

NO. 1031
SEPTEMBER 2022

800,000 Years of Climate Risk

Tobias Adrian | Nina Boyarchenko | Domenico Giannone |
Ananthakrishnan Prasad | Dulani Seneviratne | Yanzhe Xiao

800,000 Years of Climate Risk

Tobias Adrian, Nina Boyarchenko, Domenico Giannone, Ananthakrishnan Prasad, Dulani Seneviratne, and Yanzhe Xiao

Federal Reserve Bank of New York Staff Reports, no. 1031

September 2022

JEL classification: Q54, C14, C53

Abstract

We use a long history of global temperature and atmospheric carbon dioxide (CO₂) concentration to estimate the conditional joint evolution of temperature and CO₂ at a millennial frequency. We document three basic facts. First, the temperature–CO₂ dynamics are non-linear, so that large deviations in either temperature or CO₂ concentrations take a long time to correct—on the scale of multiple millennia. Second, the joint dynamics of temperature and CO₂ concentrations exhibit multimodality around historical turning points in temperature and concentration cycles, so that prior to the start of cooling periods, there is a noticeable probability that temperature and CO₂ concentrations may continue to increase. Finally, evaluating the future evolution of temperature and CO₂ concentration conditional on alternative scenarios realizing, we document that, even conditional on the net-zero 2050 scenario, there remains a significant risk of elevated temperatures for at least a further five millennia.

Key words: climate change, multimodality, NGFS scenarios

Boyarchenko: Federal Reserve Bank of New York, CEPR (email: nina.boyarchenko@ny.frb.org). Adrian: International Monetary Fund, CEPR (email: tadrian@imf.org). Prasad, Seneviratne, and Xiao: International Monetary Fund (emails: tadrian@imf.org, aprasad@imf.org, dseneviratne@imf.org, and yxiao3@imf.org). Giannone: University of Washington, Amazon.com, CEPR (email: dgiannon@uw.edu).

This paper presents preliminary findings and is being distributed to economists and other interested readers solely to stimulate discussion and elicit comments. The views expressed in this paper are those of the author(s) and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System, the International Monetary Fund or its management or its executive directors, or Amazon.com. This publication and its contents are not related to Amazon and do not reflect the position of the company and its subsidiaries. Any errors or omissions are the responsibility of the author(s).

To view the authors' disclosure statements, visit
https://www.newyorkfed.org/research/staff_reports/sr1031.html.

1 Introduction

How does climate evolve over long periods of time? Although so-called “tipping points” – where a small change in a climate fundamental can result in a disproportionate response to the overall system – have long been part of the scientific discussion of the potential future evolution of climate, research has struggled to assign probabilities to the precise nature of such tipping points, including the level at which tipping points in, for example, temperature may occur, how long it may take the overall climate to reach that tipping point, as well as the magnitude of the consequences should such a tipping point occur. In this paper, we take a non-parametric approach to modeling the joint conditional evolution of temperature and CO₂ (carbon dioxide) atmospheric concentrations over prolonged periods of time. This allows us to construct term structures of joint distributions for temperature and CO₂ at any point of the 800,000 year history of the available data on global mean temperature and atmospheric carbon concentrations, including providing forecasts of the future evolution of these two key climate variables.

More specifically, we apply the non-parametric approach developed in Adrian et al. (2021) to joint dynamics of the long time series of global temperature collected by Snyder (2016, 2019) and of global CO₂ concentration. The non-parametric approach allows us to remain agnostic about the dynamic relationship between CO₂ concentrations and temperature, allowing the data to inform us on the dynamic properties of global climate. Acknowledging the inherent uncertainties in using climate data spanning 800,000 years, our estimation uses rolling windows of 140,000 years at a time. This approach allows us to form conditional forecasts as of present using the “best” quality climate data (as there is less measurement error in temperature and carbon dioxide readings over this period), as well as implicitly allowing for time-variation in the interactions between CO₂ and temperature. For each point in time, we use the prior 140,000 years to construct out-of-sample estimates of the predicted evolution of the conditional joint distribution of CO₂ and temperature, and show that these

out-of-sample predicted distributions are well calibrated.

The estimated dynamics are startling. First, we document that the joint dynamics of temperature and CO₂ concentrations are highly non-linear, so that large deviations in either temperature or CO₂ concentrations can take millennia to correct, even if no further impulses are provided to the temperature-CO₂ system. We emphasize that these dynamics are present in the longer history of climate readings, even before the noticeable changes to the dynamics of CO₂ atmospheric concentrations follow the Industrial Revolution that have been documented in prior literature (see e.g. Kaufmann and Stock, 2006).

Second, these non-linear joint dynamics of temperature and CO₂ concentrations are exacerbated by periods of multimodality in the joint conditional distribution. We show that around historical turning points – such as the end of the Eemian period and the current period – the global climate can resolve to either a cooling of the planet and a reduction in atmospheric CO₂ concentration or to a continual warming and build-up of CO₂ concentrations. That is, while the observed glacial period that defines the end of Eemian period was not a tail outcome for the climate dynamics, neither would have continuing increases in average global temperatures been. Examining the initial conditions that correspond to the emergence of multimodalities in the joint dynamics systematic, we find that multimodality emerges systematically when CO₂ concentrations exceed their historical seventieth percentile, regardless of whether global temperatures are extreme, so that high carbon dioxide concentrations in the atmosphere today correspond to the possibility of either continuing build-up of CO₂ in the atmosphere and an associated increase in global temperatures or a decline in CO₂ concentrations and a cooling of the planet. Furthermore, multimodality is more pronounced in longer-horizon projections, suggesting that the temperature – CO₂ dynamics are inherently unstable in the long run.

Finally, we use the joint dynamics of temperature and CO₂ concentrations estimated from the most recent 140,000 years of data to construct distributional forecasts of CO₂ and temperature for the next five millennia, conditional on alternative Network for Greening the

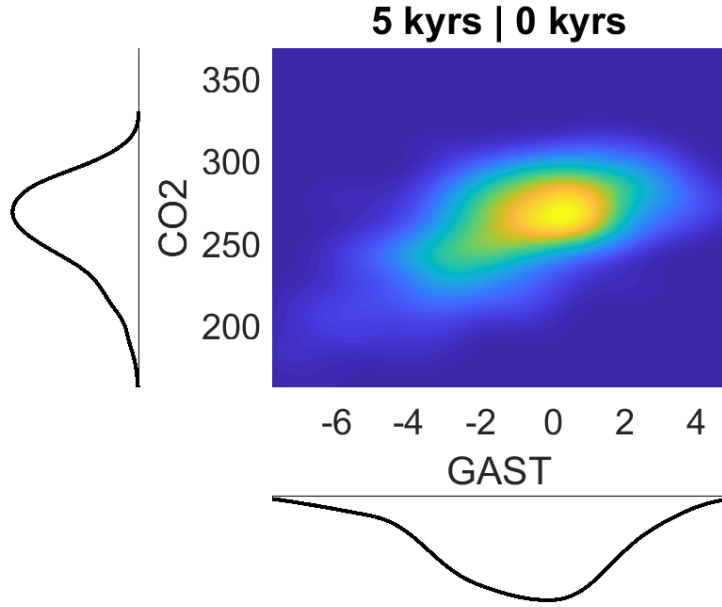


Figure 1. Five thousand years ahead forecast under the net-zero scenario. The figure shows the five thousand years ahead marginal and joint distributions of CO₂ concentrations and temperature anomalies conditional on the net-zero 2050 scenario.

Financial System (NGFS)¹ scenarios. Our estimates suggest that even under the so-called “net-zero 2050” scenario, in which global warming relative to pre-industrial averages is limited to 1.5 degrees Celsius, there is a substantial probability of global temperatures continuing to rise for the next two millennia, with positive probabilities assigned to even 4-5 degrees Celsius warming. Although our estimates predict some mean reversion in climate conditions in the medium run, Figure 1 shows a significant right tail in temperature outcomes as long as 5,000 years out, as well as emerging signs of multimodality. That is, even if the more stringent climate policies that are required for the net-zero 2050 scenario are implemented, it may take millennia for global warming to abate and the abatement may resolve into a new glacial period.

The remainder of the paper is organized as follows. Section 2 describes the data and empirical methodology. We focus on the historical non-linear relationship between CO₂

¹The NGFS consists 114 central banks and financial regulatory agencies that aim to accelerate the scaling up of green finance and develop recommendations for central banks’ role for climate change. The NGFS was created in 2017 and its secretariat is hosted by the Banque de France. See <https://www.ngfs.net/en>.

concentrations and temperature in Section 3. Section 4 then evaluates the NGFS scenarios within our framework. Section 5 reviews the literature and Section 6 concludes. We relegate the robustness exercises we conduct to the Online Appendix.

2 Data and Methodology

We are interested in constructing multi-period-ahead conditional joint distributions of global average surface temperature anomalies (GAST) and atmospheric CO₂ concentrations. In this section, we describe the source of our data and how the data is constructed. We also describe how we construct the one-period-ahead conditional joint distribution. We then use efficient Markov Chain Monte Carlo (MCMC) to construct multi-period-ahead distributions from the one-period-ahead distribution.

2.1 Data

To assess the non-linear relationship between emissions and temperature, we use global average surface temperature anomalies (GAST) and atmospheric CO₂ concentrations based on the variables constructed in Snyder (2016, 2019)² at the 1000-year frequency.³ The methodology for the global average surface temperature is described in Snyder (2016) and is constructed based on a spatially weighted proxy reconstruction process utilizing over 20,000 sea surface temperature point reconstructions. Specifically, the global average surface temperature is estimated in Snyder (2016) from proxy-based reconstructions of local Sea Surface Temperature (SST) available in a multi-proxy database. The SST database includes 61 SST proxy reconstructions from 59 ocean sediment cores: 29 using alkenone unsaturation indices, 17 using ratios of Mg/Ca in planktonic foraminifera, and 16 based on microfossil abundance. This multi-proxy approach reduces the potential biases that are unique to individual proxies. The reconstruction process uses probabilistic simulations across multiple sources

²Variables computed in Snyder (2016, 2019) are available [here](#).

³The latest observation of the sample is as of one thousand years BP.

of uncertainty to estimate credible intervals at 1000-year intervals. Global average surface temperature anomalies are constructed as the deviation from the average of 1000-5000 years ago. To put this into perspective, in radiocarbon dating, “present” typically refers to 1950 CE. Therefore, the 1st observation in the sample refers to the deviation in temperature in 950 CE from the average temperature between 1950 CE to 3050 BC. Similarly, the second observation in the sample refers to the deviation in temperature in 50 BC from the average temperature between 1950 CE to 3050 BC. Following this rationale, we can see the 6th observation for instance refers to the deviation in temperature in 4050 BC from the average temperature between 1950 CE to 3050 BC.

The methodology for the reconstruction of atmospheric CO₂ concentrations is described in Bereiter et al. (2015) based on the European Project for Ice Coring in Antarctica Dome ice core from Dome C (EDC) which reconstructed atmospheric CO₂ concentrations for the last 800,000 years. Bereiter et al. (2015) use different air extraction methods and find an analytical artifact, which increases over the deepest 200m and reaches 10.1 ± 2.4 ppm in the oldest/deepest part, based on the cracker method. Bereiter et al. (2015) present a corrected record that partly resolves the issue with a different correlation between CO₂ and Antarctic temperatures found in this oldest part of the records and provide an update of 800,000 years atmospheric CO₂ history. This is the 800,000 year history of CO₂ concentrations that we use in our study.

Figure 2 shows the time series of the global average surface temperature anomalies and atmospheric CO₂ concentrations we obtained from the Snyder dataset, together with glacier periods. Peaks in global average surface temperature anomalies tend to coincide with peak atmospheric CO₂ concentrations as reported in the Snyder dataset. It is important to note, however, that there is substantial measurement error around individual point estimates of temperature, as shown in Appendix Figure A.1. That potential measurement error is, however, larger for *declines* in global average surface temperatures, so that extreme high temperature realizations are measured more precisely than extreme low temperatures. Even taking

into account the 95% confidence interval around the point estimate of temperature, Appendix Figure A.1 shows that the most recent history of temperature increases is anomalous relative to the previous 800,000 years.




Figure 2. CO₂ concentration and change in global average surface temperature from present. The figure shows the time series of CO₂ concentration and the change in global average surface temperature from present, GAST, together with glacier periods shadings (Snyder, 2019). The vertical lines represent the 660 kyrs BP, 400 kyrs BP and 140 kyrs BP. The vertical line at 660 kyrs BP represents the beginning of the out of sample estimate using a 140 kyr-observation rolling window.

Table 1 reports summary statistics for these series. The table illustrates that the last 5,000 years of observations are unusual, with most of the prior observations of global average surface temperatures below those observed most recently.

Table 1: Summary Statistics

Variables	No. of Observation	Mean	S.D.	Min	25th Percentile	Median	75th Percentile	Max
Global average surface temperature anomalies (degrees celsius)	800	-3.35	1.80	-6.95	-4.88	-3.53	-1.99	1.99
Atmospheric CO ₂ concentrations (ppm)	800	225	26	174	203	225	243	288

Source: <http://www.carolynsnyder.com/publications.php>

We perform two types of exercises to minimize the concern that our results are driven by measurement error rather than true dynamics of temperature and CO₂ concentrations. First, in the main body of the paper, we use a pseudo-out-of-sample, rolling window approach in estimating the evolution of the joint temperature-carbon emissions dynamics. The

out-of-sample predictions allow us to construct predictive distributions from the perspective of a hypothetical scientist that existed at a given point in time and was able to access data up to their lifetime, and therefore, puts our current predictions for the future evolution of climate and temperature on an equal footing with this historical perspective. By focusing on rolling-window, rather than expanding window, estimates, the sample size remains exactly the same and the other estimation parameters including the bandwidth also remain unchanged. We emphasize that this is a pseudo-out-of-sample estimation: while we replicate the environment that was available to a researcher from 400 thousand years ago, we also have to acknowledge that estimations performed in today’s world have the privilege of using technological and methodological advances. We report the results based on the time series starting 660 thousand years ago with a rolling window of 140 thousand years, with the aim of minimizing measurement errors due to data breaks.

Second, in the Online Appendix, we perform several subsample analyses to confirm robustness of our results to account for any structural or data breaks that may exist in the reconstructed datasets of temperature and CO₂. For instance, one of our subsamples uses data through 140 thousand years ago as many individual temperature records used for temperature reconstruction process, to our understanding, end around 140 thousand years ago.

2.2 Methodology

We follow the methodology from Adrian et al. (2021), which relies on a non-parametric kernel approach to estimating the one-period-ahead conditional joint distribution and on efficient Markov Chain Monte Carlo (MCMC) techniques to then propagate that one-period-ahead distribution into multi-period-ahead distributions. The non-parametric kernel approach is particularly important in the climate setting as it allows the data to speak on the complexities of temperature and emissions interactions, allowing the possibilities of non-linear dynamics and multi-modal joint distributions. We describe here the construction of the one-period-ahead conditional joint distribution, and refer the interested reader to Adrian et al. (2021)

and the excerpt in Appendix A.2 for details on our implementation of the efficient MCMC for constructing multi-period-ahead distributions.

Consider a time series dataset of n_y endogenous variables $y_{i,t}$, $i = 1, \dots, n_y$ and denote by $y_t = (y_{1,t}, \dots, y_{n_y,t})'$ the vector of date t realizations of the n variables. In addition to the endogenous variables y_t , suppose that we have n_x exogenous predictors x_t . In this paper, we are interested in the case when the exogenous predictors are p lags of y , so that $n_x = p \times n_y$ and

$$x_t = (y'_{t-1}, \dots, y'_{t-p})'.$$

That is, we are interested in estimating a distributional equivalent to vector autoregressions.

We are now ready to write down our kernel estimator (see Li and Racine, 2007, chapter 6.2). Let T be the number of observations of y_t that we have available. Then the joint distribution of y conditional on x can be estimated as

$$\hat{p}(y|x) = \frac{\frac{1}{T-p} \sum_{t=p+1}^T \mathcal{K}_{\omega_y}^y(y - y_t) \mathcal{K}_{\omega_x}^x(x - x_t)}{\frac{1}{T-p} \sum_{t=p+1}^T \mathcal{K}_{\omega_x}^x(x - x_t)}, \quad (1)$$

where $\mathcal{K}_{\omega_y}^y$ and $\mathcal{K}_{\omega_x}^x$ are independent kernels for y and x , given by

$$\mathcal{K}_{\omega_y}^y(y - y_t) = \prod_{i=1}^{n_y} \frac{1}{\omega_{y_i}} \varphi\left(\frac{y_i - y_{i,t}}{\omega_{y_i}}\right) \equiv \prod_{i=1}^{n_y} \mathcal{K}_{\omega_{y_i}}^{y_i}(y_i - y_{i,t}) \quad (2)$$

$$\mathcal{K}_{\omega_x}^x(x - x_t) = \prod_{i=1}^{n_x} \frac{1}{\omega_{x_i}} \varphi\left(\frac{x_i - x_{i,t}}{\omega_{x_i}}\right) \equiv \prod_{i=1}^{n_x} \mathcal{K}_{\omega_{x_i}}^{x_i}(x_i - x_{i,t}). \quad (3)$$

For our baseline results, we use multivariate normal kernels, so that $\varphi(\cdot)$ is the normal probability distribution function, but the kernel estimation can easily be used with alternative kernels (such as multivariate Student kernels) or be modified to accommodate dependent kernels for the endogenous and exogenous variables.

We parameterize the bandwidths based on Li and Racine (2007), who apply a version

of the Silverman (1986) “rule-of-thumb” for joint unconditional density estimation where $w_j = 1.06\sigma_j T^{-1/(4+M+N)}$ for variable j , bandwidth w_j , standard deviation σ_j , sample size T , number of independent variables M , and number of dependent variables N . The rule is derived from minimizing asymptotic mean integrated square error for a Gaussian reference distribution. In the baseline specification, we set $T = 140$, $N = 2$ (CO₂ and temperature anomaly), $M = 2$ (lagged CO₂ and temperature anomaly).⁴

3 Baseline results

We now turn to describing the historical estimated conditional joint evolution of CO₂ concentrations and average global temperature. As described in Section 2, we take a pseudo-real time approach, using rolling windows of 140 observations (140,000 years), and construct out-of-sample estimates at each point in time. We focus on forecast horizons of up to 5 periods ahead (5,000 years ahead), though we investigate the longer-run properties of the forecast distribution towards the end of this Section. In the Online Appendix, we conduct a number of robustness checks, including alternative forecast horizons, alternative choices for the number of observations included in the rolling window, the number of lags included in the distributional VAR, and the bandwidth selected for the non-parametric kernel.

3.1 Predicted joint distribution of CO₂ and temperature

We begin with the full time series of the predicted conditional joint distributions of CO₂ and temperature, for one period and five periods ahead, plotted in Figure 3. More specifically, for each point in time, the Figure plots the contour plot associated with either the one period ahead (Figure 3a) or five periods ahead (Figure 3b) conditional joint density, with darker colors corresponding to greater probability mass. Two features are striking about these estimates. First, the joint conditional distribution is almost never Gaussian – the contour

⁴In the robustness analysis, we also consider expanded T when we include different windows, and $M = 4$ when we include two lags CO₂ and temperature anomaly.

plots are almost never disks – suggesting that the evolution of both global temperature and atmospheric CO₂ concentrations generically has fat-tailed dynamics.

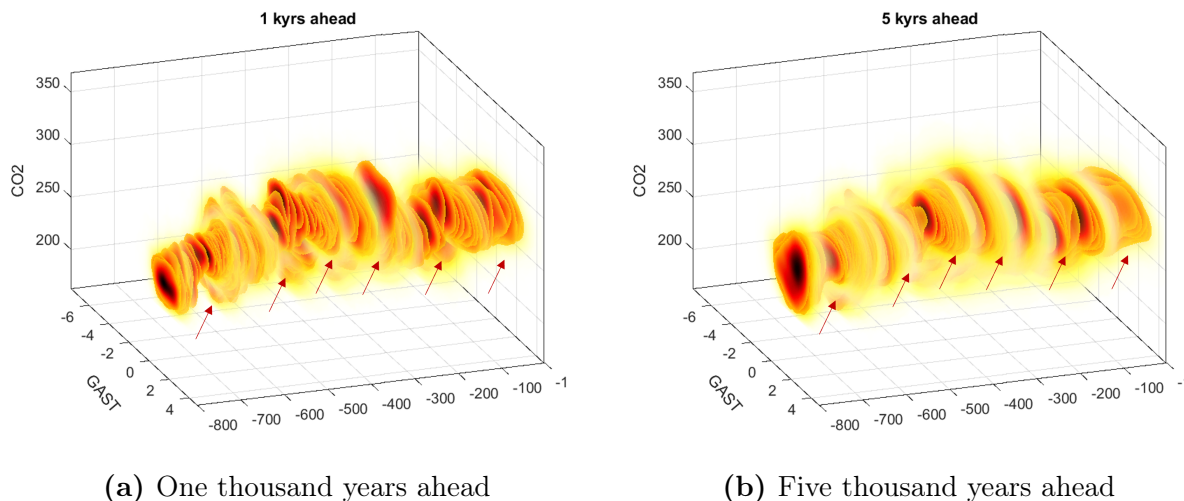

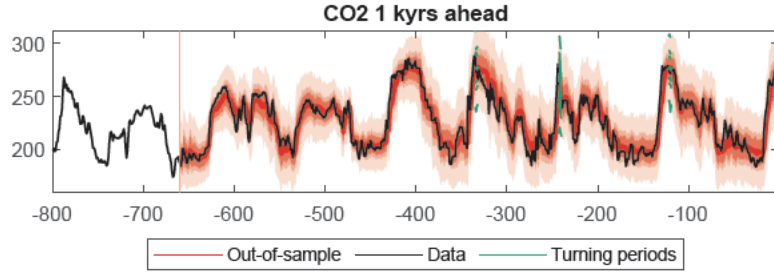
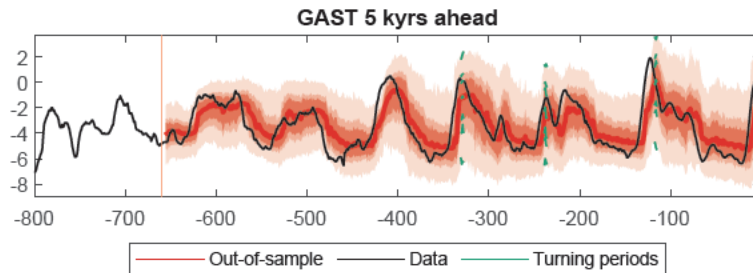



Figure 3. Joint out-of-sample distributions across horizons. The figure shows the joint distribution for one and five thousand years ahead forecasts, as contour plots of the out-of-sample joint distribution in every thousand years, with darker shades of red correspond to higher probability densities. The red arrows correspond to the selected turning periods. Estimation parameters: $k=1.06$, $p=1$, window=140, 800 sample.

Focusing on the marginal distributions one and five periods ahead, in Figure 4 we see that, although, unsurprisingly, the 95% interquartile range is narrower for the one-year-ahead distributions than for the five-year-ahead distributions, the marginal distributions are well calibrated. The realized data is most of the time close to the out-of-sample median, so that the realized data usually falls close to the most likely outcome predicted using 140 observations prior to the date of the forecast. Moreover, as should be expected from a well-calibrated predictive distribution, the realized data is almost never outside of the interquartile range (the dark red shading in Figure 4). The realized data is most likely to be outside of the interquartile range around turning points in climate, highlighted with the green conditional distribution curves, with the deviations away from the median larger for the five-thousand-year ahead distribution than for the one-thousand-year ahead distribution.



(a) One thousand years ahead



(b) Five thousand years ahead

Figure 4. Marginal out-of-sample distributions across horizons. The figure plots the marginal distributions of CO₂ concentration and GAST one and five thousand years ahead out-of-sample (red shaded area), together with the data realization (Snyder, 2019). The red shaded areas correspond to the 2.5th, 16th, 25th, 84th, 75th, and 97.5th percentiles of the out-of-sample predicted distributions. The red lines indicate the median of out-of-sample estimations. The green lines indicate estimations conditioning on selected turning periods. The vertical line at 660 kyrs BP represents the beginning of the out of sample estimate using a 140 kyr-observation rolling window. Estimation parameters: $k=1.06$, $p=1$, window=140, 800 sample.

While the marginal distributions plotted in Figure 4 already hint at the particular salience of climate turning points for the future joint evolution of global temperature and CO₂ concentrations, the joint distributions plotted in Figure 3 have a second, striking feature. Around turning points in either temperature or CO₂ concentration, the conditional joint distribution becomes multimodal and, as we will see in greater detail in Section 3.3, the multimodality becomes more pronounced the longer the forecast horizon. That is, turning points in global climate conditions have been historically associated with substantial probability mass in two potential outcomes: continued increases in both global average temperature and CO₂ concentrations or a cooling of the planet and a decline in carbon concentrations. We now turn to the out-of-sample predicted distributions around the two most recent turning points: during the Eemian (or last interglacial) period and the current turning point.

3.2 Climate turning points and multimodality

The estimated joint distributions for each observation in the time series illustrated in figure 3 suggest that the probability of multimodality is at its highest during the turning points in CO₂ concentrations and temperature anomalies. To understand the out-of-sample forecasts during these turning points in CO₂ concentrations and temperature anomalies, we plot the 1-5 thousand years ahead out-of-sample joint distributions along with their marginal distributions and the realized data (red square) in Figures 5-6 around the two most recent turning points: during the Eemian (or last interglacial) period and the current turning point.

Consider first the conditional out-of-sample joint distributions during the Eemian period, plotted in Figure 5. To put this period in context, the Eemian climate was on average around 1 to 2 degrees Celsius warmer than the current interglacial period, and the proportion of CO₂ in the atmosphere was around 280 parts per million (ppm), well below the current estimates of 421 ppm as of May 2022. The Eemian period ended with the last glacial period in around 115,000 BP, with extensive glaciation in the Northern Hemisphere and temperature declines of more than 5 degrees Celsius in some parts of the Southern Hemisphere. The bottom right

panel of Figure 5, however, shows that the period of global cooling was one of two possible outcomes for climate, with the predicted 5,000 year-ahead joint distribution as of 120,000 BP showing a substantial probability mass for continued increases in both temperature and CO₂ concentrations. This bimodality of the predicted joint distribution begins to emerge already at 123,000 BP, and persists at horizons of 3-5 thousand years ahead up until the beginning of the last glacial period.

The conditional joint distributions plotted in Figure 5 also suggest that early humans were relatively fortunate with the realized glacial period. The realized temperature appears to be at the lower bound of the *warmer* node, so that, as of 120,000 BP there was both a substantial probability of the Earth continuing to heat (corresponding to the warmer, higher CO₂ concentration node) and a noticeable probability of the Earth cooling even more than what was observed (corresponding to the cooler, lower CO₂ concentration node).

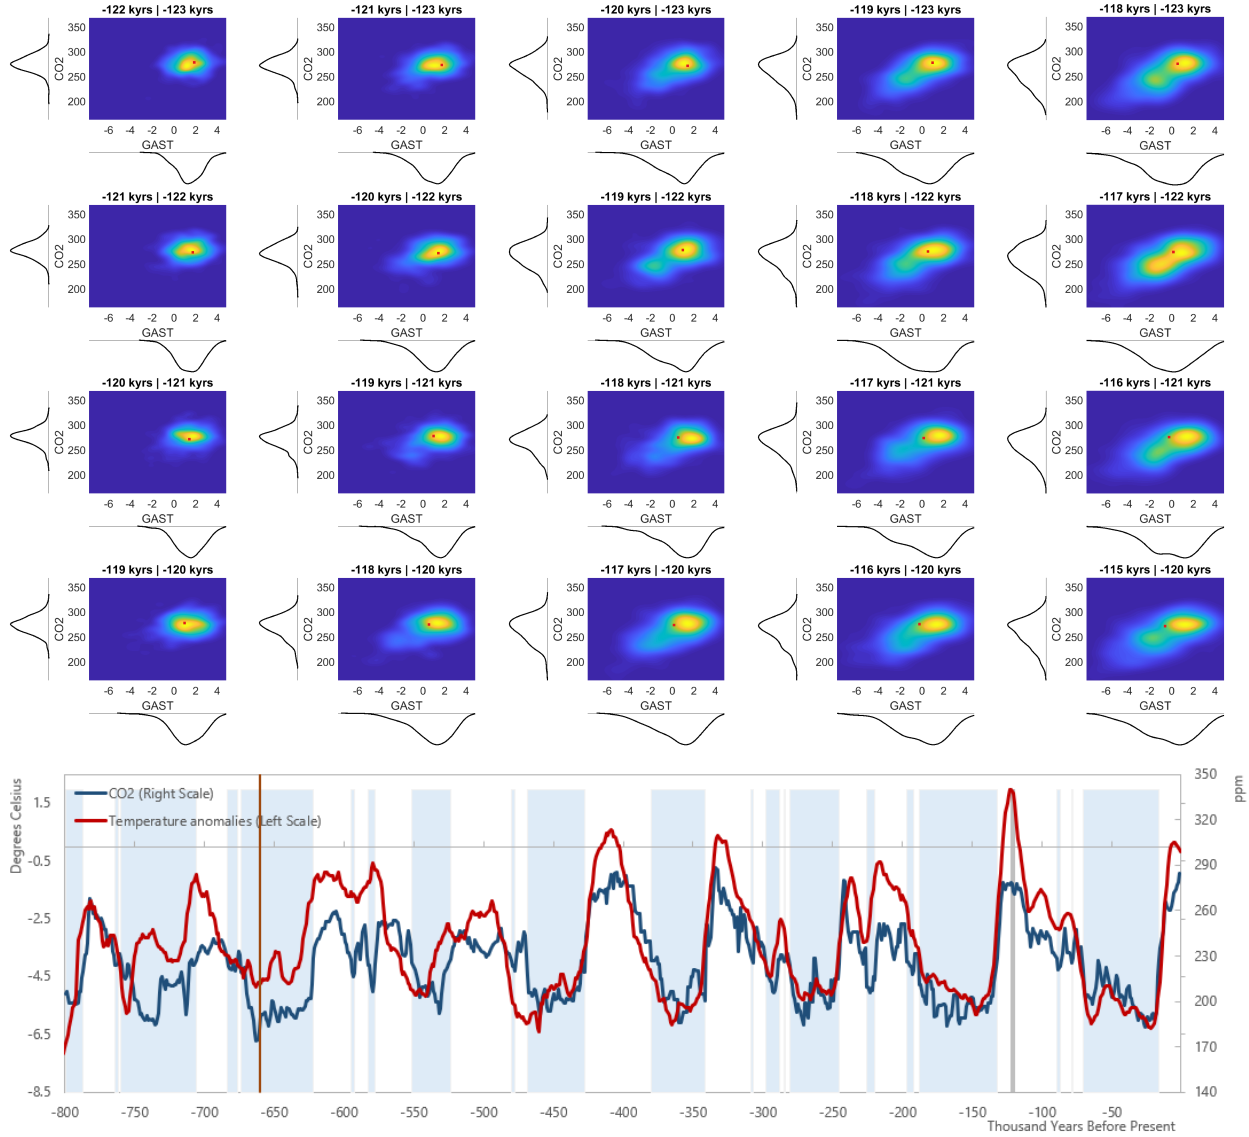


Figure 5. Out-of-sample joint distribution conditioning on 123 to 120 thousand years BP. Contour plots of 1-5 thousand-year-ahead out-of-sample density forecasts of CO₂ concentration and GAST during the selected turning periods equal to 123-120 kyrs BP, with marginal distributions display on the side. Brighter colors indicate greater probability. The red square indicates the ex post realization. The line chart shows the time series of CO₂ concentration and the change in global average surface temperature from present, GAST, together with glacier periods shadings (Snyder, 2019). The gray bar in the line chart shows the turning periods from 123-120 kyrs BP. The vertical line at 660 kyrs BP represents the beginning of the out of sample estimate using a 140 kyr-observation rolling window. Estimation parameters: $k=1.06$, $p=1$, window=140, 800 sample.

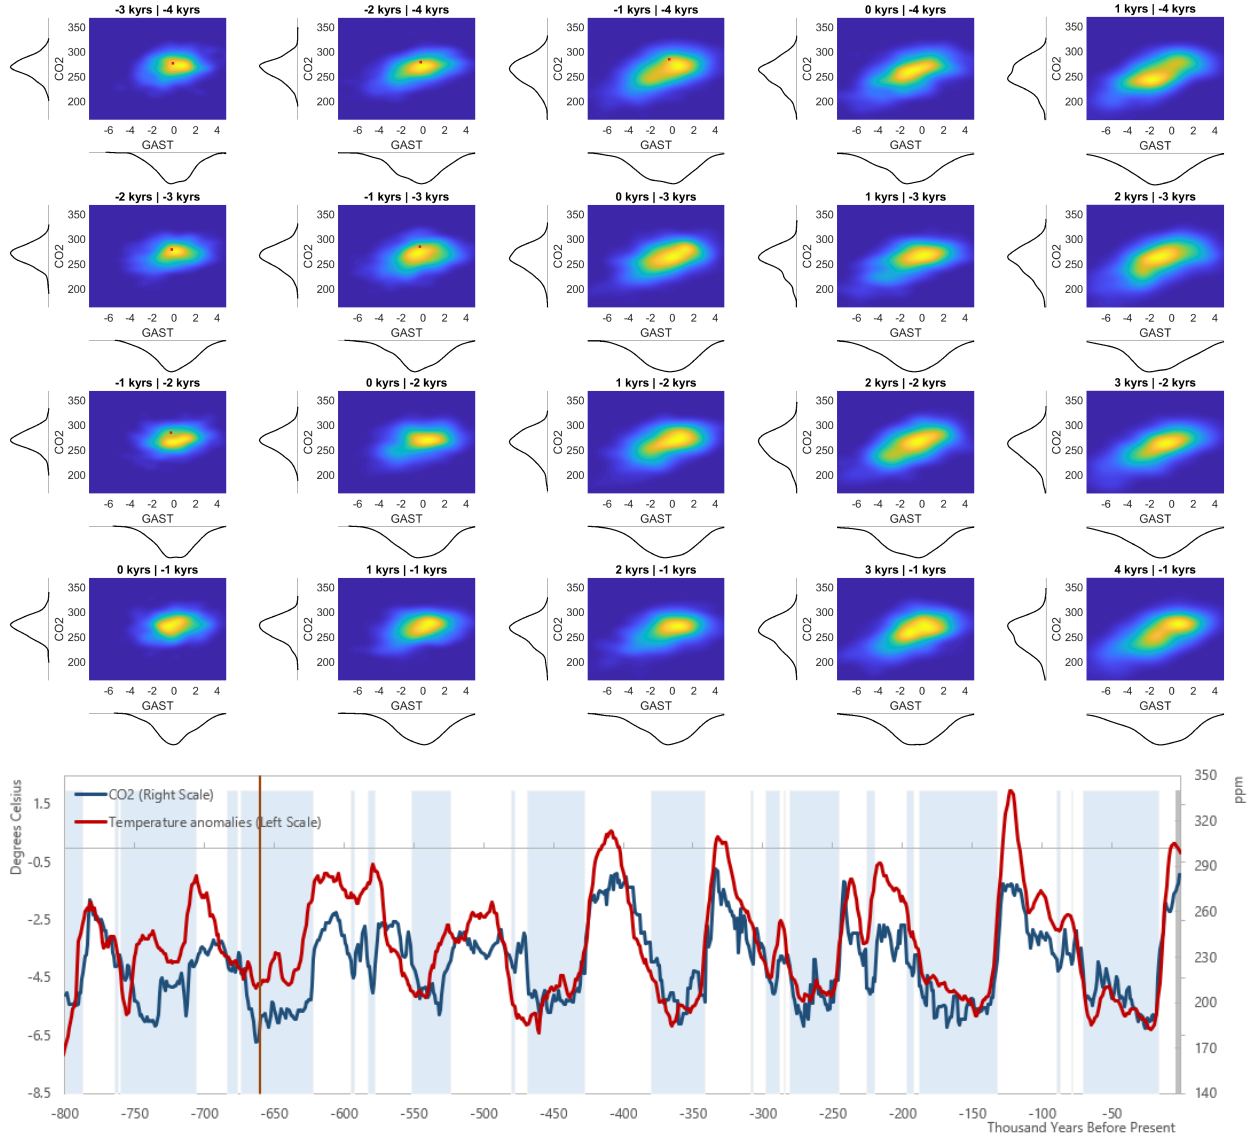


Figure 6. Out-of-sample joint distribution conditioning on 4 to 1 thousand years BP. Contour plots of 1-5 thousand-year-ahead out-of-sample density forecasts of CO₂ concentration and GAST during the selected turning periods equal to 4-1 kyrs BP, with marginal distributions display on the side. Brighter colors indicate greater probability. The red square indicates the ex post realization. The line chart shows the time series of CO₂ concentration and the change in global average surface temperature from present, GAST, together with glacier periods shadings (Snyder, 2019). The gray bar in the line chart shows the turning periods from 4-1 kyrs BP. The vertical line at 660 kyrs BP represents the beginning of the out of sample estimate using a 140 kyr-observation rolling window. Estimation parameters: $k=1.06$, $p=1$, window=140, 800 sample.

Turning to the most recent forecasts, Figure 6 plots the conditional joint distribution up to 5 periods ahead, based on data through the most recent period. Corresponding periods to Milankovitch cycles, we see the emergence of heavily elliptical 5,000 year-ahead joint

distributions as early as 4,000 BP. As the conditioning information gets closer to what we observe today (moving across rows in the Figure), we see that the possibility of extreme temperature and CO₂ concentration realizations remains and, indeed, the likelihood of those extremes realizing in the (geological) near term is increasing. It is worth emphasizing that even the bottom row of Figure 6 conditions on data through the previous millenium only, so that the increases in CO₂ concentrations associated with the Industrial Revolution are not reflected in the conditioning information. That is, even without human activity increasing atmospheric CO₂ concentrations, there was a substantial probability mass associated with substantial warming of the Earth’s atmosphere. We will return to the question of the impact of the most recent temperature and CO₂ concentration readings for these conditional distributions in Section 4.

3.3 Inspecting the mechanism

Focusing on turning points in climate allowed us to identify some points in history at which the conditional joint distribution of climate and CO₂ concentration exhibits multimodality – that is, situations in which there is substantial probability mass associated both with potential future *increases* in global average temperatures and atmospheric CO₂ concentrations and with potential future temperature and CO₂ concentration *declines*. We now examine systematically under which conditions multimodality emerges and how persistent is the predicted bifurcation.


Figure 7 plots the marginal and joint distributions five and fifty thousand years ahead for global average surface temperature anomalies and atmospheric CO₂ concentrations for the latest 140 observations in the sample, conditioning on CO₂ and global temperature realizing in various percentiles of the distribution.⁵ More specifically, each panel in Figure 7 assumes

⁵In Figure 7, we use 140 observations ($T = 140$), containing information starting from 140 kyrs BP till the most recent observation (one thousand years BP). We use one lag of CO₂ concentrations and global temperature anomalies, indicating that the number of independent variables is equal to 4 (denoted as M in this section). The bandwidth is selected by using the estimation approach mentioned above, derived from minimizing asymptotic mean integrated square error for a Gaussian reference distribution.


a particular hypothetical realization of temperature and CO₂ concentrations – for example, that temperature is in the top 30th percentile of its history and CO₂ concentrations are in the bottom 30th percentile of their history – and constructs either the 50 thousand year ahead (Figure 7a) or the five thousand year ahead (Figure 7b) joint distribution of temperature and CO₂ concentrations conditional on that hypothetical climate realization.

Figure 7a shows that in the long run – illustrated through joint 50-thousand-years ahead distributions – the joint distributions converge to its unconditional joint distributions (overlaid in red dotted lines). The 50-thousand-years ahead distributions also show that, when CO₂ concentrations are closer to the right tail of its distribution (e.g., 70th percentile) even as the temperature anomaly is closer to the left tail of its distribution (e.g., 30th percentile), there are signs of multimodalities developing (brighter colors denote greater probability): a possibility of converging into a hotter world or otherwise. More importantly, we see that when both CO₂ concentrations and temperature anomalies are high around the 70th percentile (i.e., towards the right tail of the distribution), the ex-post realized conditions at the 70th percentile tend to fall under the bad mode (i.e., a hotter world).

The probability of multimodalities developing is less pronounced in shorter forecasting horizons, such as the five thousand year ahead horizon (Figure 7b). Among these horizons, probability of multimodality appeared in several of the initial conditions. In the five thousand years ahead estimations (Figure 7b), the joint distribution exhibited signs of multimodality when both CO₂ concentrations and temperature anomalies are around the 50th percentile. The ex-post realized value corresponding to the percentile suggests that the realized temperature anomaly fell into the warmer mode.



(a) Fifty thousand years ahead



(b) Five thousand years ahead

Figure 7. Inspecting the mechanism. The figure plots the evolution of CO_2 concentration and GAST forecasts 50 thousand years and 5 thousand years ahead by varying the initial conditions (30th percentile, 50th percentile, 70th percentile). Red dashed lines correspond to the unconditional distribution. Brighter colors indicate greater probability. The red square indicates the ex-post realization corresponding to the respective percentile based on latest 140 observations. Estimation parameters: $k=1.06$, $p=1$, 140,000 year sample.

4 What do 800,000 years of data tell us about the future?

The results in Section 3 suggest that, even before the Industrial Revolution and the associated increases in carbon emissions and atmospheric CO₂ concentrations, the likelihood of extreme temperature and CO₂ realizations in our future was increasing. We now turn to the question of how the observed recent history of climate readings, and the near term projected path of temperature and CO₂ concentrations impacts the evolution of the predicted joint distribution.

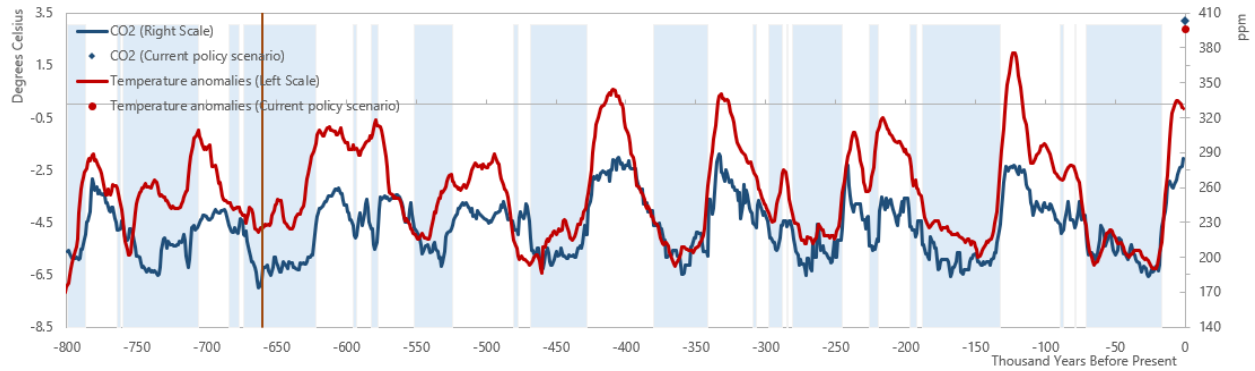
More specifically, we take advantage of the recently published Network for Greening the Financial System (NGFS) scenarios, and expand our sample to include date 0 – that is, to include an observation for the years 2000 – 2999. We consider two of the NGFS scenarios⁶: an *orderly scenario* and an extreme scenario (latter is classified under NGFS *hot house world scenarios*) to explore the impact of climate change and climate policy on global temperature in the future. As the orderly scenario, we use *Net Zero 2050* scenario, where global warming is limited to 1.5 degrees Celsius by accommodating effective climate policies and improving innovation and thereby reaching global net zero CO₂ emissions around 2050.⁷ The extreme (*hot house world*) scenario assumes only the currently implemented policies are preserved (hence named as current policy scenario), and thereby leading to high climate risks. Under this scenario, global temperature is expected to rise to above 3 degrees Celsius compared to pre-industrial levels, while CO₂ emissions do not subside compared to current levels. In either case, we assume that the temperature and CO₂ concentrations projected by the NGFS for 2050 are the average temperature and CO₂ concentrations for the millennium. Without aggressive action to combat climate change, this assumption is most likely the lower bound on climate outcomes over the next thousand years. The conditional distributions plotted in Figure 9 are then most likely an underestimate of the tail of climate outcomes.

⁶The NGFS scenarios are available at [here](#).


⁷We convert CO₂ emissions reported in Mt CO₂/yr and Gt CO₂/yr into ppm using conversion tables and information published by the Carbon Dioxide Information Analysis Center at [here](#).

Figure 8 shows the time series of the global average surface temperature anomalies and atmospheric CO₂ concentrations expanded to include current policy (Figure 8a) and net-zero (Figure 8b) scenarios.⁸ Under the current policy scenario, the latest observation of the global average surface temperature is higher than any of the previous peaks in the sample. While temperature outcomes are somewhat more moderate under the net-zero scenario, they are still well above all except one previous peak in the sample. Under both scenarios, the projected CO₂ concentrations are well above any that have been documented in the previous 800,000 years.

⁸We use last 2,000 years of annual data to construct the temperature deviations that are needed to expand the dataset. For instance, we use annual data to obtain the average pre-industrial temperature in order to calculate the point estimate corresponding to “1.5 °C above pre-industrial levels”. We then calculate the average temperature that corresponds to the period denoted as “present” in our dataset. The deviation between the two datapoints corresponds to the temperature anomaly needed for this analysis. CO₂ concentrations are constructed based on NGFS estimations under each scenario and converted into ppm using Carbon Dioxide Information Analysis Center’s conversion tables.



(a) Current policy scenario



(b) Net-zero scenario

Figure 8. CO₂ concentrations and change in global average surface temperature from present under alternative NGFS scenarios. The figure plots the evolution of CO₂ concentration and the change in global average surface temperature from present, GAST, together with NGFS current policy scenario (Figure 8a) and NGFS net zero scenario (Figure 8b). Light blue shading indicates glacier periods (Snyder, 2019). The vertical line at 660 kyrs BP represents the beginning of the out of sample estimate using a 140 kyr-observation rolling window.

Turning now to the predicted joint distributions of future evolution of temperature and CO₂ concentrations conditional on the two NGFS scenarios, Figure 9 plots the conditional joint distributions and the corresponding marginal distribution up to 5,000 years ahead (5 periods ahead). Comparing the joint distributions under the net-zero scenario (Figure 9b) to the ones under the current policy scenario (Figure 9a), we see that the risks of extreme further warming are somewhat more moderate under the net-zero scenario. However, even under the net-zero scenario, there is substantial probability mass associated with GAST outcomes as high as 4 degrees Celsius 1,000 – 2,000 years ahead (compared to the 3.5 degrees Celsius projected for date 0 even under the current policy scenario). While these risks abate at the

three to five thousand years ahead horizons, it is unclear how much irreversible ecological damage would occur if temperatures were to be 3-4 degrees Celsius above pre-industrial levels for the next three millennia. Notice also that, under the net-zero scenario, we see emerging signs of multimodality in the 5,000-years-ahead distribution, suggesting that, even if temperatures were to revert away from extreme highs in the medium-run, the result may be a new glacial period.

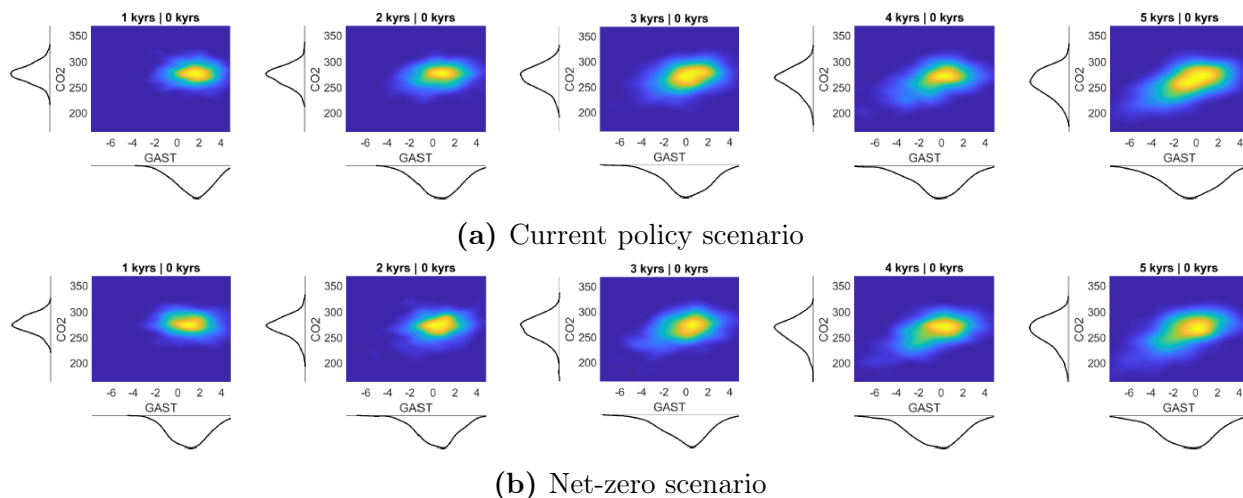


Figure 9. Forecasted joint distributions of CO₂ concentrations and change in global average surface temperature from present under alternative NGFC scenarios. The figure plots the conditional joint distributions 1 – 5 periods (1 – 5 thousand years) ahead of CO₂ concentrations and the change in global average surface temperature from present, GAST, under the NGFS current policy scenario (Figure 9a) and NGFS net zero scenario (Figure 9b). Corresponding marginal distributions plotted on the side. Brighter colors in the density plots indicate greater probability. Estimation parameters: $k = 1.06$, $p = 1$, window=140, 800,000 year sample.

5 Related Literature

We use an 800,000 year history of CO₂ emission and surface temperature. As discussed above, we rely on Snyder (2016, 2019) for our temperature time series. Snyder presents spatially weighted proxy reconstructions of global temperature over the past 2 million years estimated from a multi-proxy database of over 20,000 sea surface temperature point reconstructions. Over the past 800,000 years, polar amplification (the amplification of temperature change at the poles relative to global temperature change) has been stable over time, and global

temperature and atmospheric greenhouse gas concentrations have been closely coupled across glacial cycles. Snyder (2019) argues that stabilization at today’s greenhouse gas levels may already commit Earth to an eventual total warming of 5 degrees Celsius (range 3 to 7 degrees Celsius, 95 per cent credible interval) over the next few millennia as ice sheets, vegetation and atmospheric dust continue to respond to global warming. Our conditional joint distribution estimates in Section 4 are consistent with this prediction, though our estimated conditional distributions suggest that the modal temperature outcomes may be somewhat milder.

The studies most closely related to our paper investigate the relationship between surface temperature and CO₂, often relying on cointegration analysis. The joint conditional distributions we uncover are consistent with the cointegration hypothesis, as the prototypical distribution we estimate is elliptical, so that little probability mass is placed on (low temperature, high CO₂) and (high temperature, low CO₂) outcomes. Kaufmann (2002) use cointegration analysis to elaborate that there is a statistically meaningful relation between surface temperature and changes in the radiative forcing associated with natural variability and human activity. Kaufmann (2002) argue that the increase in global temperatures during the past 130 years can be attributed to changes in radiative forcing associated with natural variability and human activity. Kaufmann and Stock (2006) propose a simplified model of the climate system in which the effect of human activity on surface temperature is reinforced by the simultaneous relationship between surface temperature and the atmospheric concentration of CO₂. Kaufmann and Stock (2006) provide direct evidence that, since 1870, human activity is largely responsible for the increase in global surface temperature and that higher surface temperature has increased the atmospheric concentration of CO₂ and perhaps CH₄ (methane gas).

Kaufmann and Stock (2006) is part of a large literature shows that warming in recent decades is attributable to human activity. Stock (2020) studies a simple and transparent time series regression model to provide independent verification of some key conclusions from climate science models and to demonstrate that warming over the past 140 years is

primarily because of human activity. Phillips et al. (2020) employ an econometric model to address the sensitivity of climate to CO₂ from observational data taking the cooling effect of aerosol particles into consideration. The authors use time series cointegration-based methods that allows treatment of nonstationarity with cointegrated regressors at both the individual station-level and the global aggregate level. The analysis supports the notion of an underestimation of aerosol cooling in the GCMs, which would allow the GCMs to reproduce the observed temperature record with an incorrect TCSG.

Montamat and Stock (2020) propose a model to estimate transient climate response by employing instrumental variables to address the simultaneous causation bias of regressions based on CO₂ and temperature anomalies. A problem with estimating the temperature using observational data is that observed CO₂ concentrations depend in turn on temperature. Therefore, the observed concentration data are endogenous, potentially leading to simultaneous causation bias of regression estimates of the temperature. Montamat and Stock (2020) address this problem by employing instrumental variables regression, which uses changes in radiative forcing external to earth systems to provide quasi-experiments that can be used to estimate the temperature.

An important strand of the literature quantifies the impact of climate change on economic outcomes. Kiley (2021) uses a quantile regression model to account for the relationship between the temperature associated climate change and the distribution of economic growth. This allows him to sketch out the effects of temperature on downside risks to economic growth. Dietz et al. (2018) ask how climate-change mitigation impacts the aggregate consumption risk borne by future generations. To do so, the authors calculate a “climate beta” close to unity for maturities of up to one hundred years using a combination of theory and integrated assessment modeling. Dietz et al. (2018) argue that the overall the net present value of carbon emissions abatement is increasing in the climate beta. Intuitively, a large climate beta not only implies a large climate discount rate, but, also, implies large expected benefits of mitigating climate change in a growing economy. Hence this paper can be used

to illustrate the importance and the benefit of proper climate policy to the CO₂ emission abatement. Dietz et al. (2016) ask what might be the impact of climate change on asset values. The expected “climate value-at-risk” (climate VaR) of global financial assets is 1.8% along a business-as-usual emissions path, amounting to US\$2.5 trillion. Pretis (2020) estimates both the human impact on climate as well as the economic impacts of climate change. The author shows that energy-balance models of climate are equivalent to an econometric cointegrated system and shows estimates of a system of temperatures, ocean heat content, and radiative forcing including greenhouse gases. Pretis (2020) finds statistical support for the cointegrated energy balance model.

International Monetary Fund (2017) finds that increases in temperature have uneven macroeconomic effects, with adverse consequences concentrated in countries with relatively hot climates, such as most low-income countries. In these countries, a rise in temperature lowers per capita output, in both the short and medium term, by reducing agricultural output, suppressing the productivity of workers exposed to heat, slowing investment, and damaging health. Most countries will increasingly feel direct negative effects from unmitigated climate change through warming above optimal levels in currently cooler countries, more frequent natural disasters, rising sea levels, loss of biodiversity, and adverse spillovers from vulnerable countries. Looking ahead, only continued international cooperation and a concerted effort to stem the man-made causes of global warming can limit the long-term risks of climate change.

Our finding of multimodality in turning points closely relates to the literature on tipping points in climate (see e.g. Held and Kleinen, 2004; Lenton et al., 2008; Kriegler et al., 2009; Lenton, 2011, and the literature cited within). We contribute to this literature by providing a non-parametric method for estimating the full future conditional distribution based on aggregate readings only. Consistent with the findings in Lenton (2011), our estimated conditional joint distributions of temperature and CO₂ concentrations in the next few millennia suggest that the global climate may be close to an irreversible tipping point. Dietz et al.

(2021) provide unified estimates of the economic impacts of all eight climate tipping points covered in the economic literature so far using a meta-analytic integrated assessment model (IAM) with a modular structure. The model includes national-level climate damages from rising temperatures and sea levels for 180 countries, calibrated on detailed econometric evidence and simulation modeling. The tipping points with the largest effects are dissociation of ocean methane hydrates and thawing permafrost.

6 Conclusion

In this paper we study how climate evolves over the past 800,000 years applying the non-parametric approach developed in Adrian et al. (2021) to joint dynamics of the long time series of global temperature collected by Snyder (2016) and of global CO₂ concentration. This approach allows us to assign probabilities to the precise nature of climate tipping points, including the level at which tipping points in, for example, temperature may occur, how long it may take the overall climate to reach that tipping point, as well as the magnitude of the consequences should such a tipping point occur. We take a non-parametric approach, allowing us to construct term structures of joint distributions for temperature and CO₂ at any point given the available data on global mean temperature and atmospheric carbon concentrations. We also provide forecasts of the future evolution of these two key climate variables.

We find that the temperature-CO₂ dynamics are non-linear, so that large deviations in either temperature or CO₂ concentrations take a long time to correct – on the scale of multiple millennia. Furthermore, we uncover joint dynamics of temperature and CO₂ concentrations that exhibit multimodality around historical turning points in temperature and concentration cycles, so that prior to the start of cooling periods, there is a noticeable probability that temperature and CO₂ concentrations may continue to increase. Finally, evaluating the future evolution of temperature and CO₂ concentration conditional on alternative scenarios

realizing, we document that, even conditional on the net-zero 2050 scenario, there remains a significant risk of elevated temperatures for at least a further five millennia.

The key question for future research is how economic activity would interplay with these CO₂ – temperature dynamics. On the one hand, economic activity greatly impacts the level of atmospheric CO₂ and hence temperature tipping points. On the other hand, temperature directly affects economic activity. Unfortunately, the time series for economic activity is much shorter than what is available for CO₂ and temperature, and hence any study of the interplay of climate with economic activity will have to do with that shorter time series. But, we conjecture here that such approaches might worsen the tipping point dynamics, and potentially exacerbate the already high estimates of global warming.

References

- ADRIAN, T., N. BOYARCHENKO, AND D. GIANNONE (2021): “Multimodality in Macro-Financial Dynamics,” *International Economic Review*, 62, 861–886.
- BEREITER, B., S. EGGLESTON, J. SCHMITT, C. NEHRBASS-AHLES, T. F. STOCKER, H. FISCHER, S. KIPFSTUHL, AND J. CHAPPELLAZ (2015): “Revision of the EPICA Dome C CO₂ record from 800 to 600 kyr before present,” *Geophysical Research Letters*, 42, 542–549.
- DIETZ, S., A. BOWEN, C. DIXON, AND P. GRADWELL (2016): “‘Climate value at risk’ of global financial assets,” *Nature Climate Change*, 6, 676–679.
- DIETZ, S., C. GOLLIER, AND L. KESSLER (2018): “The climate beta,” *Journal of Environmental Economics and Management*, 87, 258–274.
- DIETZ, S., J. RISING, T. STOERK, AND G. WAGNER (2021): “Economic impacts of tipping points in the climate system,” *Proceedings of the National Academy of Sciences*, 118, e2103081118.
- HELD, H. AND T. KLEINEN (2004): “Detection of climate system bifurcations by degenerate fingerprinting,” *Geophysical Research Letters*, 31.
- INTERNATIONAL MONETARY FUND, R. D. (2017): “World Economic Outlook, October 2017: Seeking Sustainable Growth: Short-Term Recovery, Long-Term Challenges,” Tech. rep., International Monetary Fund.
- KAUFMANN, ROBERT K., H. K. AND J. H. STOCK (2006): “Emissions, Concentrations, & Temperature: A Time Series Analysis,” *Climatic Change*, 77, 249–278.
- KAUFMANN, R. K. (2002): “Cointegration Analysis of Hemispheric Temperature Relations,” *Journal of Geophysical Research*, 107.
- KILEY, M. T. (2021): “Growth at Risk from Climate Change,” Finance and economics discussion series 2021–054, Federal Reserve Board of Governors.
- KRIEGLER, E., J. W. HALL, H. HELD, R. DAWSON, AND H. J. SCHELLNHUBER (2009): “Imprecise probability assessment of tipping points in the climate system,” *Proceedings of the national Academy of Sciences*, 106, 5041–5046.
- LENTON, T. M. (2011): “Early warning of climate tipping points,” *Nature climate change*, 1, 201–209.
- LENTON, T. M., H. HELD, E. KRIEGLER, J. W. HALL, W. LUCHT, S. RAHMSTORF, AND H. J. SCHELLNHUBER (2008): “Tipping elements in the Earth’s climate system,” *Proceedings of the national Academy of Sciences*, 105, 1786–1793.
- LI, Q. AND J. RACINE (2007): *Nonparametric Econometrics: Theory and Practice*, Princeton University Press.

- MONTAMAT, G. AND J. H. STOCK (2020): “Quasi-Experimental Estimates of the Transient Climate Response Using Observational Data,” *Climatic Change*, 160, 361–371.
- PHILLIPS, P. C. B., 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, 256–273.
- SILVERMAN, B. (1986): *Density Estimation for Statistics and Data Analysis*, no. 9780412246203 in Monographs on Statistics and Applied Probability, Chapman and Hall.
- SNYDER, C. (2016): “Evolution of global temperature over the past two million years,” *Nature*, 538.
- (2019): “Revised estimates of paleoclimate sensitivity over the past 800,000 years,” *Climatic Change*, 156, 121–138.
- STOCK, J. H. (2020): “Climate Change, Climate Policy, and Economic Growth,” *NBER Macroeconomics Annual*, 34, 399–419.

Online Appendix

A Data and Methodology

A.1 Data - Confidence Interval Data

Figure A.1 shows the global average surface temperature anomalies along with a ± 2.5 confidence interval, accounting for estimation uncertainties.




Figure A.1. Change in global average surface temperature from present with confidence interval. The figure shows the change in global average surface temperature from present, GAST, with red shaded area corresponding to 2.5th and 97.5th percentiles range (Snyder, 2019).

A.2 Methodology - Efficient Monte Carlo

Given an estimated one-period-ahead distribution $\hat{p}(y|x)$, we can use Monte Carlo simulations to estimate h -period-ahead distributions by sequentially drawing paths of y . In principle, these draws can be made directly from the inverse CDF implied by $\hat{p}(y|x)$ by drawing u from a (multinomial) uniform distribution and finding y that solves $y = \hat{P}^{-1}(u|x)$. We increase the efficiency of this procedure by discretizing the state space as follows.

Algorithm 1. Simulating paths of y .

To estimate the H -period-ahead distribution of y , generate n_{sim} paths of y as follows.

1. Discretize the state-space. Set $\kappa = \kappa_0$. For each variable j , loop through:

- (a) Initialize grid with bound $[\min(y_j) - \kappa\sigma_{y,j}, \max(y_j) + \kappa\sigma_{y,j}]$ and grid point increments of $\Delta \times \sigma_{y,j}$.
- (b) For each grid point $y_{j,i}$, compute the kernel CDF $\Phi\left(\frac{y_j - y_{j,i}}{\omega_{y_j}}\right)$.
- (c) Verify that the kernel PDF integrates to one. Verify that the kernel CDF has a maximum of 1 within a tolerance of ε . If not, reset $\kappa = (1 + \delta) \times \kappa$ and repeat.

2. For each simulated path $k = 1, \dots, n_{sim}$, loop through each horizon $h = 1, \dots, H$ by drawing $y_{t+h}^k | y_{t+h-1}^k$ from the grid established in Step 1 and verify the normalization condition.

The choices κ_0 , ε , Δ , and δ control the speed and accuracy of the estimator, with κ_0 controlling the size of the initial state-space, Δ controlling the fineness of the discretization grid, ε the tolerance for integration error, and δ the speed with which the size of the state space grows while the estimated density is “missing” probability mass. We set $\kappa_0 = 0.1$, $\varepsilon = 10^{-5}$, $\Delta = 1/20$, and $\delta = 0.05$. Similar results shown when we include restrictive parameters.

B Subsample Analyses and Robustness Checks

This Subsample Analyses and Robustness Checks performs several analyses, validating the robustness of the results presented in the paper. These robustness checks account for any structural or data breaks that may exist in the reconstructed datasets of temperature and CO₂ concentrations. The Subsample Analyses and Robustness Checks is structured as follows: Section B.1 presents results based on a time series that starts 140 thousand years ago as opposed to 800 thousand years ago; Section B.2 presents results under a different lag length selection; Section B.3 presents results when the observation rolling window is altered from 140-observations to 400-observations; Section B.4 concludes the Subsample Analyses and Robustness Checks by presenting results under a different optimal bandwidth parameter.

B.1 140 sample, p=1, k=1.06

Section B.1 presents out of sample distributions across various horizons using subsample data from 140 thousand years ago through one thousand years ago. As mentioned in the Section 2.1, many individual temperature records used for temperature reconstruction process, to our understanding, end around 140 thousand years ago and therefore we chose a 140-observation rolling window in our paper. We use a 140-observation subsample as a robustness check to see whether the findings presented in the paper alters. Precisely, instead of using the latest 140 observations as a 140-observation rolling window for each out of sample estimation, we use a cumulative window in this robustness test. Therefore, this robustness test includes a time series starting from 140 thousand year ago till the estimation point. The other estimation parameters such as the bandwidth also change as the number of observations included in the estimation varies.

In Figures B.1a to B.1c, we compare the realized and out-of-sample predicted marginal distributions using 140 thousand years of data and a cumulative window. Figure B.1a shows the realized and out-of-sample estimates of CO₂ concentrations and temperature anomalies in one thousand years ahead. In line with the results presented in the paper, the median out-of-sample distributions are very similar to the realized conditions. When the out-of-sample medians slightly deviate from the realizations, our out-of-sample estimate is generally more conservative. Figures B.1b and B.1c show the realized and out-of-sample estimates of CO₂ concentrations and temperature anomalies in three and five thousand years ahead respectively. These results are also in line with the results presented in the paper. The

medians of out-of-sample estimates are not tightly aligned with the realized conditions, as in the one-thousand years ahead estimate. While the deviation between the medians out-of-sample estimations and realized conditions is widened, our estimation is broadly conservative. In addition, realized conditions both three and five thousand years ahead remain mostly within the out-of-sample estimation range of 25-and-75 percentile.




Figure B.1a. Out of sample quantiles in one thousand years ahead. The figure plots the marginal distributions of CO₂ concentration and GAST one thousand years ahead out-of-sample (red shaded area), together with the data realization, using 140 observation sample (Snyder, 2019). The red shaded areas correspond to the 2.5th, 16th, 25th, 75th, 84th, and 97.5th percentiles of out-of-sample estimations. The red lines indicate the median of out-of-sample estimations. The vertical line at 70 kyrs BP represents the beginning of the out of sample estimate using a cumulative window. Estimation parameters: $k=1.06$, $p=1$, 140 sample.

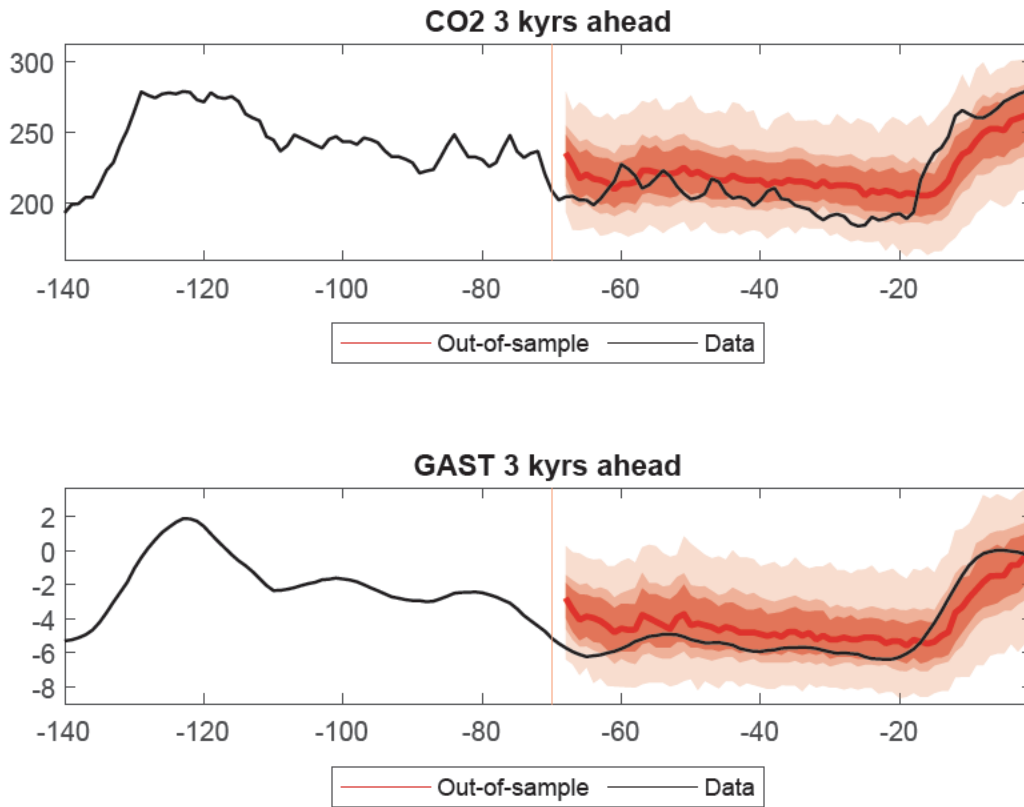


Figure B.1b. Out of sample quantiles in three thousand years ahead. The figure plots the marginal distributions of CO₂ concentration and GAST three thousand years ahead out-of-sample (red shaded area), together with the data realization, using 140 observation sample (Snyder, 2019). The red shaded areas correspond to the 2.5th, 16th, 25th, 75th, 84th, and 97.5th percentiles of out-of-sample estimations. The red lines indicate the median of out-of-sample estimations. The vertical line at 70 kyrs BP represents the beginning of the out of sample estimate using a cumulative window. Estimation parameters: $k=1.06$, $p=1$, 140 sample.




Figure B.1c. Out of sample quantiles in five thousand years ahead. The figure plots the marginal distributions of CO₂ concentration and GAST five thousand years ahead out-of-sample (red shaded area), together with the data realization, using 140 observation sample (Snyder, 2019). The red shaded areas correspond to the 2.5th, 16th, 25th, 75th, 84th, and 97.5th percentiles of out-of-sample estimations. The red lines indicate the median of out-of-sample estimations. The vertical line at 70 kyrs BP represents the beginning of the out of sample estimate using a cumulative window. Estimation parameters: $k=1.06$, $p=1$, 140 sample.

B.2 806 sample, $p=2$, $k=1.06$, window=140

In Section B.2, we present out of sample distributions across various horizons using two lags of CO₂ concentrations and temperature anomalies (i.e., changing the lag length from 1 to 2 for each variable). The other parameters such as the rolling window, sample size, and the bandwidth parameters are the same as those presented in the paper. In Figures B.2a to B.2c, we compare the realized and out-of-sample predicted marginal distributions using two lags of CO₂ concentrations and temperature anomalies. That is, each estimation point conditions on more information. Figure B.2a shows the realized and out-of-sample estimates of CO₂ concentrations and temperature anomalies one thousand years ahead. The result is very similar to what is presented in Figure 4 in the paper, which is based on one lag of CO₂ concentrations and temperature anomalies. Our median out-of-sample distributions

are aligned with the realized conditions. Figures B.2b to B.2c show the realized and out-of-sample estimates of CO₂ concentrations and temperature anomalies in three and five thousand years ahead, respectively. The out-of-sample medians tend to slightly deviate from the realized conditions. But realized conditions both three and five thousand years ahead remain mostly within the out-of-sample estimation range of 25/75 percentile. The only exception where the realized temperature anomaly conditions deviate beyond the 2.5/97.5 percentile out-of-sample is in the five thousand years ahead projection around 115 to 110 thousand years ago.




Figure B.2a. Out of sample quantiles in one thousand years ahead. The figure plots the marginal distributions of CO₂ concentration and GAST one thousand years ahead out-of-sample (red shaded area), together with the data realization, using 806 observation sample, based on two lags CO₂ concentration and GAST (Snyder, 2019). The red shaded areas correspond to the 2.5th, 16th, 25th, 75th, 84th, and 97.5th percentiles of out-of-sample estimations. The red lines indicate the median of out-of-sample estimations. The vertical line at 660 kyrs BP represents the beginning of the out of sample estimate using a 140 kyr-observation rolling window. Estimation parameters: $k=1.06$, $p=2$, window=140, 800 sample.




Figure B.2b. Out of sample quantiles in three thousand years ahead. The figure plots the marginal distributions of CO₂ concentration and GAST three thousand years ahead out-of-sample (red shaded area), together with the data realization, using 806 observation sample, based on two lags CO₂ concentration and GAST (Snyder, 2019). The red shaded areas correspond to the 2.5th, 16th, 25th, 75th, 84th, and 97.5th percentiles of out-of-sample estimations. The red lines indicate the median of out-of-sample estimations. The vertical line at 660 kyrs BP represents the beginning of the out of sample estimate using a 140 kyr-observation rolling window. Estimation parameters: $k=1.06$, $p=2$, window=140, 800 sample.




Figure B.2c. Out of sample quantiles in five thousand years ahead. The figure plots the marginal distributions of CO₂ concentration and GAST five thousand years ahead out-of-sample (red shaded area), together with the data realization, using 806 observation sample, based on two lags CO₂ concentration and GAST (Snyder, 2019). The red shaded areas correspond to the 2.5th, 16th, 25th, 75th, 84th, and 97.5th percentiles of out-of-sample estimations. The red lines indicate the median of out-of-sample estimations. The vertical line at 660 kyrs BP represents the beginning of the out of sample estimate using a 140 kyr-observation rolling window. Estimation parameters: $k=1.06$, $p=2$, window=140, 800 sample.

We also explored whether the predicted joint distributions presented in the paper, based on near-term projected path of temperature and CO₂ concentrations in line with the NGFS scenarios, are robust to lag length selection. Figures B.3a and B.3b show how the global temperatures one to five thousand years from now will change under the current policy scenario and the net-zero scenario presented in the paper. Similar to the findings presented in the paper, Figure B.3a shows a very large right tail in temperature anomalies where temperature over four degrees higher could develop. The temperature anomalies are estimated to revert back over the three to five thousand years as well, but the world might have already suffered the catastrophic events caused by the rising temperature anomaly by then. Figure B.3b shows the results when climate policies are introduced early and become gradually more stringent. In line with the paper’s findings, we can also see a right tail in the marginal distribution of temperature anomaly, but the right tail is smaller if more stringent climate

policies are implemented.




Figure B.3a. Joint distribution forecasts of CO₂ concentration and GAST given current policy scenario data. Contour plots of 1-5 thousand-year-ahead out-of-sample density forecasts of CO₂ concentration and GAST given current policy scenario data, using 806 observation sample, based on two lags CO₂ concentration and GAST, with marginal distributions display on the side. Brighter colors indicate greater probability. Estimation parameters: $k=1.06$, $p=2$, window=140, 800 sample.




Figure B.3b. Joint distribution forecasts of CO₂ concentration and GAST given net-zero scenario data. Contour plots of 1-5 thousand-year-ahead out-of-sample density forecasts of CO₂ concentration and GAST given net-zero scenario data, using 806 observation sample, based on two lags CO₂ concentration and GAST, with marginal distributions display on the side. Brighter colors indicate greater probability. Estimation parameters: $k=1.06$, $p=2$, window=140, 800 sample.

B.3 806 sample, $p=1$, $k=1.06$, window=400

In Section B.3, we present the out-of-sample results across horizons using a 400-observation rolling window as opposed to a 140-observation rolling window used in the paper. As mentioned in the paper, to our understanding, many individual temperature records used for temperature reconstruction process of our primary data source end around 140 thousand years ago. We perform this robustness test to make sure that, similar to a larger rolling window such as a 400-observation rolling window, the 140-observation rolling window also contain the enough large right tail events where both CO₂ concentrations and temperature anomalies are reaching their peak values for the model to estimate.

Figures B.4a to B.4c show the time series of CO₂ concentrations and temperature anomalies, together with out-of-sample marginal distributions using 806 thousand years of data and a 400-thousand-year rolling window. Figure B.4a shows the realized and out-of-sample estimates of CO₂ concentrations and temperature anomalies in one thousand years ahead. In line with figure 4a in the paper, the median out-of-sample distributions are similar to the realized conditions. In very few exceptions, we notice that the median out-of-sample temperature anomalies in 400-observation rolling window are lower than those of in 140-observation rolling window (for instance, the period between 128 thousand years BP to 125 thousand

years BP). During these periods, the 140-observation rolling window results perform better, as shown in the paper, where the median out-of-sample distributions are closer to the realized conditions. Figures B.4b to B.4c show the realized and out-of-sample estimates of CO₂ concentrations and temperature anomalies in three and five thousand years ahead, respectively. Even this robustness check generated results that are in line with the findings of the paper. Realized conditions both three and five thousand years ahead remain mostly within the out-of-sample estimation range of 25/75 percentile. The only exception where the realized temperature anomaly conditions deviate beyond the 2.5/97.5 percentile out-of-sample is in the five thousand years ahead projection around 115 to 110 thousand years ago.




Figure B.4a. Out of sample quantiles in one thousand years ahead. The figure plots the marginal distributions of CO₂ concentration and GAST one thousand years ahead out-of-sample (red shaded area), together with the data realization, using 806 observation sample, and 400 observation rolling window (Snyder, 2019). The red shaded areas correspond to the 2.5th, 16th, 25th, 75th, 84th, and 97.5th percentiles of out-of-sample estimations. The red lines indicate the median of out-of-sample estimations. The green lines indicate one thousand years ahead estimations conditioning on selected turning periods. The vertical line at 400 kyrs BP represents the beginning of the out of sample estimate using a 400 kyr-observation rolling window. Estimation parameters: $k=1.06$, $p=1$, window=400, 800 sample.




Figure B.4b. Out of sample quantiles in three thousand years ahead. The figure plots the marginal distributions of CO₂ concentration and GAST three thousand years ahead out-of-sample (red shaded area), together with the data realization, using 806 observation sample, and 400 observation rolling window (Snyder, 2019). The red shaded areas correspond to the 2.5th, 16th, 25th, 75th, 84th, and 97.5th percentiles of out-of-sample estimations. The red lines indicate the median of out-of-sample estimations. The green lines indicate one thousand years ahead estimations conditioning on selected turning periods. The vertical line at 400 kyrs BP represents the beginning of the out of sample estimate using a 400 kyr-observation rolling window. Estimation parameters: $k=1.06$, $p=1$, window=400, 800 sample.




Figure B.4c. Out of sample quantiles in five thousand years ahead. The figure plots the marginal distributions of CO₂ concentration and GAST five thousand years ahead out-of-sample (red shaded area), together with the data realization, using 806 observation sample, and 400 observation rolling window (Snyder, 2019). The red shaded areas correspond to the 2.5th, 16th, 25th, 75th, 84th, and 97.5th percentiles of out-of-sample estimations. The red lines indicate the median of out-of-sample estimations. The green lines indicate one thousand years ahead estimations conditioning on selected turning periods. The vertical line at 400 kyrs BP represents the beginning of the out of sample estimate using a 400 kyr-observation rolling window. Estimation parameters: $k=1.06$, $p=1$, window=400, 800 sample.

In Figures B.5a to B.5d, we evaluate the one to five thousand years ahead out-of-sample joint and marginal distributions associated with turning points under a 400-observation window. The conditioning periods are 336-333, 224-241, 123-120, and 4-1 thousand years BP. Even under a 400-observation window, our model is able to quantify the evolution of the uncertainly related CO₂ concentrations and temperature anomalies well, showing that most of the realized conditions (red square in each panel) appear within the predicted out-of-sample joint distributions. The only two exceptions where the realized conditions slightly deviate from the predicted out-of-sample distributions are in two thousand years and three thousand years ahead projections, conditioning on 123 thousand years BP. While multimodality does not emerge as pronounced as in the paper, there are some conditioning periods such as 243 thousand years BP (Figure B.5b), two to three thousand years ahead

out-of-sample predictions revealing greater probability of multimodality emerging.

Figure B.5d shows the out-of-sample joint distributions associated with the latest turning point (conditioning period of four to one thousand years BP). As with the results presented in the paper, we observe the out-of-sample joint distributions widening and greater probability of multimodality appearing as the forecasting horizon increases, indicating a greater probability of much warmer temperature could realize in one thousand to four thousand years from now.

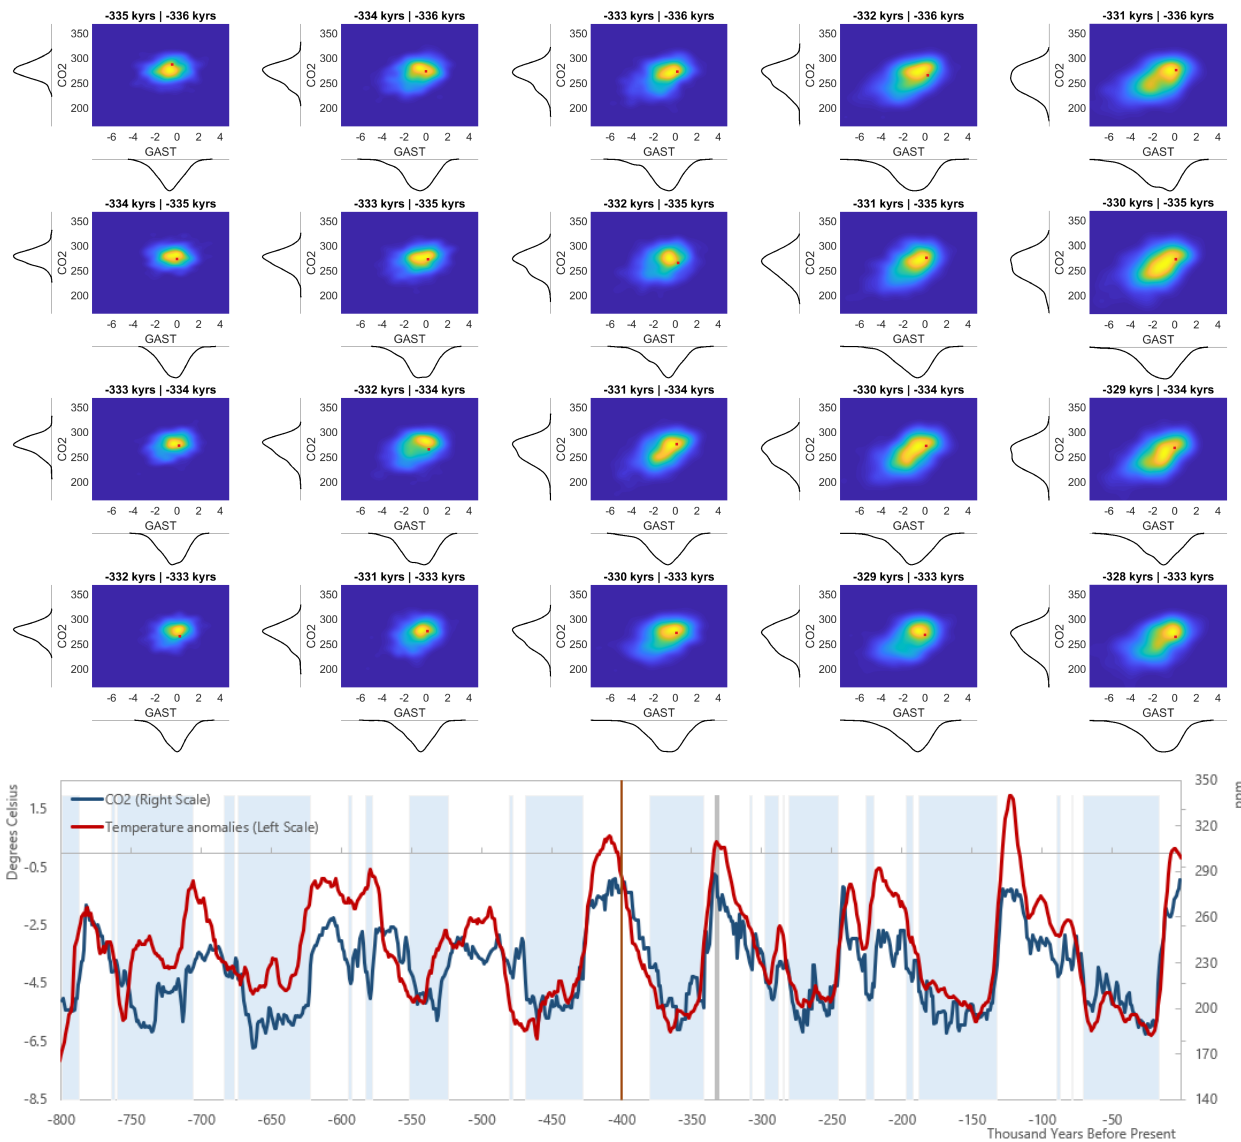


Figure B.5a. Out-of-sample joint distribution conditioning on 336 to 333 thousand years BP. Contour plots of 1-5 thousand-year-ahead out-of-sample density forecasts of CO₂ concentration and GAST during the selected turning periods equal to 336-333 kyrs BP, with marginal distributions display on the side, using 806 observation sample, and 400 observation rolling window. Brighter colors indicate greater probability. The red square indicates the ex post realization. The line chart shows the time series of CO₂ concentration and the change in global average surface temperature from present, GAST, together with glacier periods shadings (Snyder, 2019). The gray bar in the line chart shows the turning periods from 336-333 kyrs BP. The vertical line at 400 kyrs BP represents the beginning of the out of sample estimate using a 400 kyr-observation rolling window. Estimation parameters: $k=1.06$, $p=1$, window=400, 800 sample.




Figure B.5b. Out-of-sample joint distribution conditioning on 244 to 241 thousand years BP. Contour plots of 1-5 thousand-year-ahead out-of-sample density forecasts of CO₂ concentration and GAST during the selected turning periods equal to 244-241 kyrs BP, with marginal distributions display on the side, using 806 observation sample, and 400 observation rolling window. Brighter colors indicate greater probability. The red square indicates the ex post realization. The line chart shows the time series of CO₂ concentration and the change in global average surface temperature from present, GAST, together with glacier periods shadings (Snyder, 2019). The gray bar in the line chart shows the turning periods from 244-241 kyrs BP. The vertical line at 400 kyrs BP represents the beginning of the out of sample estimate using a 400 kyr-observation rolling window. Estimation parameters: $k=1.06$, $p=1$, window=400, 800 sample.




Figure B.5c. Out-of-sample joint distribution conditioning on 123 to 120 thousand years BP. Contour plots of 1-5 thousand-year-ahead out-of-sample density forecasts of CO₂ concentration and GAST during the selected turning periods equal to 123-120 kyrs BP, with marginal distributions display on the side, using 806 observation sample, and 400 observation rolling window. Brighter colors indicate greater probability. The red square indicates the ex post realization. The line chart shows the time series of CO₂ concentration and the change in global average surface temperature from present, GAST, together with glacier periods shadings (Snyder, 2019). The gray bar in the line chart shows the turning periods from 123-120 kyrs BP. The vertical line at 400 kyrs BP represents the beginning of the out of sample estimate using a 400 kyr-observation rolling window. Estimation parameters: $k=1.06$, $p=1$, window=400, 800 sample.

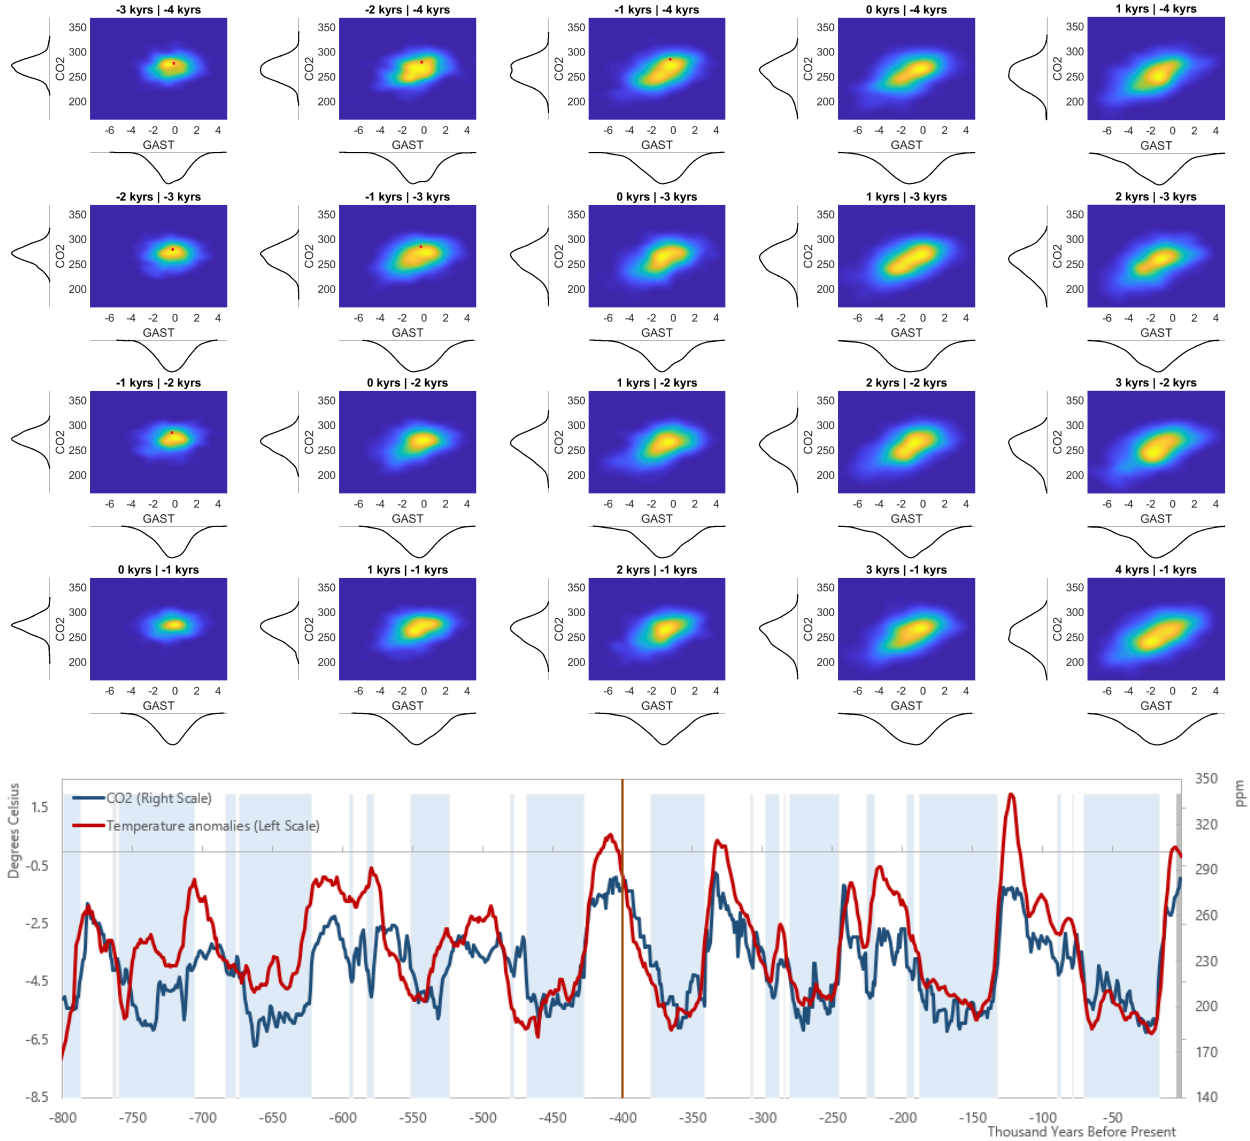


Figure B.5d. Out-of-sample joint distribution conditioning on 4 to 1 thousand years BP. Contour plots of 1-5 thousand-year-ahead out-of-sample density forecasts of CO₂ concentration and GAST during the selected turning periods equal to 4-1 kyrs BP, with marginal distributions display on the side, using 806 observation sample, and 400 observation rolling window. Brighter colors indicate greater probability. The red square indicates the ex post realization. The line chart shows the time series of CO₂ concentration and the change in global average surface temperature from present, GAST, together with glacier periods shadings (Snyder, 2019). The gray bar in the line chart shows the turning periods from 4-1 kyrs BP. The vertical line at 400 kyrs BP represents the beginning of the out of sample estimate using a 400 kyr-observation rolling window. Estimation parameters: $k=1.06$, $p=1$, window=400, 800 sample.

We also explored whether the predicted joint distributions presented in the paper, based on near-term projected path of temperature and CO₂ concentrations in line with the NGFS scenarios, are robust to a 400-observation rolling window. Figures B.6a and B.6b show how the global temperatures one to five thousand years from now will change under the NGFS current policy and net-zero scenarios. Figure B.6a shows if adequate climate policies are not

introduced within this century, the probability of much warmer temperature could increase in the next millennium. Similar to our paper’s results, the temperature anomalies are estimated to revert back over the three to five thousand years under this scenario. In Figure B.6b, we see that under the net-zero scenario, the shape of the marginal distribution in temperature anomalies changes and a greater probability of less warmer temperatures could realize in the next thousand years if more stringent climate policies are introduced. These robustness test results are very much in line with the findings of our paper.




Figure B.6a. Joint distribution forecasts of CO₂ concentration and GAST given current policy scenario data. Contour plots of 1-5 thousand-year-ahead out-of-sample density forecasts of CO₂ concentration and GAST given current policy scenario data, with marginal distributions display on the side, using 806 observation sample, and 400 observation rolling window. Brighter colors indicate greater probability. Estimation parameters: $k=1.06$, $p=1$, window=400, 800 sample.




Figure B.6b. Joint distribution forecasts of CO₂ concentration and GAST given net-zero scenario data. Contour plots of 1-5 thousand-year-ahead out-of-sample density forecasts of CO₂ concentration and GAST given net-zero scenario data, with marginal distributions display on the side, using 806 observation sample, and 400 observation rolling window. Brighter colors indicate greater probability. Estimation parameters: $k=1.06$, $p=1$, window=400, 800 sample.

B.4 806 sample $p=1$, $k=0.3$, window=140

In Section B.4, we introduce a new method of optimal bandwidth parameter selection and present out of sample distribution across various horizons using newly selected optimal bandwidth ($k=0.3$), 806 samples, and a 140-observation rolling window.

As mentioned in section two, we parameterize the bandwidths based on Li and Racine (2007) method, which applies a version of the Silverman (1986) “rule-of-thumb” for joint unconditional density estimation where $w_j = 1.06\sigma_j T^{-1/(4+M+N)}$ for variable j , bandwidth w_j , standard deviation σ_j , sample size T , number of independent variables M , and number of dependent variables N . In the main body of the paper, we use $k = 1.06$. In this section, we parameterize the bandwidths as being proportional to the in-sample unconditional standard deviation of the corresponding variable: $w_j = k\sigma_j T^{-1/(4+M+N)}$ and choose a single

proportionally constant $k = k_{y_i} = k_{x_j}$ for all $i = 1, \dots, n_y, j = 1, \dots, n_x$ to maximize the predictive accuracy of the resultant one-period-ahead conditional joint distribution.

In Figure B.7, we plot the relationship between the bandwidth proportionality constant k and the out-of-sample (log) predictive score for the one-quarter-ahead and one-year-ahead distribution. The figure shows that there is an interior solution for the optimal parameter of proportionality, and that the performance is maximized for $k \approx 0.3$. Figure B.8 shows the optimal bandwidth parameter for the one-period ahead conditional joint distribution of CO₂ concentrations and temperature anomalies starting from 400 thousand years BP. The optimal bandwidth parameter is volatile as the estimation point varies. In the latest estimation point, the out-of-sample predicted performance is maximized for $k \approx 0.3$. We also plot the optimal bandwidth parameter for the one-period-ahead joint distribution of CO₂ concentrations and temperature anomalies in the latest period (Figure B.9). We can observe that the out-of-sample log predictive scores are highest when $k \approx 0.3$.




Figure B.7. Optimal bandwidth selection. This figure plots the out-of-sample log predictive scores for the one-period-ahead conditional joint distribution of CO₂ concentration and GAST, as a function of the bandwidth proportionality constant k . Predictor variables: one lag of CO₂ concentration and GAST. Estimation parameters: $p=1$, window=140, sample=800.




Figure B.8. Optimal bandwidth selection around each period. This figure plots the optimal bandwidth for the one-period-ahead conditional joint distribution of CO₂ cocentration and GAST across estimated period, based on the out-of-sample log predictive scores. Predictor variables: one lag of CO₂ concentration and GAST. Estimation parameters: $p=1$, window=140, sample=800.




Figure B.9. Optimal bandwidth selection in latest period. This figure plots the out-of-sample log predictive scores for the one-period-ahead conditional joint distribution of CO₂ concentration and GAST in the latest period. Predictor variables: one lag of CO₂ concentration and GAST. Estimation parameters: $p=1$, window=140, sample=800.

Figures B.10a to B.10d plot the marginal and joint distributions ahead of one thousand years (which corresponds to one period ahead), three, five, and fifty thousand years for global average surface temperature anomalies and atmospheric CO₂ concentrations for the latest 140 observations in the sample, using $k = 0.3$. Figure B.10a shows that while joint distributions of CO₂ concentrations and temperature anomalies do not as tightly converge to its unconditional joint distributions (overlaid in red dotted lines) as in what is presented in our paper, there are signs of multimodality developing conditioning on all the selected percentiles.

In line with the results presented in the paper, the probability of multimodalities is less pronounced in shorter forecasting horizons such as joint distributions ahead of one thousand to five thousand years (illustrated through Figures B.10b to B.10d). In the five thousand years ahead estimations (Figure B.10b), the joint distribution exhibited signs of multimodality when CO₂ concentrations are around the 70th percentile and temperature anomalies are around the 30th percentile. The ex-post realized value corresponding to the percentile sug-

gests that the realized temperature anomaly falls into the colder mode. However, as CO₂ concentrations move from 30th percentile to 70th percentile, the probability of a warmer mode increases. In the three thousand years ahead distribution (Figure B.10c), while less pronounced, we also see signs of multimodality when CO₂ concentrations are around the 70th percentile and temperature anomalies are around the 30th percentile.




Figure B.10a. Inspecting the mechanism in 50 thousand years ahead. The figure displays the evolution of CO₂ concentration and GAST forecasts in 50 thousand years ahead by varying the initial conditions (30th percentile, 50th percentile, 70th percentile). Red dashed lines correspond to the unconditional distribution. Brighter colors indicate greater probability. The red square indicates the ex-post realization corresponding to the respective percentile based on latest 140 observations. Estimation parameters: $k=0.3$, $p=1$, sample=140.




Figure B.10b. Inspecting the mechanism in five thousand years ahead. The figure displays the evolution of CO₂ concentration and GAST forecasts in five thousand years ahead by varying the initial conditions (30th percentile, 50th percentile, 70th percentile). Brighter colors indicate greater probability. The red square indicates the ex-post realization corresponding to the respective percentile based on latest 140 observations. Estimation parameters: $k=0.3$, $p=1$, sample=140.




Figure B.10c. Inspecting the mechanism in three thousand years ahead. The figure displays the evolution of CO₂ concentration and GAST forecasts in three thousand years ahead by varying the initial conditions (30th percentile, 50th percentile, 70th percentile). Brighter colors indicate greater probability. The red square indicates the ex-post realization corresponding to the respective percentile based on latest 140 observations. Estimation parameters: $k=0.3$, $p=1$, sample=140.




Figure B.10d. Inspecting the mechanism in one thousand years ahead. The figure displays the evolution of CO₂ concentration and GAST forecasts in one thousand years ahead by varying the initial conditions (30th percentile, 50th percentile, 70th percentile). Brighter colors indicate greater probability. The red square indicates the ex-post realization corresponding to the respective percentile based on latest 140 observations. Estimation parameters: $k=0.3$, $p=1$, sample=140.

In Figures B.11a to B.11c, we compare the realized and out-of-sample predicted marginal distributions using 806 thousand years of data, a 140-thousand-year rolling window, and a 0.3 bandwidth parameter. Figure B.11a shows the realized and out-of-sample estimates of CO₂ concentrations and temperature anomalies in one thousand years ahead. As with the results in the paper, the median out-of-sample distributions are tightly aligned with the realized conditions. The range of one thousand years ahead distributions in CO₂ concentrations and temperature anomalies narrower. The out-of-sample prediction remains very similar to the realized condition when CO₂ concentrations and temperature anomalies are reaching their peak values. Figures B.11b and B.11c show the realized and out-of-sample estimates of CO₂ concentrations and temperature anomalies three and five thousand years ahead, respectively. In most periods, the median out-of-sample predictions remain similar to the realized conditions. However, as the range of distributions narrows, there are more exceptions where the realized temperature anomaly distributions deviate beyond the 2.5/97.5 percentile out-of-sample estimates (for instance, temperature anomalies in period between 130 to 125

thousand years BP for both three and five thousand years ahead).




Figure B.11a. Out of sample quantiles in one thousand years ahead. The figure plots the marginal distributions of CO₂ concentration and GAST one thousand years ahead out-of-sample (red shaded area), together with the data realization (Snyder, 2019). The red shaded areas correspond to the 2.5th, 16th, 25th, 75th, 84th, and 97.5th percentiles of out-of-sample estimations. The red lines indicate the median of out-of-sample estimations. The green lines indicate one thousand years ahead estimations conditioning on selected turning periods. The vertical line at 660 kyrs BP represents the beginning of the out of sample estimate using a 140 kyr-observation rolling window. Estimation parameters: $k=0.3$, $p=1$, window=140, 800 sample.




Figure B.11b. Out of sample quantiles in three thousand years ahead. The figure plots the marginal distributions of CO₂ concentration and GAST three thousand years ahead out-of-sample (red shaded area), together with the data realization (Snyder, 2019). The red shaded areas correspond to the 2.5th, 16th, 25th, 75th, 84th, and 97.5th percentiles of out-of-sample estimations. The red lines indicate the median of out-of-sample estimations. The green lines indicate three thousand years ahead estimations conditioning on selected turning periods. The vertical line at 660 kyrs BP represents the beginning of the out of sample estimate using a 140 kyr-observation rolling window. Estimation parameters: $k=0.3$, $p=1$, window=140, 800 sample.




Figure B.11c. Out of sample quantiles in five thousand years ahead. The figure plots the marginal distributions of CO₂ concentration and GAST five thousand years ahead out-of-sample (red shaded area), together with the data realization (Snyder, 2019). The red shaded areas correspond to the 2.5th, 16th, 25th, 75th, 84th, and 97.5th percentiles of out-of-sample estimations. The red lines indicate the median of out-of-sample estimations. The green lines indicate five thousand years ahead estimations conditioning on selected turning periods. The vertical line at 660 kyrs BP represents the beginning of the out of sample estimate using a 140 kyr-observation rolling window. Estimation parameters: $k=0.3$, $p=1$, window=140, 800 sample.

In Figures B.12a to B.12c, we plot the one to five thousand years ahead out-of-sample joint distributions along with their marginal distributions for the conditioning periods of 336-333, 224-241, 123-120, and 4-1 thousand years BP (i.e., the turning points) using 806 thousand years of data, a 140-thousand-year rolling window, and a 0.3 bandwidth parameter. Unlike the results presented in the paper, not all realized conditions appear to be within the predicted out-of-sample joint distributions (for instance, results in conditioning period of 244-243 thousand years BP) as the range of predicted distribution narrow significantly when a 0.3 bandwidth parameter is used. However, we can still observe the probability of multimodality emerging. For example, out-of-sample joint distributions associated with the turning point in CO₂ concentrations and temperature anomalies around 242-241 thousand years ago exhibit clear signs of multimodalities developing in three to five thousand years ahead out-of-sample predictions (Figure B.12b). The time series of CO₂ concentrations and

temperature anomalies in this period suggest that the CO₂ concentrations are decreasing from its peak value and temperature anomalies are increasing. This two-directional movement drives the probability of two modes, where temperature could have either realized in the warmer mode or in the colder mode. We observe that the realized temperature appears to be at the upper bound of the colder mode.

Figure B.12d shows the out-of-sample joint distributions associated with the latest turning point (conditioning period of 4-1 thousand years BP). As with the results presented in the paper, we observe that out-of-sample joint distributions widening and greater probability of multimodality appearing as the forecasting horizon widens.




Figure B.12a. Out-of-sample joint distribution conditioning on 336 to 333 thousand years BP. Contour plots of 1-5 thousand-year-ahead out-of-sample density forecasts of CO₂ concentration and GAST during the selected turning periods equal to 336-333 kyrs BP, with marginal distributions display on the side. Brighter colors indicate greater probability. The red square indicates the ex post realization. The line chart shows the time series of CO₂ concentration and the change in global average surface temperature from present, GAST, together with glacier periods shadings (Snyder, 2019). The gray bar in the line chart shows the turning periods from 336-333 kyrs BP. The vertical line at 660 kyrs BP represents the beginning of the out of sample estimate using a 140 kyr-observarion rolling window. Estimation parameters: $k=0.3$, $p=1$, window=140, 800 sample.

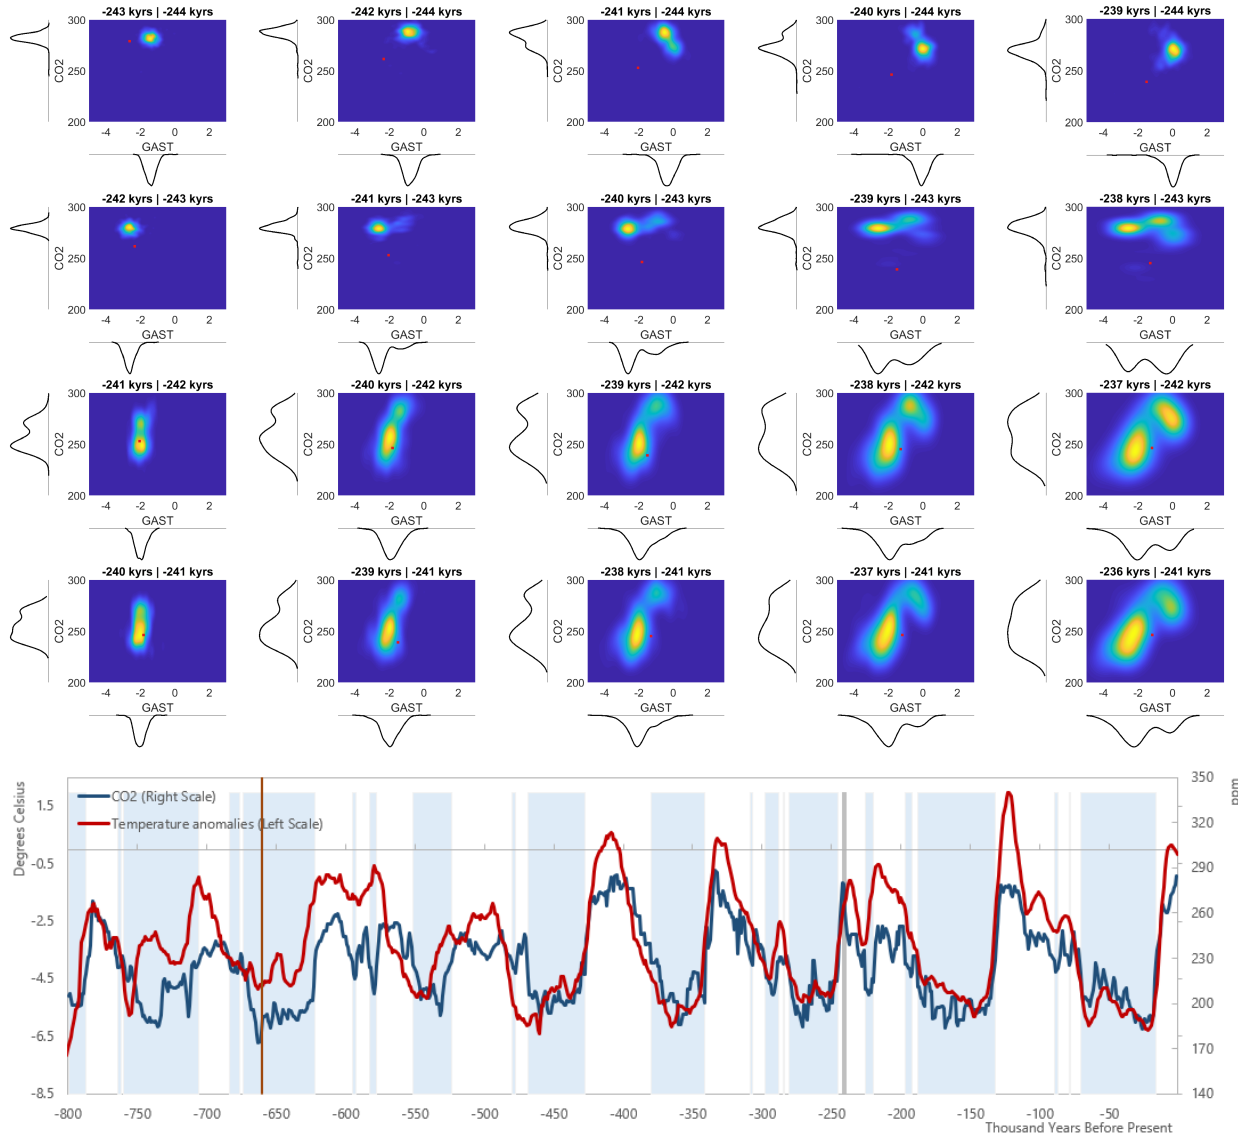


Figure B.12b. Out-of-sample joint distribution conditioning on 244 to 241 thousand years BP. Contour plots of 1-5 thousand-year-ahead out-of-sample density forecasts of CO₂ concentration and GAST during the selected turning periods equal to 244-241 kyr BP, with marginal distributions display on the side. Brighter colors indicate greater probability. The red square indicates the ex post realization. The line chart shows the time series of CO₂ concentration and the change in global average surface temperature from present, GAST, together with glacier periods shadings (Snyder, 2019). The gray bar in the line chart shows the turning periods from 244-241 kyr BP. The vertical line at 660 kyr BP represents the beginning of the out of sample estimate using a 140 kyr-observarion rolling window. Estimation parameters: $k=0.3$, $p=1$, window=140, 800 sample.

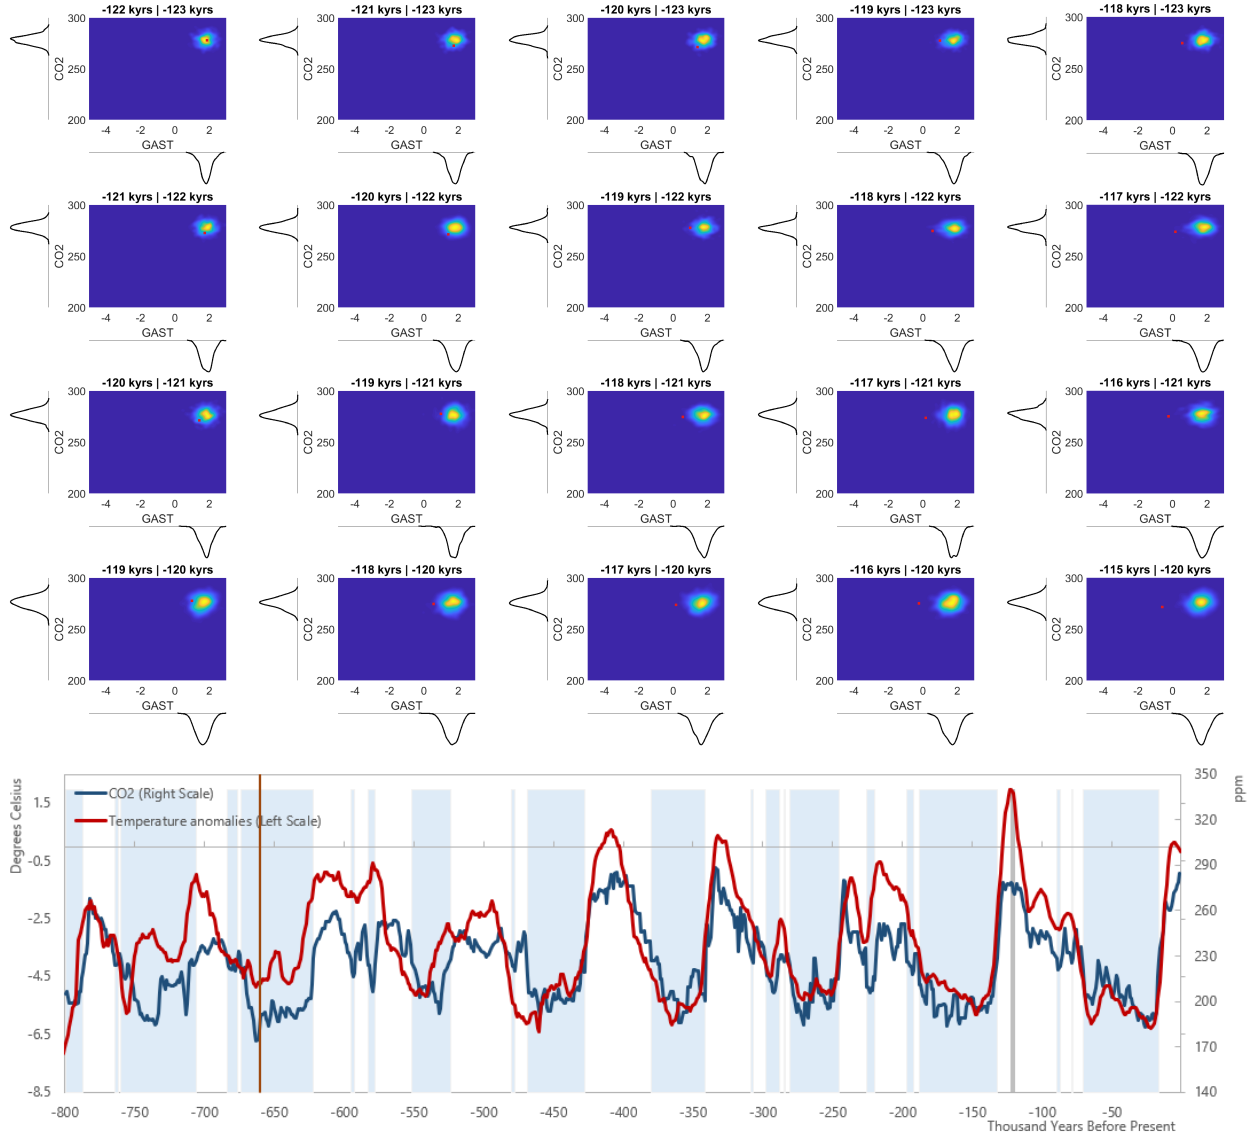


Figure B.12c. Out-of-sample joint distribution conditioning on 123 to 120 thousand years BP. Contour plots of 1-5 thousand-year-ahead out-of-sample density forecasts of CO₂ concentration and GAST during the selected turning periods equal to 123-120 kyrs BP, with marginal distributions display on the side. Brighter colors indicate greater probability. The red square indicates the ex post realization. The line chart shows the time series of CO₂ concentration and the change in global average surface temperature from present, GAST, together with glacier periods shadings (Snyder, 2019). The gray bar in the line chart shows the turning periods from 123-120 kyrs BP. The vertical line at 660 kyrs BP represents the beginning of the out of sample estimate using a 140 kyr-observation rolling window. Estimation parameters: $k=0.3$, $p=1$, window=140, 800 sample.




Figure B.12d. Out-of-sample joint distribution conditioning on 4 to 1 thousand years BP. Contour plots of 1-5 thousand-year-ahead out-of-sample density forecasts of CO₂ concentration and GAST during the selected turning periods equal to 4-1 kyrs BP, with marginal distributions display on the side. Brighter colors indicate greater probability. The red square indicates the ex post realization. The line chart shows the time series of CO₂ concentration and the change in global average surface temperature from present, GAST, together with glacier periods shadings (Snyder, 2019). The gray bar in the line chart shows the turning periods from 4-1 kyrs BP. The vertical line at 660 kyrs BP represents the beginning of the out of sample estimate using a 140 kyr-observation rolling window. Estimation parameters: $k=0.3$, $p=1$, window=140, 800 sample.

We also explored whether the predicted joint distributions presented in the paper, based on near-term projected path of temperature and CO₂ concentrations in line with the NGFS scenarios, are robust to changing the bandwidth parameter to 0.3. Figures B.13a and B.13b show how the global temperatures one to five thousand years from now will change if climate policies are introduced early and become gradually more stringent (Figure B.13b) or

on the contrary global efforts are insufficient to halt significant global warming by 2100 (Figure B.13a). Similar to the results presented in the paper, the marginal distributions in temperature anomalies are tilted towards the right under both scenarios. But in Figure B.13a under the current policy scenario, there is a greater probability of higher temperature anomalies developing. On the contrary, Figure B.13b under the net-zero scenario shows a slightly higher probability of relatively lower temperature anomalies developing when more stringent climate policies are introduced.




Figure B.13a. Joint distribution forecasts of CO₂ concentration and GAST given current policy scenario data. Contour plots of 1-5 thousand-year-ahead out-of-sample density forecasts of CO₂ concentration and GAST given current policy scenario data, with marginal distributions display on the side. Brighter colors indicate greater probability. Estimation parameters: $k=0.3$, $p=1$, window=140, 800 sample.




Figure B.13b. Joint distribution forecasts of CO₂ concentration and GAST given net-zero scenario data. Contour plots of 1-5 thousand-year-ahead out-of-sample density forecasts of CO₂ concentration and GAST given net-zero scenario data, with marginal distributions display on the side. Brighter colors indicate greater probability. Estimation parameters: $k=0.3$, $p=1$, window=140, 800 sample.