathbb{R}^k . In particular, let us assume that $[\lambda'_{0i} \ \mathbf{0}_{1 \times (k-r)}] \in \mathcal{A}^k$ for all i . Then let $\lambda^k_i, f^k_i \in \mathbb{R}^k$ for all i, t , and write $\theta^k = [\lambda^k_1, \dots, \lambda^k_N, f^k_1, \dots, f^k_T]'$, $\Lambda^k = [\lambda^k_1, \dots, \lambda^k_N]'$, $F^k = [f^k_1, \dots, f^k_T]'$. Consider the normalizations for factors and loadings discussed above; define $\hat{\Lambda}^k = [\hat{\lambda}^k_1, \dots, \hat{\lambda}^k_N]'$, and write

$$(\hat{\Lambda}^k)' \hat{\Lambda}^k / N = \text{diag}(\hat{\sigma}^k_{N,1}, \dots, \hat{\sigma}^k_{N,k}).$$

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

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

where P_{NT} is a sequence that goes to 0 as $N, T \rightarrow \infty$. In other words, \hat{r}_{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 cut-off value determining that asymptotic rank. In particular, Chen *et al.* (2021) find that the following choice of P_{NT} works well in practice:

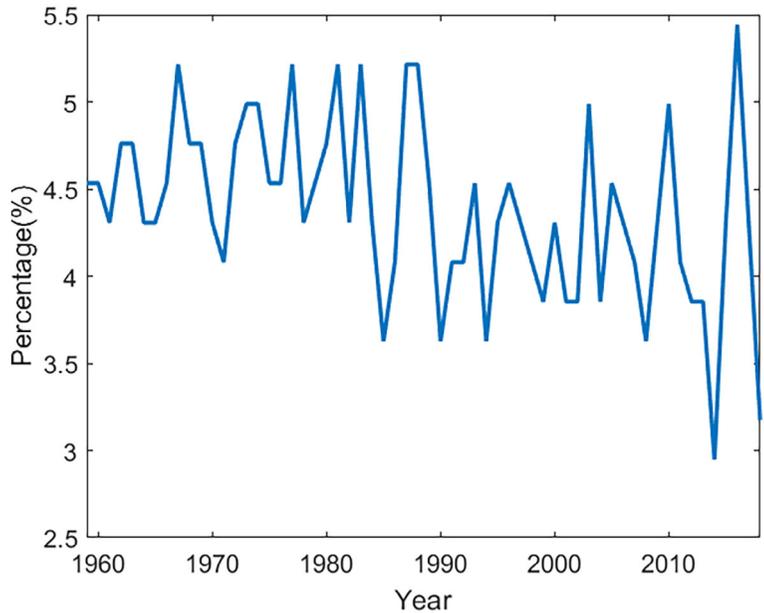
$$P_{NT} = \hat{\sigma}^k_{N,1} \left(\frac{1}{L^2_{NT}} \right)^{1/3}.$$

3.3 | Relative performance of PCA and QFA in a DGP with outliers

Two of the main characteristics of climate change are 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 3).¹⁰ Therefore this subsection studies how robust our novel proposed methodology is 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 other popular estimators used in the AFM literature, such as the information criteria (IC)-based method of Bai and Ng (2002) and the eigenvalue ratio (ER) estimator proposed by Ahn and Horenstein (2013). In what follows, we confirm these two claims by means of a few Monte Carlo simulations.

FIGURE 3 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 $f_{1t} = 0.8f_{1,t-1} + \varepsilon_{1t}$, $f_{2t} = 0.5f_{2,t-1} + \varepsilon_{2t}$ and $f_{3t} = 0.2f_{3,t-1} + \varepsilon_{3t}$, the $\lambda_{ji}, \varepsilon_{jt}$ 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 mean 0.98, and $\text{Cauchy}(0, 1)$ 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 the principal components PC_{p1} and information criteria IC_{p1} classes of Bai and Ng (2002), the ER estimator of Ahn and Horenstein (2013), and the above-mentioned Chen *et al.* (2021) rank-minimization estimator with P_{NT} chosen as in the previous subsection. We set $k = 8$ for all four estimators, and consider $N, T \in \{50, 100, 200, 500\}$.

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

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

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

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 find the PCA estimator (denoted \hat{F}_{PCA}), and then obtain the QFA estimator at $\tau = 0.5$ (denoted $\hat{F}_{QFA}^{0.5}$) using the IQR algorithm. Next, each of the true factors is regressed on \hat{F}_{PCA} and $\hat{F}_{QFA}^{0.5}$ separately, and their average (adjusted) R^2 values 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

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

N	T	$PC_{\hat{\rho}_1}$ of BN	$IC_{\hat{\rho}_1}$ of BN	Eigenvalue ratio	Rank estimator
50	50	[0.00]	[0.00]	[0.26]	[0.47]
50	100	[0.00]	[0.00]	[0.33]	[0.40]
50	200	[0.00]	[0.00]	[0.41]	[0.33]
50	500	[0.00]	[0.00]	[0.56]	[0.29]
100	50	[0.00]	[0.00]	[0.34]	[0.39]
100	100	[0.00]	[0.00]	[0.41]	[0.10]
100	200	[0.00]	[0.00]	[0.48]	[0.06]
100	500	[0.00]	[0.00]	[0.65]	[0.02]
200	50	[0.00]	[0.00]	[0.45]	[0.37]
200	100	[0.00]	[0.00]	[0.48]	[0.10]
200	200	[0.00]	[0.00]	[0.63]	[0.00]
200	500	[0.00]	[0.00]	[0.76]	[0.00]
500	50	[0.00]	[0.00]	[0.57]	[0.36]
500	100	[0.00]	[0.00]	[0.68]	[0.05]
500	200	[0.00]	[0.00]	[0.76]	[0.00]
500	500	[0.00]	[0.00]	[0.80]	[0.00]

Notes: N stands for Bai and Ng (2002).

The DGP considered in this table is $X_{it} = \sum_{j=1}^3 \lambda_{jt} f_{jt} + u_{it}$, where $f_{1T} = 0.8f_{1,t-1} + \varepsilon_{1t}$, $f_{2T} = 0.5f_{2,t-1} + \varepsilon_{2t}$, $f_{3T} = 0.2f_{3,t-1} + \varepsilon_{3t}$, $\lambda_{jt}, \varepsilon_{jt} \sim \text{i.i.d. } \mathcal{N}(0, 1)$, and $u_{it} \sim \text{i.i.d. } B_{it} \cdot \mathcal{N}(0, 1) + (1 - B_{it}) \cdot \text{Cauchy}(0, 1)$, where $B_{it} \sim \text{i.i.d. Bernoulli}(0.98)$. For each estimation method,

[proportion of $\hat{\tau} < 3$, proportion of $\hat{\tau} = 3$, proportion of $\hat{\tau} > 3$] are reported from 1000 replications.

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

N	T	Regress F on \hat{F}_{PCA}			Regress F on $\hat{F}_{QFA}^{0.5}$		
		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

Notes: The DGP considered in this table is as for Table 1. 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).

evidence in favour of using QFA instead of PCA in those cases where the idiosyncratic error terms in AFM exhibit heavy tails and outliers.

4 | CLIMATE CHANGE AND CO₂ CONCENTRATIONS

4.1 | Data description

For the empirical analysis, we use data from the Climatic Research Unit (CRU) at the University of East Anglia. In principle, the CRU provides monthly and annual data of land and sea temperatures in the northern and southern 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 USA, 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 are 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).

Data on CO₂ are obtained from the Global Carbon Budget (GCB) series.¹¹ The specific information on atmospheric CO₂ 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. The same series has been used recently by Bennedsen *et al.* (2023) in the estimation of a multivariate dynamic model involving the main variables included in the GCB; Figure 4 displays this series, whereas its corresponding first differences were plotted in Figure 2 in the Introduction. For the

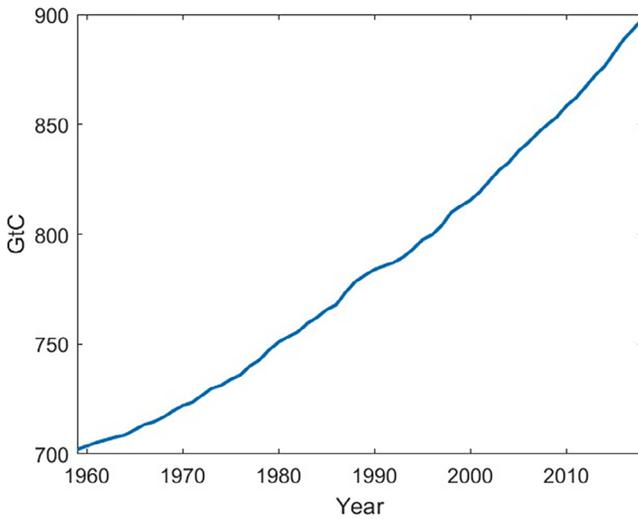


FIGURE 4 Atmospheric CO₂ concentrations (GCB).

early period 1959–80, estimations are based on Mauna Loa and South Pole stations as observed by the CO₂ programme at Scripps Institution of Oceanography; from 1980 onwards, they correspond to global averages estimated from multiple stations run by the National Oceanic and Atmospheric Administration (NOAA) and Earth System Research Laboratory (ESRL).

In addition, our choice of sample period is also determined by the characteristics of the data on CO₂. In fact, before 1958, CO₂ 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. Other papers analysing the relationship between temperature and CO₂ over the same time window include Pretis (2020) and Bennedsen *et al.* (2022).

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

$$C_t = \alpha + \beta t + v_t, \quad t = 1, \dots, T, \quad (7)$$

where C_t denotes a particular distributional characteristic of interest (e.g. a given quantile). The asymptotic properties of the OLS estimator in equation (7) 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 (7) becomes

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

where

$$\sum t C_t = T^{2+k} \frac{1}{T} \sum \left(\frac{t}{T} \frac{C_t}{T^k} \right) \quad \text{and} \quad \sum t^2 = T^3 \frac{1}{T} \sum \left(\frac{t}{T} \right)^2,$$

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 $t_{\beta=0} \rightarrow N(0, 1)$, so that the linear test based on equation (7) can detect any type of trend even if it is non-linear. Notice that we have assumed v_t to be $I(0)$, but nothing changes when it is *a priori* unknown whether v_t is $I(0)$ or $I(1)$. In effect, Perron and Yabu (2009) propose an estimation procedure based on feasible quasi-generalized least squares for this agnostic case, and show that the limiting distribution of a similar t -statistic for $\beta = 0$ remains $N(0, 1)$, the intuition being that the divergence rate of a deterministic trend is faster than that of a stochastic trend. In any case, our temperature data—either by stations or by characteristics—do not contain a unit root, as discussed in the second subsection of the Appendix.

The 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, standard deviation, quantiles, etc.) are estimated for the sample period 1959–2018. Figure 5 presents 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 interquantile range, 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.

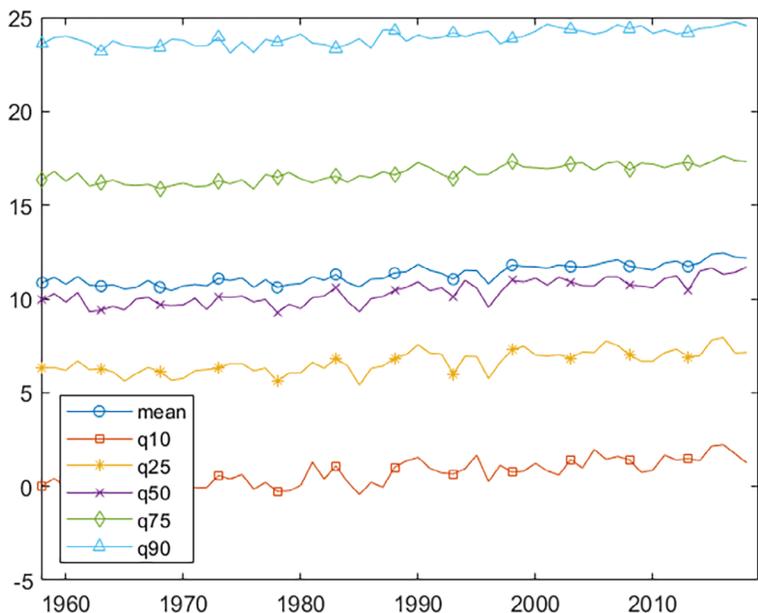


FIGURE 5 Estimated distributional characteristics.

TABLE 3 Gadea–Gonzalo Trend Test (1959–2018)

Characteristic	Test statistic	<i>p</i> -value
Mean	0.0264	0.0000
S.D.	−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

Notes: Here, 'kur' refers to kurtosis, and 'skw' to skewness. Annual distributional characteristics are estimated using the cross-sectional distribution at each year (1959–2018). OLS estimates and heteroscedasticity- and autocorrelation-consistent $t_{\beta=0}$ *p*-values from regression (7) are reported.

4.3 | Quantile factor analysis

The previous evidence opens the door to analyse heterogeneous association patterns between local (station-level) temperatures and atmospheric CO₂ 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 CO₂ concentrations.

A key requirement for extracting the QFA (and the PCA) factors is the absence of stochastic and deterministic trends in the individual temperature processes. The second subsection of the Appendix reports the results of applying standard augmented Dickey–Fuller (ADF) tests in regressions with a linear trend component. As was mentioned before, 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 the second subsection of Section 2 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 $\tau = 0.01, 0.05, 0.10, 0.95, 0.90, 0.99$), 3 (at $\tau = 0.25, 0.75$) and 4 (at $\tau = 0.5$). For illustrative purposes, Figure 6 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 is chosen according to the PC_{p1} criterion of Bai and Ng (2002), which selects 8.

FIGURE 6 Estimated QFA factors.

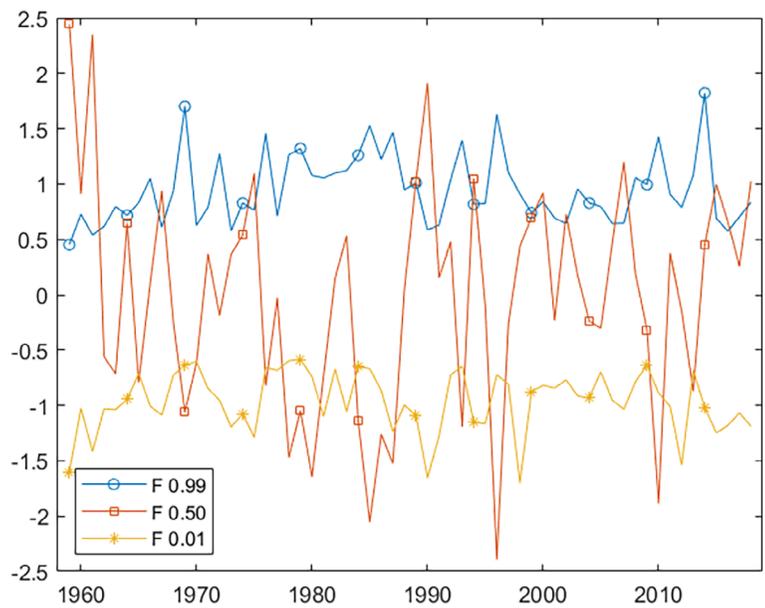


TABLE 4 Comparison of \hat{F}_{QFA}^τ and \hat{F}_{PCA}

τ	Elements of \hat{F}_{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			

Notes: This table reports the R^2 of regressing each element of \hat{F}_{QFA}^τ on \hat{F}_{PCA} . For \hat{F}_{QFA}^τ , the number of estimated factors is obtained using the rank-minimization criterion, while for \hat{F}_{PCA} , the number of estimated factors is equal to 8 in all datasets.

This is the maximum number imposed in the IQR computational algorithm described in the first subsection of the Appendix.

To compare the QFA factors (denoted as \hat{F}_{QFA}^τ) with the PCA factors (denoted as \hat{F}_{PCA}), we regress each element of \hat{F}_{QFA}^τ on the eight \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 centre of the distribution ($\tau = 0.25, 0.5, 0.75$), the estimated factors are highly correlated with the PCA factors, with all the R^2 values 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. $\tau = 0.01, 0.05, 0.95, 0.99$) exhibit much lower correlations with the PCA factors, with R^2 values 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 to identify different features of the temperature distribution that the factors at the medium quantiles are unable to capture.

4.4 | Temperature factors and CO₂ concentrations

Bivariate tests characterizing the predictive association between the estimated QFA temperature factors and a suitable transformation of the atmospheric CO₂ concentrations are implemented next. The statistical properties of the latter determine how it should be transformed. The plot in Figure 4 makes it clear that the CO₂ 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 Figure 2 points to an acceleration over the sample period following a linear trend. These features are confirmed by an ADF test reported in the second subsection of the Appendix, which suggest that the level of CO₂ concentrations has a unit root, while the growth rate is trend-stationary. As discussed in Bennedsen *et al.* (2023), these properties are consistent with a dynamic statistical model relating CO₂ concentrations, anthropogenic CO₂ emissions, and absorption of CO₂ by the terrestrial, ocean and marine biospheres. The same statistical properties hold for the (logged) CO₂ concentrations from Mauna Loa and the CO₂ radiative forcing, series that we use in the robustness exercises undertaken in the third subsection of the Appendix.

Building on the previous observations, the tests to be implemented rely on separate OLS regressions of the estimated QFA factors at each relevant quantile ($\hat{F}_{QFA,t}^\tau$) on p own lags and q lags of the linearly detrended changes in atmospheric CO₂ concentrations ($\hat{\Delta}CO_2$). In other words, this approach sheds light on whether past fluctuations in CO₂ concentration changes around a linear trend have a predictive power on the common quantile-dependent factors of the station-level temperature fluctuations, again around a linear trend. The specific regressions are given by

$$\hat{F}_{QFA,t}^\tau = \alpha + \sum_{i=1}^p \beta_i \times \hat{F}_{QFA,t-i}^\tau + \sum_{j=1}^q \gamma_j \times \hat{\Delta}CO_{2,t-j} + u_t, \quad (8)$$

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 condition similar to that used in Bai and Ng (2006) is required to replace the true quantile factors by the QFA estimated ones in equation (8), namely, $\sqrt{T}/N \rightarrow 0$. 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, $j = 1, \dots, q$, 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 CO₂ concentrations have a predictive power on the current values of QFA factors of the temperature at the lower and middle quantiles (τ from 0.01 to 0.75). Yet this is not the case for the QFA factors at the extreme upper quantiles ($\tau = 0.90, 0.95, 0.99$), 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 the third subsection of the Appendix shows—happens to be robust to alternative sources, units of measurement, and transformations of the CO₂ concentrations series.

As a complementary exercise, we study the predictive association between the eight PCA mean factors selected with the PC_{p1} criterion, and the suitable transformation of the CO₂ concentrations. Notice that under the standard factor model conditions, the average temperature across stations is equivalent to a linear combination of these eight 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 5 p -values of the Predictive Association Tests for the QFA Factors and transformed CO₂

Regressor	Elements of \hat{F}_{QFA}^{τ}				
	τ	(1)	(2)	(3)	(4)
$\hat{\Delta}CO_2$	0.01	0.0713			
	0.05	0.0317			
	0.10	0.0546			
	0.25	0.0140	0.5626	0.7714	
	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			

Notes: This table reports the p -values of the proposed F -test for the joint significance of the coefficients $\gamma_j, j = 1, \dots, q$, in equation (8). Lag lengths are chosen following a general-to-specific approach. p -values smaller than 0.1 are in bold.

TABLE 6 p -values of the Predictive Association Tests for the PCA Factors

Regressor	Elements of \hat{F}_{PCA}							
	(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

Notes: See Table 5.

5 | DISCUSSION

The physical link between temperature and CO₂ concentrations dates back from the late 1800s (see Arrhenius 1896) in the climate science literature. Our analysis of climate sensitivity departs from the standard approach that is typically quantified as warming per doubling of CO₂ (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 CO₂ on temperature, where the former is measured in (detrended) levels, and the latter as (detrended) changes in global concentrations. In this fashion, our approach complements other research on heterogeneous climate sensitivity, such as Shindell and Faluvegi (2009), where the sensitivity of regional climate to changes in CO₂ (spatial heterogeneity) is investigated. Moreover, this non-uniformity seems to be a finding associated with only CO₂. In effect, when we consider the radiative forcing of other greenhouse gases (see the final subsection of the Appendix), such as methane, ozone and nitrous oxide, as well as natural forcings, such as solar irradiance and volcanic activity, we find predictive association for all quantile factors rather than for specific ones associated with the lower and central parts of the temperature distribution, as with CO₂.

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 CO₂ concentrations will have unforeseen consequences (i.e. whatever happens at 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 (CO₂ and methane) from the permafrost melting. In this respect, we highlight that non-uniform climate sensitivity is not concentrated regionally but rather affects all the regions around the globe. In particular, the growth rate of CO₂ emissions predicts (positively) the periods where the temperature decreases or does not increase much.

In view of this evidence, further research should aim at analysing jointly the heterogeneous predictive power of CO₂ 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 those 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) CO₂ concentration changes and temperatures from 441 weather stations in the northern and southern hemispheres over the period 1959–2018. Using the QFA methodology proposed by Chen *et al.* (2021), we retrieve the quantile-dependent common factors and their number at different quantiles. We apply predictive association tests of different CO₂ 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 than standard PCA in the presence of outliers, as is the case in climate data.

Our main finding is that CO₂ concentration 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 of 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 being complementary to those available on climate sensitivity, as, for example, in Sherwood *et al.* (2020).

ACKNOWLEDGMENTS

This paper was 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.

NOTES

¹ When sunlight reaches the 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 with which they interact, CO₂ and other greenhouse gases absorb energy at a variety of wavelengths whose ranges do overlap that of infrared energy. As CO₂ 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 the Earth as heat, contributing to GW through the so-called 'greenhouse effect', first discovered by Fourier (1824), verified experimentally by Foote (1856) and Tyndall (1863), and quantified by Arrhenius (1896).

- 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 Note that a simple approach in that direction consisting of the estimation of a standard quantile regression for the mean temperature in each period, $Q_\tau[\bar{X}_t|Z_t] = \bar{\beta}_0(\tau) + \bar{\beta}_1(\tau)Z_t$, $0 < \tau < 1$, would still disregard all the information on the distribution of temperature across different stations.
- ⁴ Similarly, we also use the alternative notation $Q_X[\tau|Z]$ to denote the conditional quantile of X given Z .
- ⁵ Ando and Bai (2020) use a similar setup with an unobservable factor structure that 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.
- ⁶ Note that the simplifying assumption of a known number of factors in this specific example is relaxed later.
- ⁷ Note that since f_{1t} can be estimated consistently 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.
- ⁸ 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).
- ⁹ A complete theoretical as well as empirical trend analysis can be found in Gadea and Gonzalo (2020).
- ¹⁰ 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.
- ¹¹ The GCB series is compiled by Friedlingstein *et al.* (2022) and available at <https://www.icos-cp.eu/science-and-impact/global-carbon-budget/2021> (accessed 2 July 2023).
- ¹² Note that these properties of the dependent variables preclude the use of cointegration in a bivariate setup.
- ¹³ Recall that PC_{p1} is chosen to select the number of PCA factors estimated in these regressions to play conservative.
- 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.
- ¹⁵ By definition, $ERF_t^{CO_2} = 5.35 \times \ln(CO_{2t}/CO_{2base})$, where CO_{2t} are the CO_2 concentrations at a given year t , and CO_{2base} are the CO_2 concentrations at a given base year.

REFERENCES

- Ahn, S. C. and Horenstein, A. R. (2013). Eigenvalue ratio test for the number of factors. *Econometrica*, **81**(3), 1203–27.
- Ando, T. and Bai, J. (2020). Quantile co-movement in financial markets: a panel quantile model with unobserved heterogeneity. *Journal of the American Statistical Association*, **115**(529), 266–79.
- 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), 237–76.
- Bai, J. (2003). Inferential theory for factor models of large dimensions. *Econometrica*, **71**(1), 135–71.
- and Ng, S. (2002). Determining the number of factors in approximate factor models. *Econometrica*, **70**(1), 191–221.
- and ——— (2006). Confidence intervals for diffusion index forecasts and inference for factor-augmented regressions. *Econometrica*, **74**(4), 1133–50.
- Bennedsen, M., Hillebrand, E. and Koopman, S. J. (2023). A multivariate dynamic statistical model of the global carbon budget 1959–2020. *Journal of the Royal Statistical Society Series A: Statistics in Society*, **86**(1), 20–42.
- , ——— and Zhou-Lykke, J. (2022). Global temperature projections from a statistical energy balance model using multiple sources of historical data. arXiv:2205.10269.
- Castle, J. and Hendry, D. (2020). Climate econometrics: an overview. *Foundations and Trends in Econometrics*, **10**(3–4), 145–322.
- Chang, Y., Kaufmann, R., Sik-Kim, C., Miller, I., Park, Y. and Park, S. (2020). Evaluating trends in time series of distributions: a spatial fingerprint of human effects on climate. *Journal of Econometrics*, **214**(1), 274–94.
- Chapman, S., Stainforth, D. and Watkins, N. (2013). On estimating local long-term climate trends. *Philosophical Transactions of the Royal Society A*, **371**(1991), 20120287.
- Chen, L., Dolado, J. and Gonzalo, J. (2021). Quantile factor models. *Econometrica*, **89**(2), 875–910.
- , Gao, J. and Vahid, F. (2022). Global temperatures and greenhouse gases: a common features approach. *Journal of Econometrics*, **230**, 240–54.
- Davy, T., Esau, I., Chernokulsky, A., Outten, S. and Zilitinkevich, S. (2017). Diurnal asymmetry to the observed global warming. *International Journal of Climatology*, **37**(1), 79–93.
- Dlugokencky, E. and Tans, P. (2020). Trends in atmospheric carbon dioxide. Technical Report, National Oceanic and Atmospheric Administration, Earth System Research Laboratory.
- Estrada, F., Perron, P. and Martínez-López, B. (2013). Statistically derived contributions of diverse human influences to twentieth-century temperature changes. *Nature Geoscience*, **6**(12), 1050–5.

- Foote, E. (1856). Circumstances affecting the heat of the sun's rays. *American Journal of Science and Arts*, **22**(2), 382–3.
- Fourier, J. (1824). On the temperatures of the terrestrial sphere and interplanetary space. *Chimie et de Physique*, **27**, 136–67.
- Friedlingstein, P., Jones, M., *et al.* (2022). Global carbon budget 2021. *Earth System Science Data*, **14**(4), 1917–2005.
- Gadea, M. and Gonzalo, J. (2020). Trends in distributional characteristics: existence of global warming. *Journal of Econometrics*, **214**(1), 153–74.
- and ——— (2023). Climate change heterogeneity: a new quantitative approach. arXiv:2301.02648.
- Gao, J. and Hawthorne, K. (2006). Semiparametric estimation and testing of the trend of temperature series. *Journal of Econometrics*, **9**(2), 332–55.
- Gay-Garcia, C., Estrada, F. and Sánchez, A. (2009). Global and hemispheric temperatures revisited. *Climatic Change*, **94**(3), 333–49.
- Golub, G. H. and Van Loan, C. F. (2013). *Matrix Computations: 3*. Baltimore, MD: Johns Hopkins University Press.
- Hansen, J., Sato, M., Kharecha, P. and von Schuckmann, K. (2011). Earth's energy imbalance and implications. *Atmospheric Chemistry and Physics*, **11**, 13421–49.
- Huber, P. (1981). *Robust Statistics*. London: John Wiley and Sons.
- Ji, F., Wu, Z., Huang, J. and Chassignet, E. (2014). Evolution of land surface air temperature trend. *Nature Climate Change*, **4**, 462–6.
- Kaufmann, R., Kauppi, K. and Stock, J. (2006). Emissions, concentrations, and temperature: a time series analysis. *Climatic Change*, **77**, 249–78.
- Koenker, R. (2005). *Quantile Regression*. Cambridge: Cambridge University Press.
- Leduc, M., Matthews, H. D. and de Elía, R. (2016). Regional estimates of the transient climate response to cumulative CO₂ emissions. *Nature Climate Change*, **6**, 474–8.
- Miller, I. (2023). Local climate sensitivity: what can time series of distributions reveal about spatial heterogeneity of climate change? *Advances in Econometrics*, **45B**, 319–50.
- Montamat, G. and Stock, J. (2020). Quasi-experimental estimates of the transient climate response using observational data. *Climatic Change*, **160**, 361–71.
- Perron, P. and Yabu, T. (2009). Estimating deterministic trends with an integrated or stationary noise component. *Journal of Econometrics*, **151**, 56–69.
- Phillips, P., Leirvik, T. and Storelvmo, T. (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), 256–73.
- and Hendry, D. (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**, 375–84.
- Previdi, M., Smith, K. and Polvani, L. (2021). Arctic amplification of climate change: a review of underlying mechanisms. *Environmental Research Letters*, **16**(9), 1–26.
- Sherwood, S., Webb, M., Annan, J., Armour, K., Forster, P., Hargreaves, J., Hegerl, G., Klein, K., Marvel, S., Rohling, E., Watanabe, M., Andrews, T., Braconnot, P., Bretherton, C., Foster, G., Hausfather, Z., von der Heydt, A., Knutti, R., Mauritsen, T., Norris, J., Proistosescu, C., Rugenstein, M., Schmidt, G., Tokarska, K. and Zelinka, M. (2020). An assessment of Earth's climate sensitivity using multiple lines of evidence. *Reviews of Geophysics*, **58**(4), e2019RG000678.
- Shindell, D. (2014). Inhomogeneous forcing and transient climate sensitivity. *Nature Climate Change*, **4**, 274–7.
- and Faluvegi, G. (2009). Climate response to regional radiative forcing during the twentieth century. *Nature Climate Change*, **2**, 294–300.
- Stips, A., Macias, D., Coughlan, C., Garcia-Gorritz, E. and Liang, X. (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), 200–6.

How to cite this article: Chen, L., Dolado, J. J., Gonzalo, J. and Ramos, A. (2023). Heterogeneous predictive association of CO₂ with global warming. *Economica*, 1–25. <https://doi.org/10.1111/ecca.12491>

APPENDIX

IQR algorithm

Using the notation in the first subsection of Section 2, let $\Lambda = [\lambda_1, \dots, \lambda_N]'$, $F = [f_1, \dots, f_T]'$, and define the averages

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

Note that we have

$$\mathbb{M}_{NT}(\theta) = N^{-1} \sum_{i=1}^N \mathbb{M}_{i,T}(\lambda_i, F) = T^{-1} \sum_{i=1}^T \mathbb{M}_{t,N}(\Lambda, f_i).$$

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 solved efficiently by various linear programming methods (see Koenker 2005, ch. 6). Based on this observation, the following iterative procedure is proposed.

Iterative quantile regression (IQR)

1. Choose random starting parameters $F^{(0)}$.
2. Given $F^{(l-1)}$, solve $\lambda_i^{(l-1)} = \arg \min_{\lambda} \mathbb{M}_{i,T}(\lambda, F^{(l-1)})$ for $i = 1, \dots, N$. Given $\Lambda^{(l-1)}$, solve $f_t^{(l)} = \arg \min_f \mathbb{M}_{t,N}(\Lambda^{(l-1)}, f)$ for $t = 1, \dots, T$.
3. For $l = 1, \dots, L$, iterate step 2 until $\mathbb{M}_{NT}(\theta^{(L)})$ is close to $\mathbb{M}_{NT}(\theta^{(L-1)})$, where $\theta^{(l)} = (\text{vech}(\Lambda^{(l)})', \text{vech}(F^{(l)})')'$.
4. Normalize $\Lambda^{(L)}$ and $F^{(L)}$ so that they satisfy the normalizations in (6).

In the general case where $r \geq 1$, if we replace the check function in the IQR algorithm by the least squares loss function and normalize $F^{(l-1)}, \Lambda^{(l-1)}$ to satisfy (6) at step 2, then 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' .

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 A1, the null hypothesis of unit root is rejected in 98.66% of the stations; it is not rejected in only 5 stations out of 441. In a similar direction, the ADF test suggests no unit roots on the different distributional characteristics. This piece of evidence suggests that temperature series do not contain unit roots, but rather follow trend-stationary processes.

TABLE A1 ADF Unit Root Tests for the Temperature Related Series

ADF test by stations		
Percentage of rejections	98.66%	
Number of non-rejections	5	
ADF test by characteristics		
Characteristic	Test statistic	<i>p</i> -value
Mean	-5.6261	0.0001
S.D.	-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

Notes: Here, 'kur' refers to kurtosis, and 'skw' to skewness. Annual distributional characteristics estimated using the cross-sectional distribution at each year (1959–2018). Significance level 5% is considered in the individual tests. ADF test equations include intercept and trend. Lag selection conducted using the structural Bayesian information criterion.

TABLE A2 ADF Unit Root Tests CO₂ and ERF Related Variables

Variable	Level series		First differences	
	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

Notes: This table reports the test statistic of the ADF test on the corresponding variable when the test includes only a constant or a constant and an intercept. Lags are selected using the Bayesian information criterion. The values in bold indicate that the null of the test is rejected.

In Table A2, the *p*-values of the ADF test implemented on the series of CO₂ and effective radiative forcing (ERF) are reported. The three series related to CO₂ ($CO_{2,GCB}$, $\log(CO_{2,MLO})$ and ERF^{CO_2}) contain a unit root in their levels, while the first differences are trend-stationary. Regarding the ERF series, the test results indicates that the non-CO₂ 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).

Robustness

The main finding of the paper is the heterogeneous predictive association between a suitable transformation of the atmospheric CO₂ 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 CO₂ concentrations as regressors. The first series is the mean annual CO₂ concentrations series measured in parts per million by volume as obtained from direct measurements at Mauna Loa ($CO_{2,MLO}$) and available at <https://gml.noaa.gov/ccgg/trends/data.html> (accessed 2 July 2023). The second series corresponds to the ERF from the $ERF_t^{CO_2}$ series,¹⁵ measured in watts per square metre ($W m^{-2}$) and obtained from Hansen *et al.* (2011) (available at <http://www.columbia.edu/~mhs119/Forcings>, accessed 2 July 2023). From Figures A1 and A2, it is clear that the dynamics of both series are similar to the dynamics of the CO₂ concentrations from the GCB.

Provided that the series in Figures A1(b) and A2(b) are trend-stationary according to the ADF tests for unit roots described in the previous subsection, in equation (8) we consider as regressors the linearly detrended growth rate of CO₂ concentrations from Mauna Loa ($\hat{\Delta} \log(CO_{2,MLO})$) and the linearly detrended first differences of the ERF from CO₂ ($\hat{\Delta} ERF_t^{CO_2}$). Results reported in Table A3 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

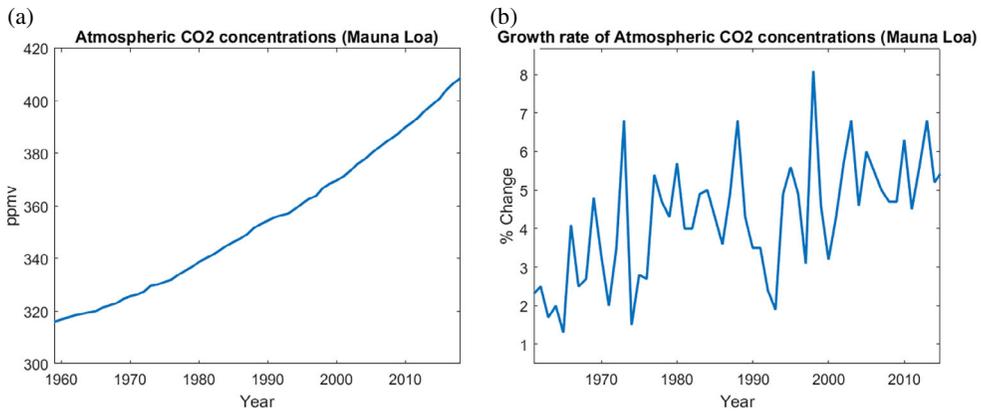


FIGURE A1 Atmospheric CO₂ concentrations (Mauna Loa).

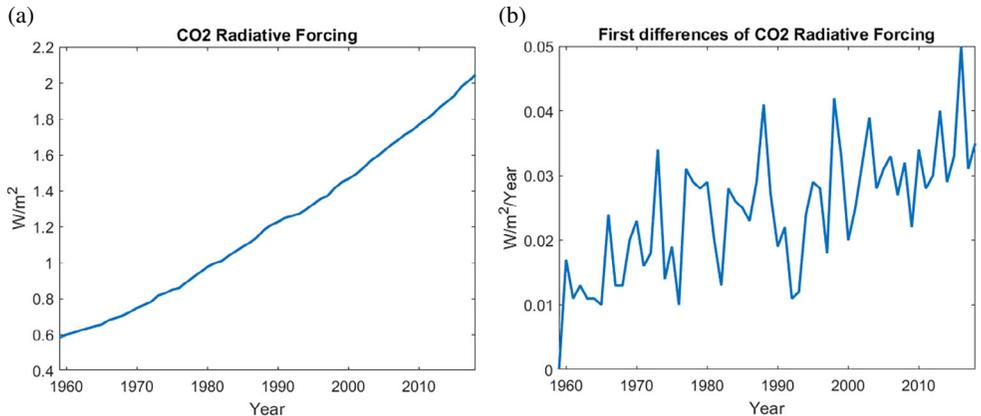
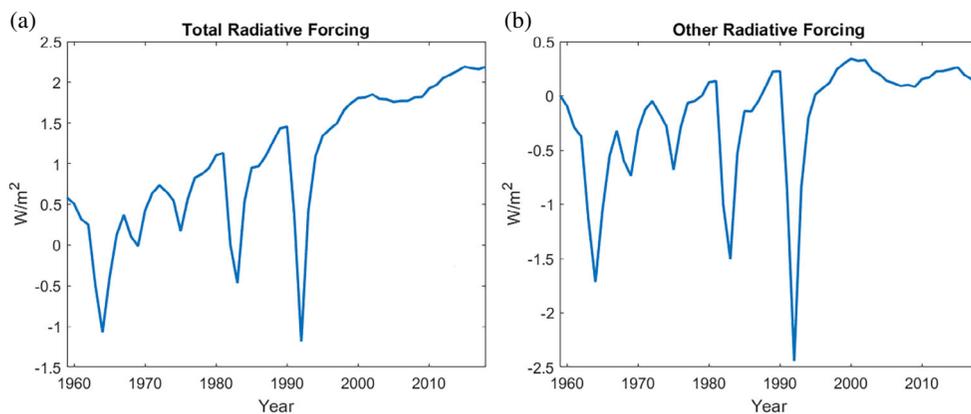


FIGURE A2 Effective radiative forcing from CO₂.

TABLE A3 p -values of the Predictive Association Tests for the QFA Factors and Transformed CO_2 (Robustness)

Regressor	Elements of \hat{F}_{QFA}^{τ}				
	τ	(1)	(2)	(3)	(4)
$\hat{\Delta} \log(CO_{2,MLO})$	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			

Notes: This table reports the p -values of the proposed F -test for the joint significance of the coefficients $\gamma_j, j = 1, \dots, q$, in the corresponding version of equation (8) for $\hat{\Delta} \log(CO_{2,MLO})$ 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.

**FIGURE A3** Effective radiative forcing from all sources.

middle quantiles (τ from 0.01 to 0.75), while the null of the test is not rejected at upper quantiles ($\tau = 0.90, 0.95, 0.99$).

Other warming sources

Even though the interest of the paper is on the bivariate association between CO_2 and station-level temperatures, in this subsection we briefly examine the association with other greenhouse gases such as methane, ozone and nitrous oxide, as well as natural forcings such as solar irradiance and volcanic activity such as the Fuego (Guatemala), El Chinchón (Mexico) and Pinatubo (Philippines) eruptions that explain the big dips. In Figure A3, we present the total ERF (ERF^{Tot})

TABLE A4 p -values of the Predictive Association Tests for the QFA Factors and Other Warming Sources

Regressor	Elements of \hat{F}_{QFA}^τ				
	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			

Notes: This table reports the p -values of the proposed F -test for the joint significance of the coefficients $\gamma_j, j = 1, \dots, q$, in the corresponding version of equation (8) 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.

and the ERF from sources other than CO_2 (ERF^{nonCO_2}) as taken from Hansen *et al.* (2011). The ADF tests for unit root described in the second subsection of this Appendix indicate that ERF^{Tot} is $I(1)$ and ERF^{nonCO_2} is $I(0)$.

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 A4, 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.