

Sceptic priors and climate consensus

Grant R. McDermott *

This version: May 2019

Abstract

How much evidence would it take to convince sceptics that they are wrong about climate change? I explore this question within a Bayesian framework. I consider a group of stylised sceptics and examine how these individuals update their beliefs in the face of current and continuing climate change. I find that available evidence in the form of instrumental climate data tends to overwhelm all but the most extreme priors. Most sceptics form updated beliefs about climate sensitivity that correspond closely to estimates from the scientific literature. However, belief convergence is a non-linear function of prior strength. It becomes increasingly difficult to convince the remaining pool of sceptics that they are wrong. I discuss the necessary conditions for consensus formation under Bayesian learning and show how apparent deviations from the Bayesian ideal still be accommodated within the same conceptual framework. I argue that a generalized Bayesian model thus provides a bridge between competing theories of climate scepticism as a social phenomenon.

JEL Codes: Q54, Q58, C11

Keywords: climate change, climate scepticism, social cost of carbon, Bayesian econometrics

*University of Oregon. Email: grantmcd@uoregon.edu. An earlier draft of this paper was circulated under the title, "Sceptic priors and climate policy". I am indebted to various people for helpful comments and suggestions, including Øivind A. Nilsen, Jonas Andersson, Ken Rice, Christopher Costello, Matthew Burgess, Roman Mureika, Linda Nøstbakken, Timo Goeschl and Corbett Grainger, as well as participants at various workshops and seminars held at the Norwegian School of Economics, University of California at Santa Barbara, and the University of Guelph. I am especially grateful to Chris Hope for making the PAGE09 model available to me.

1 Introduction

Climate change has come to represent a defining policy issue of our age. Yet support for comprehensive climate policy at the global scale remains elusive. Many policy makers and citizens are openly sceptical about the human role in our changing climate, despite decades of accumulated research and an overwhelming scientific consensus.¹ What are we to make of this scepticism? And just how much evidence would it take to convince climate sceptics that they are wrong? In the present paper, I seek to answer these questions within a Bayesian framework that combines a range of sceptic beliefs (i.e. priors) with available climate data. My goal is to pin down plausible rates of convergence with the scientific consensus, by examining how different sceptics update their beliefs in the face of current and continuing climate change. In so doing, I hope to shed light on our current policy impasse and offer some remarks about the possibility for finding common ground in the future.

Many studies have explored the cultural and psychological factors underlying climate scepticism. These include [Kahan et al. \(2011, 2012\)](#); [McCright and Dunlap \(2011a,b\)](#); [Corner et al. \(2012\)](#); [Ranney et al. \(2012\)](#); [Clark et al. \(2013\)](#) — see [Hornsey et al. \(2016\)](#) for a recent literature review. My present concern is less with the origins of scepticism than what it represents, namely a set of *beliefs* about the likely causes of global warming, which will in turn affect how new information about those causes is interpreted. A convenient way to model such beliefs is by defining scepticism in terms of climate sensitivity, i.e. the temperature response to a doubling of CO₂. More precisely, we can map sceptic beliefs directly to subjective estimates of climate sensitivity, because they both describe the likely causes and probability distribution of future warming. The particular measure of climate sensitivity that I focus on here is the transient climate response (TCR). Formally, TCR describes the warming at the time of CO₂ doubling — i.e. after 70 years — in a 1% per year increasing CO₂ experiment ([IPCC, 2013](#), Box 12.2). For the purposes of this paper, however, it will simply be thought of as the contemporaneous change in global temperature that results from a steady doubling of atmospheric CO₂.

According to the the Intergovernmental Panel on Climate Change ([IPCC, 2013](#)), TCR is “likely” to be somewhere in the range of 1.0–2.5 °C. This corresponds roughly to a 66–100% probability interval in IPCC terminology. The IPCC further emphasises the

¹For review and further discussion, see: [Oreskes \(2004\)](#); [Anderegg et al. \(2010\)](#); [Doran and Zimmerman \(2011\)](#); [Cook et al. \(2013\)](#); [Verheggen et al. \(2014\)](#); [Tol \(2014\)](#); [Cook et al. \(2016\)](#); [Saad \(2019\)](#).

inherently Bayesian nature of climate sensitivity estimates, going so far as to state:

“[T]he probabilistic estimates available in the literature for climate system parameters, such as ECS [i.e. equilibrium climate sensitivity] and TCR have all been based, implicitly or explicitly, on adopting a Bayesian approach and therefore, even if it is not explicitly stated, involve using some kind of prior information.” (IPCC, 2013, p. 922)

To understand why classical (i.e. frequentist) methods are ill-suited for the task of producing credible estimates of climate sensitivity, recall that frequentism interprets probability as the limiting frequency in a large number of repeated draws. Such a narrow definition holds little relevance to the question of climate sensitivity, for which there exists but one unique value. There is no population of “sensitivities” to draw samples from. I too adopt a Bayesian approach to the question of climate sensitivity and its concomitant policy implications. However, my approach differs from the previous literature along several dimensions.

The most obvious point of departure is the fact that I deliberately focus on the beliefs of sceptics. Priors for determining climate sensitivity are usually based on paleo data, the judgements of scientific experts, or noninformative methods. Such approaches may possess obvious scientific merit for establishing a best estimate of climate sensitivity. Yet, they are of limited relevance for understanding people’s motivations and voting behaviour when it comes to actual climate policy. My approach is to take sceptics at their word and work through to the conclusions of their stated priors. In other words, my goal is to recover posterior probabilities about the rate and causes of climate change that are logically consistent with the initial beliefs of these sceptics.

Contrarian climate beliefs have also been largely ignored in the economic policy literature to date. The handful of studies that *do* consider policy options from the sceptic perspective have tended to emphasise edge scenarios like climate catastrophe and irreversibility. For example, Kagan (2014) introduces various levels of scepticism into an integrated assessment model (IAM) and solves for optimal climate policy under the threat of catastrophe. He shows that even vehement sceptics will not follow an unbounded emissions path. The possibility of catastrophe — which to their minds is exogenously determined — causes agents to draw down on their capital stock in preemptive fashion, leaving insufficient investment for continued economic growth be-

yond some point. Similarly, [Kiseleva \(2016\)](#) introduces an IAM of heterogeneous agents that includes various degrees of climate scepticism. She too shows that a world comprised only of sceptical policy makers will make sufficient investments in mitigation measures to avoid catastrophic outcomes. In this case, the mechanism is a dominant subset of “weak” sceptics who are sufficiently concerned by anthropogenic climate change that they reduce their emissions accordingly. Neither [Kagan \(2014\)](#) or [Kiseleva \(2016\)](#) allow for learning in their numerical simulations.² However, theoretical work by [Van Wijnbergen and Willems \(2015\)](#) show that climate sceptics actually have an incentive to reduce emissions, since it will facilitate learning about the true causes of climate change. While it is possible for an increase in emissions to yield similar learning effects, the irreversibility of climate change renders this an inferior strategy in their model. From a methodological perspective, the present paper differs from these earlier studies in that it combines Bayesian learning with an empirical framework. Unlike the existing numerical ([Kagan, 2014; Kiseleva, 2016](#)) and game-theoretic ([Van Wijnbergen and Willems, 2015](#)) approaches, I am not attempting to prescribe an optimal emissions strategy or learning paths for climate sceptics under future uncertainty.³ Rather, my goal is to establish some ground rules for thinking about climate policy today, given the information that is already available to us.

Another distinguishing feature of this paper is that the results are derived via conceptually straightforward time-series regression analysis. While climate scientists have typically relied on complex computer models to simulate TCR, a growing body of research is aimed at understanding the link between human activities and climate change through the purview of time-series econometrics. Much of this literature has concerned itself with the apparent non-stationarity of climate data over time. Suffice it to say that the present paper takes as its foundation newer studies by [Gay-Garcia et al. \(2009\)](#) and [Estrada et al. \(2013a,b\)](#), who argue convincingly that global surface temperatures and anthropogenic forcings can best be described as trend-stationary processes, incorporat-

²There is a literature on Bayesian learning within IAMs that originates with [Kelly and Kolstad \(1999\)](#). Other prominent examples include [Leach \(2007\)](#) and [Lemoine and Traeger \(2014\)](#). However, I am not aware of any studies within this literature that attempt to model learning among climate sceptics.

³A more recent empirical paper by [Kaufmann et al. \(2017\)](#) shows that spatial heterogeneity in local climate change effects and temperatures can at least partially explain persistent scepticism in different regions of the United States. A related paper by [Moore \(2017\)](#) does not deal with sceptics *per se*, but characterises learning about climate as a (potentially) Bayesian process where individuals make inferences based on local weather shocks. This builds off of earlier work by [Deryugina \(2013\)](#), who finds that longer spells of abnormal local weather patterns are consistent with Bayesian updating about climate beliefs.

ing at least one structural change.⁴ The upshot is to permit the use of level terms within an ordinary least squares (OLS) regression framework. Such matters notwithstanding, virtually all econometric studies of climate change attribution to date have been carried out in the frequentist paradigm. They do not consider the influence of priors, nor are they able to yield the probabilistic estimates that are characteristic of Bayesian analysis. A noteworthy and early exception is that of [Tol and De Vos \(1998\)](#), who are motivated to adopt a Bayesian approach because of multicollinearity in their anthropogenic emissions data. Such multicollinearity does not plague newer datasets, which employ common measurement units in a way that allows for the aggregation of various forcing agents (see Section 2). This allows us to pin down the influence of specific forcings with greater confidence, and so obtain a more precise estimate of climate sensitivity. Furthermore, [Tol and De Vos](#) do not consider the influence of overtly contrarian priors as a basis for affecting policy.

2 Data

The various data sources for this paper are summarised in Table 1.⁵ Global mean surface temperature data (1850–2017) are taken from the HadCRUT4 dataset, jointly compiled by the UK Met Office and the Climatic Research Unit at the University of East Anglia. Two alternate global temperature reconstructions — one provided by [Cowtan and Way \(2014a,b\)](#) and the other by the NASA Goddard Institute for Space Studies (GISTEMP) — are used as a check against coverage issues and other uncertainties. Radiative forcing data, covering both historic estimates (1765–2005) and future scenarios (2006–2300), are taken from the Representative Concentration Pathway (RCP) database, hosted by the Potsdam Institute for Climate Impact Research. These data include anthropogenic sources of radiative forcing like industrial greenhouse gas emissions, as well as natural sources like solar irradiance and volcanic eruptions. It is important to note that the forcings are defined in terms of a common energy unit, Watts per square metre (Wm^{-2}). This allows for aggregation into a composite series of total

⁴Another group of researchers, beginning with [Stern and Kaufmann \(2000\)](#), has argued that the instrumental temperature record contains a stochastic trend that is imparted by (and therefore cointegrates with) the time-series data of radiative forcings. The reader is referred to [Estrada and Perron \(2013\)](#) for a useful overview of this debate.

⁵The compiled dataset and *R* code for reproducing all of the results in this paper can be found at the companion GitHub repository: <https://github.com/grantmcdermott/sceptic-priors>.

Table 1: Data sources

Variable	Product	Description	Period
GMST	HadCRUT4 ^a	Global mean surface temperature. Primary series. Compiled by the UK Met Office and the Climatic Research Unit at the University of East Anglia.	1850–2017
	CW2014 ^b	Secondary series. Compiled by Cowtan and Way (2014a,b) . Corrects for coverage bias in HadCRUT4.	1850–2017
	GISTEMP ^c	Secondary series. Compiled by the NASA Goddard Institute for Space Studies.	1880–2015
RF	RCP ^d	Total radiative forcing due to anthropogenic and natural factors (excluding volcanic aerosols). Compiled by Meinshausen et al. (2011) . Historical data until 2005, simulated scenarios thereafter.	1765–2300
AER	RCP ^d	Radiative forcing due to volcanic stratospheric aerosols. Compiled by Meinshausen et al. (2011) .	1750–2005
AMO	NOAA ^e	Atlantic Multidecadal Oscillation.	1856–2017
SOI	NCAR ^f	Southern Oscillation Index.	1866–2017

^a <http://www.metoffice.gov.uk/hadobs/hadcrut4/data/current/download.html>

^b <http://www-users.york.ac.uk/~kdc3/papers/coverage2013/series.html>

^c <http://data.giss.nasa.gov/gistemp/>

^d <http://www.pik-potsdam.de/~mmalte/rcps/>

^e <http://www.esrl.noaa.gov/psd/data/timeseries/AMO/>

^f <http://www.cgd.ucar.edu/cas/catalog/climind/soi.html>

radiative forcing, which circumvents the attributional problems that would otherwise arise due to severe multicollinearities in the various sources of anthropogenic forcing. Data for two major oceanic-atmospheric phenomena, the Atlantic Multidecadal Oscillation (AMO, 1856–2017) and the Southern Oscillation Index (SOI, 1866–2017), are taken from the U.S. National Oceanic and Atmospheric Administration (NOAA) and National Center for Atmospheric Research (NCAR). Summarising the common historic dataset for which data are available across all series, we have 140 annual observations running over 1866–2005. RCP scenarios until 2100 will also be considered for making future predictions later in the paper.

3 Econometric approach

3.1 Bayesian regression overview

The Bayesian regression framework is less familiar to many researchers than the frequentist paradigm that is commonly taught in universities. For this reason, I provide a brief overview of the key principles of the Bayesian method and highlight some important distinctions versus the frequentist approach. A more in-depth discussion may be found in [Koop \(2006\)](#), [Kruschke \(2014\)](#), and [Gelman et al. \(2014\)](#) among others.

A Bayesian regression model uses the logical structure of Bayes' theorem to estimate probable values of a set of parameters θ , given data X :

$$p(\theta|X) = \frac{p(X|\theta)p(\theta)}{p(X)}. \quad (1)$$

Here, $p(\theta|X)$ is known as the *posterior* and serves as the fundamental criterion of interest in the Bayesian framework. The posterior asks, "What are the probable values of our parameters, given the observed data?" This stands in direct contrast to the first term in the right-hand numerator, $p(X|\theta)$, which is the familiar *likelihood function* from frequentist statistics. The likelihood essentially reverses the question posed by the posterior and instead asks how likely we are to observe some data for a given set of parameters (e.g. based on an assumption about the data generating process). The second term in the numerator is the *prior*, $p(\theta)$. While the prior can take on any distributional form, it should in principle encapsulate our knowledge about the parameters before we have observed the data. Insofar as we are interested in learning about θ , it is common practice to ignore the term in the denominator, $p(X)$. This is simply the marginal probability of the data and can be thought of as a normalisation constant, which helps to ensure that the posterior is a proper probability distribution (i.e. integrates to one) and can be calculated *ad hoc* if needed. For this reason, equation (1) is typically re-written as

$$p(\theta|X) \propto p(X|\theta)p(\theta). \quad (2)$$

Equation (2) embodies the mantra of Bayesian statistics: "The posterior is proportional to the likelihood times the prior." Solving for the posterior typically involves the com-

combination of various integrals, which cannot be calculated analytically.⁶ Fortunately, the advent of increased computing power and software allow us to simulate the posterior density with relative ease using Markov Chain Monte Carlo (MCMC) routines. This can be done for virtually any combination of prior and likelihood function. Obtaining a valid posterior is then simply a matter of (i) choosing a prior distribution for our regression parameters, i.e. regression coefficients and variances, and (ii) specifying a likelihood function to fit the data. For ease of exposition — how we map parameter values to beliefs about TCR will be determined by the specification of the regression model — I begin with the likelihood function.

3.2 Likelihood function

The likelihood function is governed by the choice of empirical model. Following [Estrada and Perron \(2012\)](#); [Estrada et al. \(2013a\)](#), I model global temperatures using the regression equation

$$GMST_t = \alpha_0 + \beta_1 RF_t + \gamma_2 VOLC_t + \delta_3 SOI_t + \eta_4 AMO_t + \epsilon_t, \quad (3)$$

where $\epsilon_t = \phi\epsilon_{t-1} + \nu_t$ is a first-order autoregressive, or AR(1), error process.

Here, $GMST$ is the global mean surface temperature anomaly relative to the pre-industrial period (defined as the 1871–1900 average); RF is total radiative forcing due to both anthropogenic and natural factors (excluding volcanic eruptions); $VOLC$ is the radiative forcing due to volcanic stratospheric aerosols; and SOI and AMO are scaled indices of these respective climatic phenomena. The subscript t denotes time. Specifying that the error term ϵ follows an AR(1) process allows us to account for dynamic elements such as potential autocorrelation.

It is worth dwelling on the composite RF variable before continuing. This variable combines both anthropogenic *and* natural forcings. Upon first blush, such a choice may seem at odds with the goal of the paper, which is to separate out and interrogate scepticism about the human role in climate change. However, recall that the underlying forcings in my dataset are all expressed in terms of a common unit (i.e. Wm^{-2}). The

⁶A well-known exception occurs in the case of conjugate priors, i.e. prior distributions that are of the same family as the posterior distribution. However, this places strong restrictions on the questions that can be asked of the data.

model must therefore constrain these forcings to have the same effect on temperature, or else it risks becoming unphysical. The use of a single forcing series is thus a necessary step to ensure that the model remains physically consistent.⁷ Later in the paper, I will show in an alternate specification that relaxing this constraint nonetheless leads to near identical conclusions as the physically correct specification (see Section 4.2) I will also show that the core results do not hinge on the imposition of a common efficacy among different forcing agents.

Returning to my primary regression model, equation (3) implies a likelihood function that is multivariate normal,

$$p(GMST|\boldsymbol{\beta}, \sigma^2) = \frac{1}{(2\pi\sigma^2)^{T/2}} \exp \left[-\frac{(GMST - \mathbf{X}\boldsymbol{\beta})'(GMST - \mathbf{X}\boldsymbol{\beta})}{2\sigma^2} \right], \quad (4)$$

where \mathbf{X} is the design matrix of explanatory variables; $\boldsymbol{\beta}$ is the coefficient vector; $\sigma^2 = \text{Var}(\epsilon)$ is the variance of the error term; and $T = 140$ is the observed number of time periods. Equation (4) can also be written more simply as $GMST|\boldsymbol{\beta}, \sigma^2 \sim \mathcal{N}_T(\mathbf{X}\boldsymbol{\beta}, \sigma^2\mathbf{I})$.

An important feature of equations (3) and (4) is that they define how we should map probabilities about the regression parameters to beliefs about climate sensitivity. Recall that TCR describes the contemporaneous change in temperature that will accompany a steady doubling of atmospheric CO₂ concentrations. It follows trivially that

$$\text{TCR} = \beta_1 * F_{2\times}, \quad (5)$$

where β_1 is the regression coefficient describing how responsive global temperatures are to a change in total radiative forcing; and $F_{2\times}$ is the change in forcing that results from a doubling of CO₂. For the latter, I use the IPCC's best estimate of $F_{2\times} = 3.71 \text{ Wm}^{-2}$ and further assume an additional $\pm 10\%$ variation to account for uncertainties over spatial heterogeneity and cloud formation (Schmidt, 2007; IPCC, 2001, Chapter 6).⁸ The key point is that assigning a distribution over the parameter β_1 will necessarily imply a distribution for TCR, and vice versa. We therefore have a direct means of

⁷Volcanic aerosols are an exception because they impart only a transitory level of forcing. This explains why *VOLC* may be included as a separate component in the regression equation (Estrada et al., 2013a).

⁸It is worth noting that a number of studies which rely on time-series methods to derive an estimate of climate sensitivity — e.g. Kaufmann et al. (2006); Mills (2009); Estrada and Perron (2012) — do so under the assumption that $F_{2\times} = 4.37 \text{ Wm}^{-2}$. This outdated figure appears to be based on early calculations by Hansen et al. (1988). The climate sensitivity estimates of these studies may consequently be regarded as inflated.

linking prior and posterior probabilities of the regression parameters to beliefs about TCR. It also means that the primary goal of the regression analysis will be to determine probable values of β_1 . The rest of the parameters will take a backseat in the analysis that follows, acting largely as controls.

3.3 Priors

Climate scepticism is a matter of degree. I account for this fact by defining a simple typology of sceptics as per Table 2. I assume that these sceptics have priors on TCR that are normally distributed. Each sceptic “type” is then distinguished by the parameters of their normal distribution, i.e. by a specific mean and variance combination.

- The prior **mean** represents a type’s best guess for the true value of TCR. I distinguish between two basic cases. A *lukewarmer’s* prior for TCR is centered around a 1 °C mean, while a *denier’s* prior is centered around a 0 °C mean. A lukewarmer believes that humans are having an effect on climate, but only at the very low range of the scientific consensus (as measured by the IPCC’s “likely” range). A denier, on the other hand, believes that there is probably no relationship between human activity and climate change.
- The prior **variance** represents the certainty or confidence that each type has in their belief. I again distinguish between two basic cases. A *moderate* level of certainty corresponds to a prior distribution where the 95% probability interval extends over a full degree centigrade. In contrast, a *strong* level of certainty would see that interval shrink to just a quarter of a degree.

Following equation (5), obtaining priors over β_1 is a simple matter of dividing the respective TCR distributions by $F_{2\times} = 3.71 \text{ Wm}^{-2}$. These are the parameters that actually enter the Bayesian regression model and are also shown in Table 2.

In addition to the subjective priors of our stylised sceptics, a noninformative prior provides a useful reference case for the analysis. Loosely speaking, noninformative priors are vague and do not privilege any parameter values before observing the data. Following standard practice, I use a normal distribution with large variance to represent noninformative priors. Experimenting with other options shows that this choice has a negligible impact on my results.

Table 2: Sceptic priors

Type	TCR ($^{\circ}\text{C}$)	Implied β_1
Moderate lukewarmer	$\mathcal{N}(1, 0.250^2)$	$\mathcal{N}(0.27, 0.0674^2)$
Strong lukewarmer	$\mathcal{N}(1, 0.065^2)$	$\mathcal{N}(0.27, 0.0175^2)$
Moderate denier	$\mathcal{N}(0, 0.250^2)$	$\mathcal{N}(0.00, 0.0674^2)$
Strong denier	$\mathcal{N}(0, 0.065^2)$	$\mathcal{N}(0.00, 0.0175^2)$

Priors types are defined according to the mean and variance parameters of normal distributions over TCR. *Lukewarmers* believe that TCR is distributed around a 1°C mean (the lower bound of the IPCC’s “likely” range), while *Deniers* believe that TCR is distributed around a zero mean. Similarly, the variance describes uncertainty so that a person with “moderate” convictions believes that the true value of TCR lies within a 1°C interval of their prior mean (95% probability), while that interval falls to just 0.25°C for someone with “strong” convictions. The implied priors for β_1 are obtained using the simple formula described in equation (5), i.e. $\beta_1 = \text{TCR}/3.71$.

Two points merit further discussion before continuing on to the posterior results. The first is that our group of sceptics only hold subjective priors about TCR and, thus, β_1 . Noninformative priors are always assumed for the remaining parameters in the regression equation. The second point is to acknowledge that these sceptics are, of course, highly stylised caricatures. Their priors are simply taken as given. We are not concerned with where these priors come from and why they are of a particular strength. However, such abstractions are ultimately unimportant given the objectives of this study. My goal is to explore how climate sceptics would respond to evidence for climate change, provided that they update their beliefs rationally. Moreover, it gives a sense of just how strong someone’s prior beliefs need to be, so as to preclude the acceptance of any policy interventions.

4 Results

4.1 Regression results and updated TCR beliefs

The posterior regression results for the various prior types are presented in Table 3. Beginning with the noninformative case in the first column, all of the regression coefficients are credibly different from zero and of the anticipated sign. For example, GMST is negatively correlated with SOI. This is to be expected since the El Niño phenomenon

Table 3: Posterior regression results and implied TCRDependent variable: Global Mean Surface Temperature anomaly ($^{\circ}\text{C}$)

	Noninformative	Lukewarmer		Denier	
		Moderate	Strong	Moderate	Strong
Total radiative forcing	0.419 [0.393, 0.444]	0.413 [0.388, 0.438]	0.360 [0.336, 0.383]	0.403 [0.376, 0.428]	0.108 [0.069, 0.146]
Volcanic aerosols	0.051 [0.008, 0.094]	0.051 [0.008, 0.095]	0.051 [0.005, 0.097]	0.051 [0.008, 0.094]	0.050 [-0.048, 0.147]
SOI	-0.028 [-0.039, -0.016]	-0.028 [-0.040, -0.016]	-0.032 [-0.045, -0.019]	-0.029 [-0.041, -0.017]	-0.050 [-0.077, -0.023]
AMO	0.473 [0.399, 0.547]	0.473 [0.399, 0.546]	0.467 [0.387, 0.547]	0.472 [0.395, 0.547]	0.443 [0.274, 0.613]
Constant	-0.102 [-0.124, -0.080]	-0.098 [-0.120, -0.077]	-0.063 [-0.083, -0.042]	-0.092 [-0.113, -0.07]	0.107 [0.069, 0.148]
Implied TCR	1.6 [1.4, 1.7]	1.5 [1.4, 1.7]	1.3 [1.2, 1.5]	1.5 [1.3, 1.7]	0.4 [0.3, 0.6]

Mean values are given, with square brackets denoting 95% Bayesian credible intervals. Columns are distinguished by different sets of prior covariates. Column (1) specifies noninformative priors over all regression parameters. Columns (2)–(5) specify subjective priors over the coefficient on total radiative forcing (see Table 2), with noninformative priors specified over the remaining parameters. Total radiative forcing and volcanic stratospheric aerosols are measured in Wm^{-2} . The Southern Oscillation Index (SOI) and Atlantic Multidecadal Oscillation (AMO) are measured as scaled indices. The implied TCR values at the bottom of the table are measured in $^{\circ}\text{C}$ and are obtained by multiplying the coefficient on total radiative forcing by $F_{2\times}$ per equation (5). The dataset consists of 140 annual observations over the period 1866–2005.

is defined by SOI moving into its negative phase. The posterior density of our main parameter of interest, the coefficient on RF , shows that global temperature will rise by an average of 0.4°C for every Wm^{-2} increase in total radiative forcing. Of greater interest, however, is the fact that the posterior estimates yielded by the group of sceptic priors are very similar to this noninformative case. With the exception of the Strong Denier, there is a clear tendency to congregate towards the noninformative parameter values.

Of course, the exact values of the regression parameters are themselves of somewhat limited interest. Rather, their primary usefulness is to enable the recovery of posterior beliefs about TCR. These are summarised at the bottom of Table 3, while the full prior and posterior distributions are plotted in Figure 1. We see that the posterior TCR distributions are generally clustered around a best estimate of 1.5°C , with a 95% credible interval somewhere in the region of $1.2\text{--}1.7^{\circ}\text{C}$, depending on the prior. Excepting the Strong Denier, these posterior beliefs about TCR fall comfortably within the IPCC “likely” range. However, the derived probability intervals are decidedly narrower and

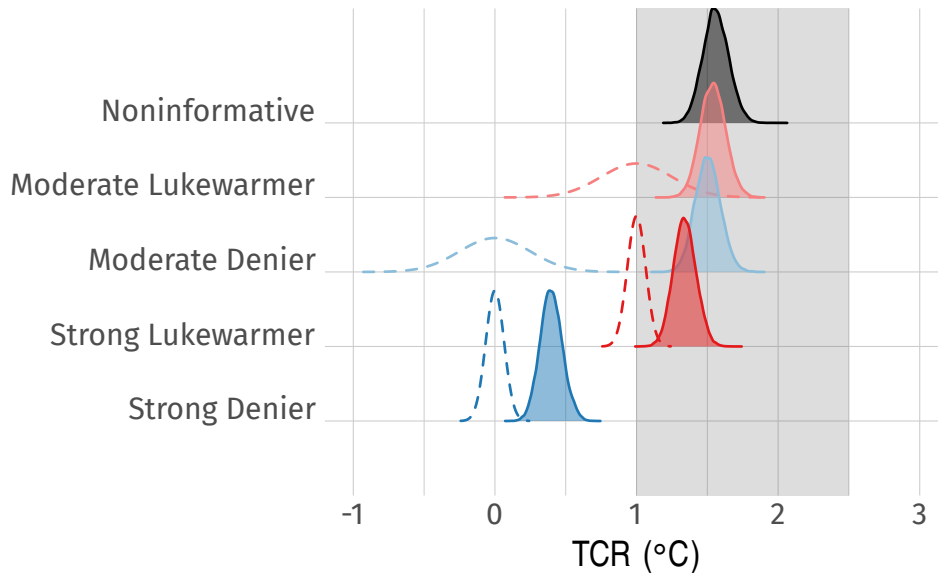


Figure 1: TCR densities

Notes: Dashed lines denote priors, solid lines denote posteriors. The grey shaded region denotes the IPCC “likely” TCR range of 1.0–2.5 °C.

TCR values at the upper end of the spectrum are discounted accordingly.

Further insight into the updating behaviour of our stylised sceptics is provided by the recursive TCR estimates shown in Figure 2. It is apparent that stronger convictions about one’s prior beliefs (in the form of a smaller prior variance) have a greater dampening effect on posterior outcomes than the prior mean. For example, the Moderate Denier converges more rapidly to the noninformative distribution than the Strong Lukewarmer. However, most sceptics will converge to the noninformative distribution only after “observing” data from a number of decades. And, yet, this does not alter the conclusions that we are able to draw from our Bayesian analysis. For as long as we have fully specified a prior that encapsulates a person’s initial beliefs, then we should in principle treat the full historical dataset as new information for updating those beliefs.⁹ But it does highlight the importance of using all the available instrumental climate data for

⁹As a corollary, concerns over the use of the full historical dataset would only hold sway in cases where priors already incorporate information that has been obtained from applying the same model on a sub-sample of the dataset. In that case, we would need to exclude the sub-sample from the analysis to derive a valid posterior that avoids double counting.

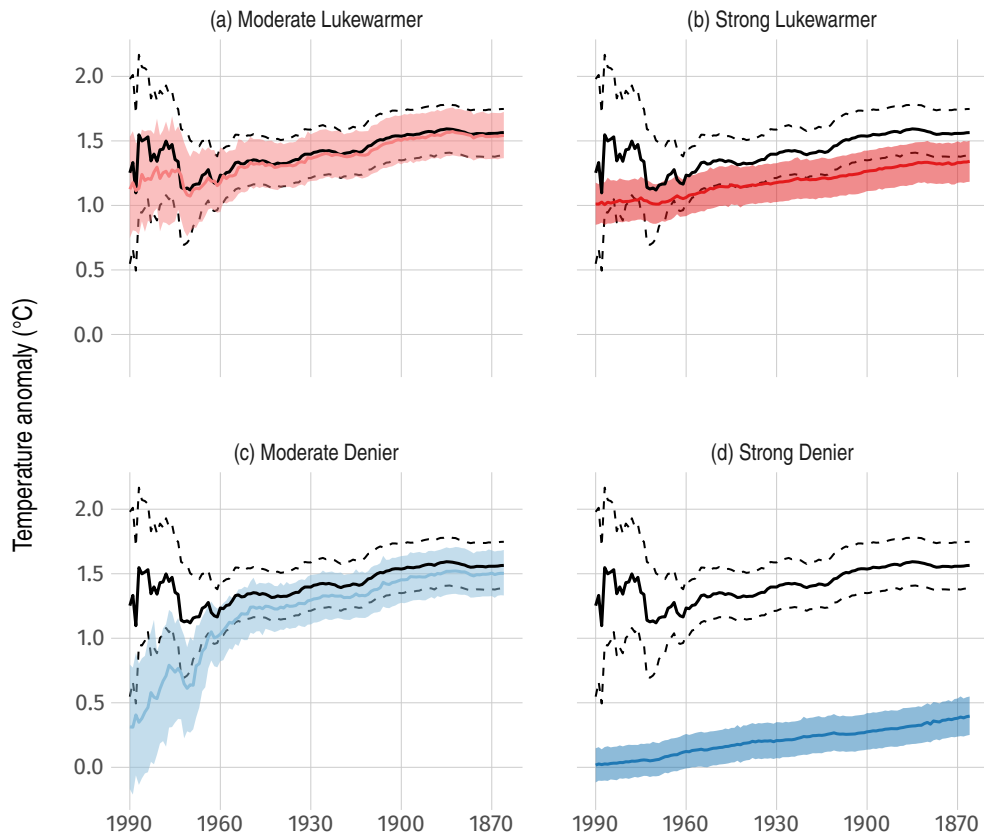


Figure 2: Recursive TCR estimates

Notes: Solid lines denote means, shaded regions (or dashed lines) denote 95% credible intervals. Recursive estimates are obtained in reverse chronological order, i.e. starting nearest to the present and moving backwards in time. In each panel, the posterior TCR estimate resulting from a sceptical prior is contrasted with the noninformative case (in black).

building any kind of policy consensus. Limiting the sample period under observation to, say, the last 35 years would largely preclude the possibility of consensus formation. The tendency of some prominent sceptics to rely on satellite records of global temperatures — which only stretch back as far as 1979 — could be seen as anecdotal evidence in support of this claim (e.g. [Mooney, 2016](#)).

Returning, then, to the question posed at the beginning of this paper: How much evidence would it take to convince climate sceptics that they are wrong about global warming? One way to reframe this question is to think about how much data a sceptic needs to observe before their best estimate of climate sensitivity begins to look reasonable to a




Figure 3: When do sceptic beliefs about TCR converge with mean posterior targets?

Notes: Axes denote the means and standard deviations of a range of normally-distributed sceptic priors about TCR. Convergence occurs when the mean posterior TCR for a particular prior equals the relevant target value, either (a) 1.3 °C or (b) 1.5 °C. The year of convergence assumes a starting date of 1866 to coincide with the common historical dataset used in this study. Blue shading indicates that convergence is feasible with historically available data. Red shading indicates that convergence can only occur once additional data has been accumulated in the future. Climate data are simulated according to the parameters obtained from the noninformative Bayesian regression in Table 3.

mainstream climate scientist. For example, how long would it take before they obtained a mean posterior TCR of 1.3 °C or 1.5 °C? While it is possible to look at the sceptics' recursive TCR estimates using only historical data, we run into problems with the more extreme priors. In short, there is simply not enough historical data to overcome higher orders of scepticism. I therefore simulate over 200 years' worth of global temperature and climate data using parameters obtained from the noninformative Bayesian regression in Table 3. I then use this simulated data to run a set of secondary regressions that are distinguished by a range of different sceptic priors on TCR. (This range is much more granular than my original four-sceptic typology.) Each regression is estimated recursively, incrementing forwards one year at a time, until I obtain a posterior TCR distribution that has a mean value equal to the relevant target.

The results are shown in Figure 3. While the instrumental climate record constitutes enough data to convince many sceptics in this hypothetical pool, it does not suffice in all cases. Similarly, although we expect that many present-day sceptics will eventually

acquiesce their beliefs if climate change continues into the future, there remains a small group of hardcore sceptics who defiantly reject convergence with the mainstream view — even if we project as far ahead as 2100. Such is the strength of their priors. Note further that the year of convergence is a non-linear function of prior strength, so that it becomes increasingly difficult to convince the marginal sceptic. The steady accumulation of evidence over time will inexorably bring more sceptics into the mainstream fold. But the delay between each round of new converts is increasing.

An implication of this thought experiment is as follows. If a someone is unconvinced of the human influence on climate today — despite all of the available evidence — then there is a high probability that they will remain unconvinced for many years hence. Such extreme cases are perhaps in the minority of the population. However, it is striking to think that these individuals are already out of reach from the perspective of supporting comprehensive climate policy. Even the accumulation of evidence over the next several decades may not be enough to convince them.

4.2 Robustness checks

I consider a number of alternative specifications to demonstrate the robustness of my results. These are summarized in Table 4. For example, the HadCRUT4 dataset is known to suffer from potential coverage biases due to incomplete placement of *in situ* thermometers. I therefore rerun the analysis with two alternate reconstructions of GMST. [Cowtan and Way \(2014a,b\)](#), hereafter CW2014, correct for the gaps in the HadCRUT4 dataset by using an interpolation algorithm based on the “kriging” method. Similarly, the NASA Goddard Institute for Space Studies uses an extrapolation algorithm to overcome coverage bias in GISTEMP, its own GMST reconstruction. Running the Bayesian regression model on these alternative series yields moderately higher TCR values compared to HadCRUT4. Under a noninformative prior, the posterior TCR means (and 95% Bayesian credible intervals) are 1.6 °C (1.4–1.8 °C) for CW2014 and 1.7 °C (1.5–1.9 °C) for GISTEMP. While I omit them for brevity, the posterior results for the group of climate sceptics are similarly nudged higher towards the new noninformative distributions. Given that the explicit goal of this paper is to evaluate policy options from the perspective of climate sceptics, I continue using the results from the HadCRUT4 series as a default. Yet, it should be noted that this is a conservative choice that may, at least marginally, understate the true level of warming.

Table 4: TCR: Alternative specifications and robustness checks

Key	Mean (°C)	95% C.I.	Comment
HadCRUT4	1.6	[1.4, 1.7]	Primary GMST series. For comparison only.
CW2014	1.6	[1.4, 1.8]	Alternative GMST series.
GISTEMP	1.7	[1.5, 1.9]	Alternative GMST series.
Meas. error	1.5	[1.4, 1.7]	Specifying measurement error in HadCRUT4.
Marvel I	2.1	[1.9, 2.4]	Adjusted forcing efficacies (means).
Marvel II	1.9	[0.4, 3.4]	Adjusted forcing efficacies (distributions).
Anthro [†]	1.5	[1.4, 1.8]	Separate anthropogenic from natural forcings.
CO ₂ [†]	1.7	[1.4, 2.0]	Separate CO ₂ from all other forcings.

All estimates are computed using noninformative priors. See main text for additional details.

[†] See Table SM2 for a summary of the posterior TCR values across the different sceptic prior types.

All three GMST reconstructions also provide estimates of measurement error. The Bayesian framework is ideally suited to incorporate such knowledge, since the nested model structure allows us to fully specify measurement error on the dependent variable within the regression model itself. Doing so in the present setup yields coefficient estimates that are effectively identical to those presented in Table 3. This is unsurprising once we recall that measurement error on the dependent variable is absorbed by the disturbance term of the regression model.¹⁰ Since the Bayesian regression framework is primarily concerned with total model uncertainty, specifying the relative contribution of such measurement error to the overall disturbance doesn’t meaningfully alter the analysis — though it may be useful for incorporating known sources of heteroscedasticity.¹¹ The primary regression results already have the measurement error “baked in” to the estimation, regardless of whether we define it explicitly or not.

Thus far, I have assumed that the different physical drivers that make up total radiative forcing have the same per-unit effect on GMST. Forcing agents that yield a similar radiative imbalance in Wm^{-2} are expected to result in similar feedbacks and responses in GMST. However, recent research has suggested that the warming efficacy of different forcing agents can, in fact, vary with factors like geography. Aerosol emissions, for example, are primarily concentrated in the mid-to-high latitudes of the Northern Hemisphere. The disproportionately large land mass in this region causes aerosol forcing

¹⁰For example, see [Greene \(2007, p. 326\)](#). To illustrate with a simple univariate case: The regression model can be written as $y_t \sim \mathcal{N}(\beta X_t, \sigma^2 + \omega_t^2)$, where $\sigma^2 = \text{Var}(\epsilon)$ is the variance of the error term and $\omega_t^2 = \text{Var}(v_t)$ is the variance of the measurement error on y_t . Together, ϵ and v_t make up the overall disturbance of the regression.

¹¹See [Lewis and Linzer \(2005\)](#) for a related discussion in a frequentist setting.

to exhibit stronger feedback effects and an accelerated temperature response than if it were uniformly distributed across the globe (Shindell, 2014). The implications of such forcing inhomogeneity on climate sensitivity estimates are more fully explored by Marvel et al. (2016). The authors use a controlled set of climate model runs to estimate the relative efficacies of individual forcings on both TCR and ECS. I adapt their results to construct an adjusted series of total radiative forcing, where various inputs have been pre-multiplied by the appropriate efficacy coefficients (see Table SM1). More specifically, I consider two approaches. The first takes Marvel et al.'s mean efficacy estimates as given and ignores all distributional uncertainty when re-running the main Bayesian regression. The second accounts for uncertainty in a Monte Carlo framework, repeatedly sampling from the t distributions that describe each efficacy estimate and then combining the posterior results from many regressions into a single meta-distribution at the end.¹² Consistent with Marvel et al.'s findings, both approaches lead to a pronounced increase in the posterior TCR mean. Note, however, that the Monte Carlo sampling approach yields significantly larger posterior credible intervals. That being said, the large distributional uncertainties in Marvel et al.'s efficacy estimates for certain forcing agents (e.g. land use) result primarily from the small changes that these drivers have experienced over the historical period. Furthermore, the authors obtain their estimates using a single climate model rather than a multi-model ensemble. The wider TCR uncertainty bounds resulting from my Monte Carlo approach should hence be regarded with some caution.

As final robust check, I relax the constraint that all sources of radiative forcing have to be included in the regression model under the same composite RF term. Recall that this decision was motivated by the fact that the forcing agents in my dataset are all defined in Wm^{-2} . Separating out individual forcings and then placing different priors on them thus causes the model to become physically inconsistent. Such admonishments notwithstanding, I implement two version of this unphysical model. The first separates out anthropogenic forcings (e.g. GHGs) from natural forcings (e.g. solar radiation). The second separates out CO_2 forcing from all other sources. In both cases, the subjective priors from Table 2 are placed on the isolated anthropogenic component, while all other variables take noninformative priors. Both sets of regressions yield very similar results to the main, physically-correct specification. If anything, isolating CO_2

¹²This probabilistic approach is often used in Bayesian regression models to account for measurement error in explanatory variables. Indeed, deriving consistent estimators in the presence of explanatory variable measurement error is much more problematic under the frequentist framework (Greene, 2007).

on its own yields a higher posterior TCR for certain prior types. However, this latter implementation should be treated with double caution as it likely suffers from multicollinearity problems in addition to physical inconsistency.¹³ The full posterior TCR results for these two models — for all prior types — are presented in Table SM2.

4.3 Future temperature predictions

Climate policy is largely predicated upon the risks to future generations. As such, any policy discussion must consider predictions that run many years into the future. TCR estimates are one means of gaining an insight into how global temperatures will evolve over the coming decades. A more explicit way of demonstrating this is by predicting temperatures until the end of the century. While the trajectory of future radiative forcings is subject to much uncertainty, some guidance is available in the form of the IPCC's Representative Concentration Pathways (Van Vuuren et al., 2011). These so-called "RCPs" describe a family of emissions scenarios, where total anthropogenic forcings evolves according to various economic, demographic and technological assumptions. Each RCP includes a core component of atmospheric CO₂ concentrations, while they all share a common prediction for radiative forcing due to solar activity. I take these series as the basis for constructing covariate vectors to predict temperatures until the year 2100. For the remaining explanatory variables — stratospheric aerosols, SOI and AMO — I take the mean historical values from my dataset. A summary of covariate vectors in 2100 for each RCP scenario is provided in Table SM3.

Figure 4 shows the temperature evolution for each RCP under the noninformative case, which we again take as the benchmark. The principal message is that CO₂ concentrations must be constrained to RCP 4.5 or lower, if we are to avoid a 2 °C rise in global temperature. Given the prominence of this particular threshold in international climate treaties and the popular narrative, the result is a reinforcement of commonly cited emissions targets such as 450 and 540 ppmv. On the other hand, we can expect to breach even 3 °C by the year 2100 if we continue along a truly unconstrained emissions path à la RCP 8.5.

What of the predictions yielded by our group of climate sceptics? While it is straightfor-

¹³In particular, other anthropogenic forcings such as CH₄ and N₂O, follow very similar trends over time. The model will therefore struggle to correctly attribute warming across these different sources.

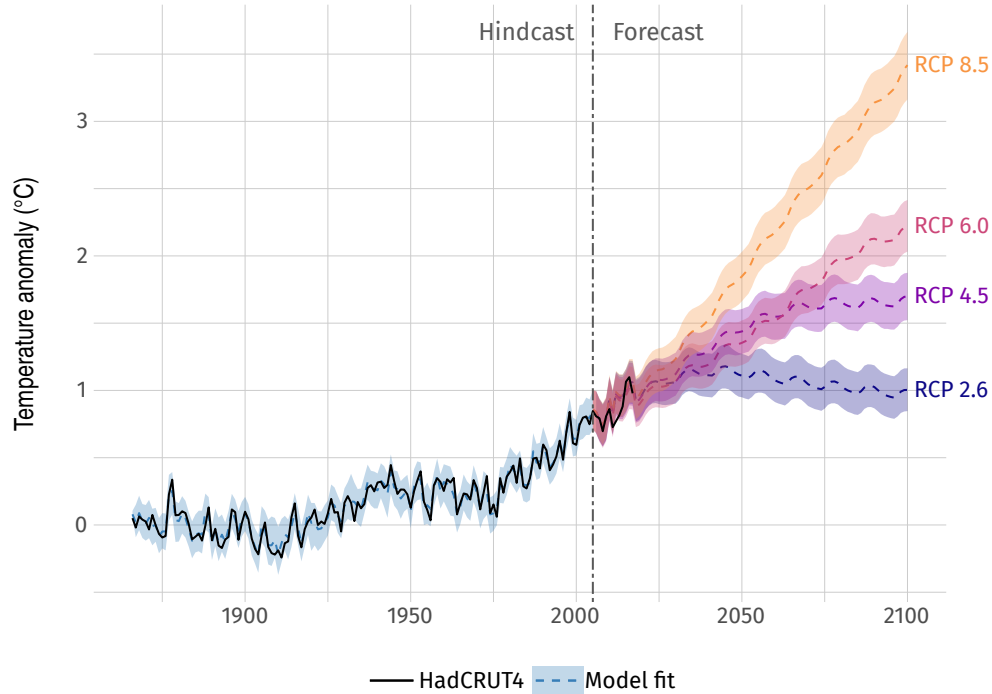


Figure 4: Model fit and prediction: Noninformative priors

Notes: Temperature anomaly relative to the 1871–1900 average. Shaded regions denote 95% credible intervals.

ward to redraw Figure 4 for each prior type, a more intuitive comparison can be made by looking at the total warming that each person expects by the end of the century. Figure 5 plots the predictive temperature density in the year 2100 for all prior types by RCP scenario. Again, the data have a clear tendency to overwhelm even reasonably staunch forms of climate scepticism. Nearly all of the stylised sceptics would expect to breach the 2 °C threshold by 2100 under RCP 6.0, while a temperature rise of more than 3 °C is likely under under RCP 8.5. An exception can only be found in the form of the Strong Denier, whose extreme prior dominates the posterior in a way that obviates nearly all concern about large temperature increases.

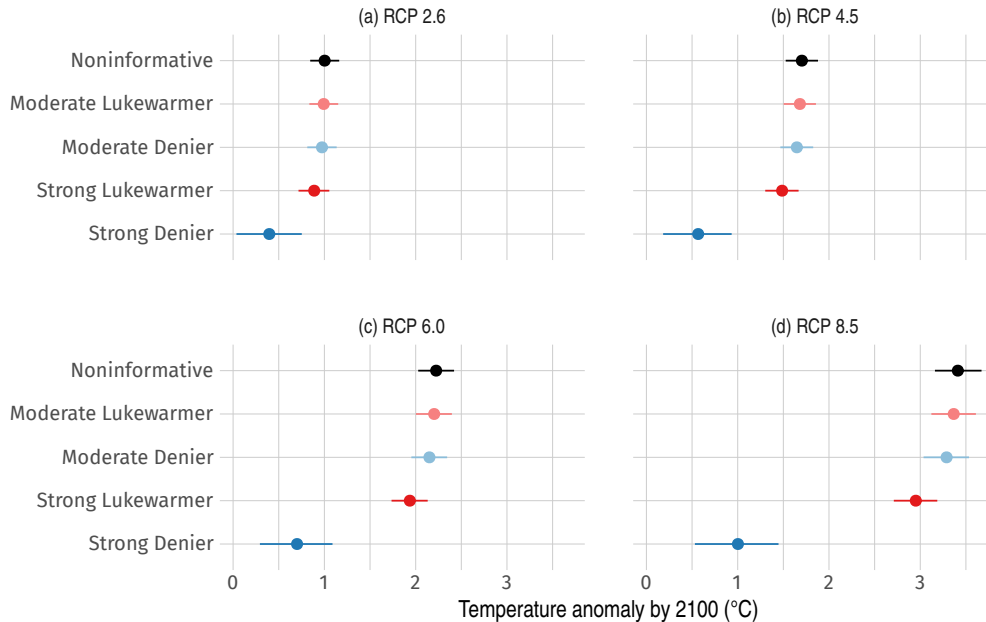


Figure 5: Predicted temperature anomaly by 2100: All priors

4.4 Welfare implications and the social cost of carbon

Provided they consider enough data, we have seen that most climate sceptics should be able to agree that a 2 °C target requires limiting CO₂ concentrations to around 540 ppmv. However, whether someone actually subscribes to policy measures aimed at achieving the 2 °C goal is dependent on many things; their choice of discount rate, beliefs about the efficacy of policy, damage expectations, to name but a few. Such issues are largely beyond the scope of this paper. Nonetheless, we may still gain a deeper insight into the welfare implications of our posterior TCR values by analysing their effect on the social cost of carbon (SCC). The SCC represents the economic costs associated with a marginal unit of CO₂ emissions. It can therefore be thought of as society’s willingness to pay for the prevention of future damages associated with human-induced climate change.

Obtaining SCC estimates generally requires the use of integrated assessment models (IAMs), which are able to solve for optimal climate policy along a dynamic path by simulating across economic and climate systems. The PAGE09 model developed by [Hope \(2011a\)](#) is ideally suited to our present needs. It is widely used as one of the major IAMs for evaluating climate policy ([IWG, 2010, 2013](#); [Nordhaus, 2014](#)). At a more

Table 5: Social cost of carbon (US\$2005 per tonne)

	Mean	Median	95% Probability interval
Noninformative	78.97	40.54	[11.74, 168.03]
Moderate Lukewarmer	64.86	39.32	[11.15, 160.45]
Strong Lukewarmer	47.88	26.31	[6.80, 106.97]
Moderate Denier	67.72	36.03	[10.34, 148.55]
Strong Denier	1.76	1.47	[−0.17, 5.35]

Results for each agent type are obtained from 10,000 simulation runs of PAGE09 (Hope, 2011a). Posterior TCR distributions serve as key inputs to the model, while the remaining parameters are set to the PAGE09 model defaults.

specific level, PAGE09 accepts random variables as inputs and yields the type of probabilistic output that is consistent with the rest of this paper. I take the posterior TCR distributions yielded by our Bayesian regression model and use these as inputs for calculating the SCC. The PAGE09 defaults are used for the remaining parameters (Hope, 2011a,b).

Table 5, together with Figure SM1 in the Supplementary Material, present probability distributions for the SCC across all prior groups in 2005 US dollars. The resulting distributions are highly skewed and characterised by extremely long upper tails. This is partly due to the fact that PAGE09, like most IAMs, assumes that economic damages are convex in temperature. More importantly, however, PAGE09 allows for the possibility of large-scale discontinuities (e.g. melting of the Greenland ice sheet) at temperatures above 3 °C. Such low probability, high impact events would yield tremendous economic losses and result in some extreme SCC values as a consequence (Hope, 2011a,b). The frequency of these events are more common in my adapted version of page, since I replace the default triangular (i.e. bounded) TCR distribution with distributions that are approximately normally distributed (thus allowing weight in the tails). For this reason, I provide both the mean and median SCC values alongside the 95% probability interval.

Excepting the Strong Denier, the SCC for all prior types is comfortably larger than zero. The mean value ranges from \$40 to \$71 per tonne (2005 prices), while the 95% probability interval extends from around \$10 to upwards of \$100 per tonne. These results are consistent with the SCC estimates found within the literature. For example, an influential synthesis review conducted by the United States government under the Obama administration established a mean SCC value of \$11–\$52 per tonne (2007 prices), de-

pending on the preferred discount rate ([IWG, 2010, 2013](#)). The encouraging point from a policy perspective is that such congruency exists despite the fact that the analysis proceeds from an initial position of scepticism. Another way to frame the SCC estimates presented here is to imagine that each prior type represents an equal segment of a voting population. We would then expect to see broad support for a carbon tax of at least \$20–\$25. While such a thought experiment clearly abstracts from the many complications that would arise from free-riding and so forth, again we see that nominal climate scepticism does not correspond to a mechanical dismissal of climate policy.

5 Discussion

We have seen that a non-trivial carbon price is consistent with a range of contrarian priors once we allow for updating of beliefs and, crucially, consider enough of the available data. An optimist might interpret these findings as a sign that common ground on climate policy is closer than many people think. On the other hand, they may also help to explain why the policy debate is so polarised in the first place. As all intermediate positions are absorbed into the mainstream, only the most hardcore sceptics will remain wedded to their priors. Such a group is unlikely to brook any proposals for reduced carbon emissions and virtually no amount of new information will convince them otherwise. Taken together with the persistent scepticism that one sees in actual polling data (e.g. [Saad, 2019](#)), it then becomes reasonable to ask whether real-life climate sceptics hold such extreme views? For that matter, are they numerous enough to prevent political action? Such considerations are reinforced by the idealized nature of the analysis until now. Irrespective of the scientific merit of working through such a set-up, normal people clearly do not update their priors in lockstep with a Bayesian regression model, supported by large dataset of time-series observations.

A natural starting point for thinking about these issues is to take a closer look at the mechanisms underlying posterior agreement formation. The notion that partisans should converge toward consensus with increasing information has long been taken as a logical consequence of Bayes' theorem. Indeed, empirical evidence to the contrary has been cited as a weakness of the Bayesian paradigm and its relevance to real-life problems (e.g. [Kahneman and Tversky, 1972](#)). This is a misconception. Nothing in the Bayesian paradigm precludes the possibility of diverging opinions in the face of shared informa-

tion (Jaynes, 2003; Bullock, 2009). It may even be the case that the same information has a polarising effect on individuals, pushing them towards opposite conclusions. This is perhaps most easily shown by incorporating perceptions of trust and source credibility into our Bayesian model. That is to say, we must broaden our conception of the *prior* so that it describes not only our existing beliefs about some phenomenon S , but also the credibility that we assign to different sources of information about S .

Consider an example, which is closely adapted from a related discussion in Jaynes (2003). Al, Bob and Christie hold different beliefs about climate change. Al is a “warmist”, Bob is a “lukewarmer” and Christie is a “denier”. These labels are encapsulated by the prior probabilities that each person assigns to climate sensitivity S , which we assume for simplicity is either high or low: $S \in S_L, S_H$. Denote by I an individual’s prior information about the world. Then, indexing by the first letter of their names, we summarise their prior beliefs about climate change as the following probabilities: $P(S_H|I_A) = 0.90$, $P(S_H|I_B) = 0.40$, and $P(S_H|I_C) = 0.10$.

Suppose that the IPCC now publishes its latest assessment report, wherein it claims that climate sensitivity is high. How do Al, Bob and Christie respond to this new data, $D = D_H$? It turns out that the answer depends on the regard that each individual holds for the IPCC itself. For example, let us say that all three individuals agree the IPCC would undoubtedly present data supporting a high climate sensitivity if that were indeed the true state of the world, i.e. $P(D_H|S_H, I_A) = P(D_H|S_H, I_B) = P(D_H|S_H, I_C) = 1.00$. However, they disagree on whether the IPCC can be trusted to disavow the high sensitivity hypothesis if the scientific evidence actually supported a low climate sensitivity. Despite their different beliefs about climate sensitivity, Al and Christie both regard the IPCC as an upstanding institution that can be trusted to accurately represent the science on climate change. In contrast, Bob is dubious about the motives of the IPCC and believes that the organisation is willing to lie in advancement of a preconceived agenda. Representing these beliefs in terms of probabilities, we have $P(D_H|S_L, I_A) = 0.05$, $P(D_H|S_L, I_B) = 0.89$, and $P(D_H|S_L, I_C) = 0.05$.

Recovering the posterior beliefs about climate sensitivity for our three individuals is now a simple matter of modifying Bayes’ theorem to account for each person’s relative

trust in the IPCC. For Al, we have

$$\begin{aligned}
 P(S_H|D_H, I_A) &= \frac{P(D_H|S_H, I_A)P(S_H|I_A)}{P(D_H|S_H, I_A)P(S_H|I_A) + P(D_H|S_L, I_A)P(S_L|I_A)} \\
 &= \frac{1.0 \times 0.9}{1.0 \times 0.9 + 0.05 \times 0.1} \\
 &\approx 0.98.
 \end{aligned}$$

Similarly, we obtain $P(S_H|D_H, I_B) \approx 0.43$ for Bob and $P(S_H|D_H, I_C) \approx 0.69$ for Christie.

Taking a step back, Al is now even more of a believer in the high sensitivity hypothesis than, having raised his subjective probability for S_H from 90% to 98%. Christie has experienced a still greater effect and has updated her subjective probability for S_H from 10% to 69%. She now attaches a larger probability to the high sensitivity hypothesis than the low sensitivity alternative. However, the same cannot be said of Bob, who has not been swayed by the IPCC report in the slightest. Both his prior and posterior probabilities suggest that S_H only has an approximately 40% chance of being true. Bob’s extreme mistrust has effectively led him to discount the IPCC’s high sensitivity claim in its entirety.

Extending the above framework to account for increasing granularity is conceptually straightforward. The principal insight remains the same: Trust is as much a determinant of whether beliefs are amenable to data — and whether individuals converge towards consensus — as the precision of the data itself. Such an extension seems especially relevant to the climate change debate given the sense of scientific distrust that pervades certain segments of society (Malka et al., 2009; Gauchat, 2012; Leiserowitz et al., 2013; Fiske and Dupree, 2014; Hmielowski et al., 2014). Indeed, recent research supports the notion that distrust of scientists is causing belief polarization about climate change in some demographic groups, even as scientific evidence may increase consensus in others (Cook and Lewandowsky, 2016; Zhou, 2016). Similar “backfire” effects have been well documented in other fields (Nyhan and Reifler, 2010; Harris et al., 2015).

Perhaps the most important feature of generalising the Bayesian framework in this way is that it offers a bridge between competing explanations of climate scepticism as a social phenomenon. Whereas the “deficit model” posits a lack of scientific knowledge and understanding as key drivers of scepticism, advocates of the “cultural cognition” theory argue that group identity and value systems are more relevant (Kahan et al., 2011, 2012; Ranney and Clark, 2016). A Bayesian model that incorporates perceptions of source

credibility is able to accommodate both camps. Exposure to new scientific evidence can ameliorate a person’s scepticism, but only if their priors allow for it. This includes factors like cultural identity and whether they cause us to discount some sources of information more than others.

6 Concluding remarks

The goal of this paper has been to explore the way in which prior beliefs affect our responsiveness to information about climate change. The Bayesian paradigm provides a natural analytical framework and I have proposed a group of stylised sceptics to embody the degrees of climate scepticism that one encounters in the real world. The primary finding is that subjective priors are generally overwhelmed by the empirical evidence for climate change. Once they have updated their beliefs in accordance with the available data, most sceptics demonstrate a clear tendency to congregate towards the noninformative case that serves as an objective reference point for this study. Depending on the preferred reconstruction of global temperatures, the 95% posterior probability range for TCR under a noninformative prior is somewhere between 1.4 °C and 1.9 °C. This implies a tighter bound on climate sensitivity than suggested by the IPCC (2013), whose “likely” range for TCR is 1.0–2.5 °C.

While these findings would seem to offer comfort against the most probability of extreme future warming, they should not be taken as evidence against the need for climate mitigation. Indeed, the updated beliefs of our various sceptics are shown to be consistent with a carbon price that is substantially greater than zero. I obtain a mean SCC range of approximately \$40–\$70 per tonne using the PAGE09 model of Hope (2011a). Only those with extreme *a priori* sceptic beliefs will find themselves in disagreement, or feel any confidence in the notion that unfettered emissions growth won’t lead to sizeable future warming. This suggests that a rational climate sceptic — even one that holds relatively strong prior beliefs — could embrace policy measures to constrain CO₂ emissions once they have seen all of the available data. Yet, those who still remain unconvinced are unlikely to converge with the mainstream consensus for many years hence. My results imply that their priors must be of such a strength that even decades more of accumulated evidence may not be enough to convince them. Fully disentangling the root causes of such information immunity — whether climate sceptics are extremely sure of

their priors, distrustful of scientists and other experts, or some combination thereof — remains an important area for future research.

References

- Anderegg, William R. L., James W. Prall, Jacob Harold, and Stephen H. Schneider (2010). "Expert credibility in climate change," *Proceedings of the National Academy of Sciences of the United States of America (PNAS)*, 27: 12107–12109.
- Bullock, John G. (2009). "Partisan bias and the Bayesian ideal in the study of public opinion," *The Journal of Politics*, 71(03): 1109–1124.
- Clark, Dav, Michael Andrew Ranney, and Jacqueline Felipe (2013). "Knowledge Helps: Mechanistic Information and Numeric Evidence as Cognitive Levers to Overcome Stasis and Build Public Consensus on Climate Change," in N. Sebanz & I. Wachsmuth M. Knauff, M. Pauen (ed.), "Proceedings of 35th Annual Meeting of the Cognitive Science Society," pp. 2070–2075.
- Cook, John and Stephan Lewandowsky (2016). "Rational irrationality: Modeling climate change belief polarization using Bayesian networks," *Topics in Cognitive Science*, 8(1): 160–179.
- Cook, John, Dana Nuccitelli, Sarah A. Green, Mark Richardson, Bärbel Winkler, Rob Painting, Robert Way, Peter Jacobs, and Andrew Skuce (2013). "Quantifying the consensus on anthropogenic global warming in the scientific literature," *Environmental Research Letters*, 8(2): 024024.
- Cook, John, Naomi Oreskes, Peter T. Doran, William R.L. Anderegg, Bart Verheggen, Ed W. Maibach, J. Stuart Carlton, Stephan Lewandowsky, Andrew G. Skuce, and Sarah A. Green (2016). "Consensus on consensus: a synthesis of consensus estimates on human-caused global warming," *Environmental Research Letters*, 11(4): 1–7.
- Corner, Adam, Lorraine Whitmarsh, and Dimitrios Xenias (2012). "Uncertainty, scepticism and attitudes towards climate change: biased assimilation and attitude polarisation," *Climatic Change*, 114(3–4): 463–478.
- Cowtan, Kevin and Robert G. Way (2014a). "Coverage bias in the HadCRUT4 temperature series and its impact on recent temperature trends," *Quarterly Journal of the Royal Meteorological Society*, 140(683): 1935–1944.
- Cowtan, Kevin and Robert G. Way (2014b). "Coverage bias in the HadCRUT4 temperature series and its impact on recent temperature trends. UPDATE: Tempera-

- ture reconstruction by domain: version 2.0 temperature series." Available: <http://www-users.york.ac.uk/~kdc3/papers/coverage2013/update.140106.pdf>.
- Deryugina, Tatyana (2013). "How do people update? The effects of local weather fluctuations on beliefs about global warming," *Climatic change*, 118(2): 397–416.
- Doran, Peter T. and Maggie K. Zimmerman (2011). "Examining the Scientific Consensus on Climate Change," *Eos, Transactions, American Geophysical Union (EOS)*, 90(3): 22–23.
- Estrada, Francisco, Pierre Perron, , and Benjamín Martínez-López (2013a). "Statistically derived contributions of diverse human influences to twentieth-century temperature changes," *Nature Geoscience*, 6: 1050–1055.
- Estrada, Francisco and Pierre Perron (2012). "Breaks, trends and the attribution of climate change: a time-series analysis." Working paper, Department of Economics, Boston University. Available: <http://people.bu.edu/perron/workingpapers.html>.
- Estrada, Francisco and Pierre Perron (2013). "Detection and attribution of climate change through econometric methods." Working paper, Department of Economics, Boston University. Available: <http://people.bu.edu/perron/workingpapers.html>.
- Estrada, Francisco, Pierre Perron, Carlos Gay-García, and Benjamín Martínez-López (2013b). "A Time-Series Analysis of the 20th Century Climate Simulations Produced for the IPCC's Fourth Assessment Report," *PLOS ONE*, 8(3): e60017.
- Fiske, Susan T. and Cydney Dupree (2014). "Gaining trust as well as respect in communicating to motivated audiences about science topics," *Proceedings of the National Academy of Sciences*, 111(4): 13593–13597.
- Gauchat, Gordon (2012). "Politicization of science in the public sphere a study of public trust in the United States, 1974 to 2010," *American Sociological Review*, 77(2): 167–187.
- Gay-Garcia, Carlos, Francisco Estrada, and Armando Sánchez (2009). "Global and hemispheric temperatures revisited," *Climatic Change*, 94(3-4): 333–349.
- Gelman, Andrew, John B. Carlin, Hal S. Stern, David B. Dunson, Aki Vehtari, and Donald B. Rubin (2014). *Bayesian Data Analysis*, Chapman & Hall / CRC, 3rd edition,
- Greene, William H. (2007). *Econometric Analysis*, Pearson Education, 6th edition.

- Hansen, J., I. Fung, A. Lacis, D. Rind, S. Lebedeff, R. Ruedy, G. Russell, and P. Stone (1988). "Global climate changes as forecast by Goddard Institute for Space Studies three-dimensional model," *Journal of Geophysical Research*, 93: 9341–9364.
- Harris, Adam J.L., Ulrike Hahn, Jens K. Madsen, and Anne S. Hsu (2015). "The appeal to expert opinion: Quantitative support for a Bayesian network approach," *Cognitive Science*, 40: 1–38.
- Hmielowski, Jay D., Lauren Feldman, Teresa A. Myers, Anthony A. Leiserowitz, and Edward W. Maibach (2014). "An attack on science? Media use, trust in scientists, and perceptions of global warming," *Public Understanding of Science*, 23(7): 866–883.
- Hope, Chris (2011a). "The PAGE09 Integrated Assessment Model: A Technical Description," Cambridge Judge Business School, Working Paper Series (4/2011). Available: http://www.jbs.cam.ac.uk/fileadmin/user_upload/research/workingpapers/wp1104.pdf.
- Hope, Chris (2011b). "The Social Cost of CO2 from the PAGE09 model," Cambridge Judge Business School, Working Paper Series (5/2011). Available: http://www.jbs.cam.ac.uk/fileadmin/user_upload/research/workingpapers/wp1105.pdf.
- Hornsey, Matthew J., Emily A. Harris, Paul G. Bain, and Kelly S. Fielding (2016). "Meta-analyses of the determinants and outcomes of belief in climate change," *Nature Climate Change*, 6(6): 622.
- IPCC (2001). *Climate change 2001: The Scientific Basis. Contribution of Working Group I to the Third Assessment Report of the Intergovernmental Panel on Climate Change*, [Houghton, J.T., Y. Ding, D.J. Griggs, M. Noguer, P.J. van der Linden, X. Dai, K. Maskell, and C.A. Johnson (eds.)], Cambridge: Cambridge University Press, 1st edition.
- IPCC (2013). *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, [Stocker, T.F., D. Qin, G.-K. Plattner, M. Tignor, S.K. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex and P.M. Midgley (eds.)], Cambridge University Press, 1st edition.
- IWG (2010). "Social Cost Of Carbon For Regulatory Impact Analysis Under Executive Order 12866." Technical Support Document, Interagency Working Group on Social Cost of Carbon, United States Government, Washington, DC.

- IWG (2013). "Technical Update of the Social Cost Of Carbon For Regulatory Impact Analysis Under Executive Order 12866." Technical Support Document, Interagency Working Group on Social Cost of Carbon, United States Government, Washington, DC.
- Jaynes, Edwin T. (2003). *Probability Theory: The Logic of Science*, Cambridge University Press.
- Kagan, Mark (2014). "Climate Change Skepticism in the Face of Catastrophe," Tinbergen Institute Discussion Paper. Available: <http://papers.tinbergen.nl/12112.pdf>.
- Kahan, Dan M., Hank Jenkins-Smith, and Donald Braman (2011). "Cultural cognition of scientific consensus," *Journal of Risk Research*, 14(2): 147–174.
- Kahan, Dan M., Ellen Peters, Maggie Wittlin, Paul Slovic, Lisa Larrimore Ouellette, Donald Braman, and Gregory Mandel (2012). "The polarizing impact of science literacy and numeracy on perceived climate change risks," *Nature Climate Change*, 2: 732–735.
- Kahneman, Daniel and Amos Tversky (1972). "Subjective probability: A judgment of representativeness," *Cognitive Psychology*, 3(3): 430 – 454.
- Kaufmann, Robert K., Heikki Kauppi, and James H. Stock (2006). "Emissions, Concentrations, & Temperature: A Time Series Analysis," *Climatic Change*, 77(3-4): 249–278.
- Kaufmann, Robert K., Michael L. Mann, Sucharita Gopal, Jackie A. Liederman, Peter D. Howe, Felix Pretis, Xiaojing Tang, and Michelle Gilmore (2017). "Spatial heterogeneity of climate change as an experiential basis for skepticism," *Proceedings of the National Academy of Sciences*, 114(1): 67–71.
- Kelly, David L. and Charles D. Kolstad (1999). "Bayesian learning, growth, and pollution," *Journal of Economic Dynamics and Control*, 23(4): 491 – 518.
- Kiseleva, Tatiana (2016). "Heterogeneous beliefs and climate catastrophes," *Environmental and resource economics*, 65(3): 599–622.
- Koop, Gary (2006). *Bayesian Econometrics*, John Wiley & Sons Ltd, 2nd edition.
- Kruschke, John K. (2014). *Doing Bayesian Data Analysis: A Tutorial with R, JAGS, and Stan*, Academic Press, 2nd edition.

- Leach, Andrew J. (2007). "The climate change learning curve," *Journal of Economic Dynamics and Control*, 31(5): 1728–1752.
- Leiserowitz, Anthony A., Edward W. Maibach, Connie Roser-Renouf, Nicholas Smith, and Erica Dawson (2013). "Climategate, public opinion, and the loss of trust," *American behavioral scientist*, 57(6): 818–837.
- Lemoine, Derek and Christian Traeger (2014). "Watch your step: Optimal policy in a tipping climate," *American Economic Journal: Economic Policy*, 6(1): 137–66.
- Lewis, Jeffrey B. and Drew A. Linzer (2005). "Estimating Regression Models in Which the Dependent Variable Is Based on Estimates," *Political Analysis*, 13(4): 345–364.
- Malka, Ariel, Jon A. Krosnick, and Gary Langer (2009). "The association of knowledge with concern about global warming: Trusted information sources shape public thinking," *Risk Analysis*, 29(5): 633–647.
- Marvel, Kate, Gavin A. Schmidt, Ron L. Miller, and Larissa S. Nazarenko (2016). "Implications for climate sensitivity from the response to individual forcings," *Nature Climate Change*, 6(4): 386–389.
- McCright, Aaron M. and Riley E. Dunlap (2011a). "Cool dudes: The denial of climate change among conservative white males in the United States," *Global Environmental Change*, 21(4): 1163–1172.
- McCright, Aaron M. and Riley E. Dunlap (2011b). "The Politicization of Climate Change and Polarization in the American Public's Views of Global Warming 2001–2010," *Sociological Quarterly*, 52(2): 155–194.
- Meinshausen, Malte, Steven J. Smith, K. Calvin, John S. Daniel, M.L.T. Kainuma, J.F. Lamarque, K. Matsumoto, S.A. Montzka, S.C.B. Raper, K. Riahi, et al. (2011). "The RCP greenhouse gas concentrations and their extensions from 1765 to 2300," *Climatic Change*, 109(1-2): 213–241.
- Mills, Terence C. (2009). "How robust is the long-run relationship between temperature and radiative forcing?" *Climatic Change*, 94(3-4): 351–361.
- Mooney, Chris (2016). "Ted Cruz keeps saying that satellites don't show global warming. Here's the problem." *The Washington Post*.

- Moore, Frances C. (2017). "Learning, adaptation, and weather in a changing climate," *Climate Change Economics*, 8(04): 1750010.
- Nordhaus, William D. (2014). "Estimates of the Social Cost of Carbon: Concepts and Results from the DICE-2013R Model and Alternative Approaches," *Journal of the Association of Environmental and Resource Economists*, 1(2): 273–312.
- Nyhan, Brendan and Jason Reifler (2010). "When Corrections Fail: The Persistence of Political Misperceptions," *Political Behavior*, 32(2): 303–330.
- Oreskes, Naomi (2004). "The Scientific Consensus on Climate Change," *Science*, 306: 1686.
- Ranney, Michael Andrew and Dav Clark (2016). "Climate Change Conceptual Change: Scientific Information Can Transform Attitudes," *Topics in Cognitive Science*, 8(1): 49–75.
- Ranney, Michael Andrew, Dav Clark, Daniel Lee Reinholz, and Sarah Cohen (2012). "Changing global warming beliefs with scientific information: Knowledge, attitudes, and RTMD (Reinforced Theistic Manifest Destiny theory)," in & R.P. Cooper N. Miyake, D. Peebles (ed.), "Proceedings of the 34th Annual Conference of the Cognitive Science Society. Austin, TX: Cognitive Science Society," pp. 2228–2233.
- Saad, Lydia (2019). "Americans as Concerned as Ever About Global Warming," *Gallup*, 25 March 2019. Available: <https://news.gallup.com/poll/248027/americans-concerned-ever-global-warming.aspx>.
- Schmidt, Gavin (2007). "The CO₂ problem in 6 easy steps," *RealClimate.org*, 6 August 2007. Available: <http://www.realclimate.org/index.php/archives/2007/08/the-co2-problem-in-6-easy-steps/>.
- Shindell, Drew T. (2014). "Inhomogeneous forcing and transient climate sensitivity," *Nature Climate Change*, 4: 274–277.
- Stern, David I. and Robert K. Kaufmann (2000). "Detecting a global warming signal in hemispheric temperature series: A structural time series analysis," *Climatic Change*, 47(4): 411–438.
- Tol, Richard S. J. (2014). "Quantifying the consensus on anthropogenic global warming in the literature: A re-analysis," *Energy Policy*, 73: 701–705.

- Tol, Richard S. J. and Aart F. De Vos (1998). "A Bayesian Statistical Analysis of the Enhanced Greenhouse Effect," *Climatic Change*, 38: 87–112.
- Van Vuuren, Detlef P., Jae Edmonds, Mikiko Kainuma, Keywan Riahi, Allison Thomson, Kathy Hibbard, George C. Hurtt, Tom Kram, Volker Krey, Jean-Francois Lamarque, et al. (2011). "The representative concentration pathways: an overview," *Climatic Change*, 109(1-2): 5–31.
- Van Wijnbergen, Sweder and Tim Willems (2015). "Optimal Learning on Climate Change: Why climate skeptics should reduce emissions," *Journal of Environmental Economics and Management*, 70: 17–33.
- Verheggen, Bart, Bart Strengers, John Cook, Rob van Dorland, Kees Vringer, Jeroen Peters, Hans Visser, and Leo Meyer (2014). "Scientists' Views about Attribution of Global Warming," *Environmental Science & Technology*, 48 (16): 8963–8971.
- Zhou, Jack (2016). "Boomerangs versus Javelins: How Polarization Constrains Communication on Climate Change," *Environmental Politics*, 25(5): 1–24.

Supplementary Material

Table SM1: TCR efficacies

Forcing agent	Mean	95% C.I.
Aerosols	1.55	(1.05, 2.05)
GHGs	1.17	(1.07, 1.28)
Land use	3.82	(-2.16, 9.80)
Ozone	0.66	(0.34, 0.98)
Solar	1.68	(-1.27, 4.63)
Volcanic	0.61	(0.33, 0.89)

Adapted from [Marvel et al. \(2016, Table S1\)](#). Confidence intervals on the sample means are constructed from a t distribution with 4 degrees of freedom.

Table SM2: Posterior TCRs when separating anthropogenic sources from other forcings (robustness check)

	Noninformative	Lukewarmer		Denier	
		Moderate	Strong	Moderate	Strong
TCR (anthropogenic)	1.5 [1.4, 1.8]	1.5 [1.3, 1.7]	1.2 [1.1, 1.4]	1.4 [1.2, 1.6]	0.3 [0.2, 0.5]
TCR (CO ₂)	1.7 [1.4, 2]	1.5 [1.2, 1.8]	1.1 [1, 1.3]	1.3 [1, 1.5]	0.2 [0.1, 0.3]

Implied TCR values are derived from an alternative implementation of regression equation 3. In the first case, anthropogenic forcings are entered separately from natural forcings. In the second, CO₂ forcings are entered separately from all other sources. See main text for details. While the underlying regression coefficients are omitted for brevity, the notes from Table 3 all apply.

Table SM3: Covariate vectors for prediction in the year 2100

	RCP 2.6	RCP 4.5	RCP 6.0	RCP 8.5
RF_{2100}	2.626	4.281	5.522	8.340
CO_2 component	85%	83%	86%	78%
Solar component	7%	4%	3%	2%
\overline{AER}	0.017	0.017	0.017	0.017
\overline{SOI}	-0.079	-0.079	-0.079	-0.079
\overline{AMO}	-0.002	-0.002	-0.002	-0.002

Notes: Covariates are used to predict the global mean surface temperature anomaly in the year 2100. The Representative Concentration Pathways (RCPs) are a family of forcing scenarios developed for the IPCC ([Van Vuuren et al., 2011](#)). Each RCP has a core component of atmospheric CO_2 concentrations, measured in parts per million volume (ppmv). With regard to the covariates in the regression model, total radiative forcing (RF) and volcanic aerosols ($VOLC$) are measured in Wm^{-2} . The Southern Oscillation Index (SOI) and Atlantic Multidecadal Oscillation (AMO) are measured as scaled indices. Future values for RF are taken from the RCP database. For the rest, historical mean values are used.

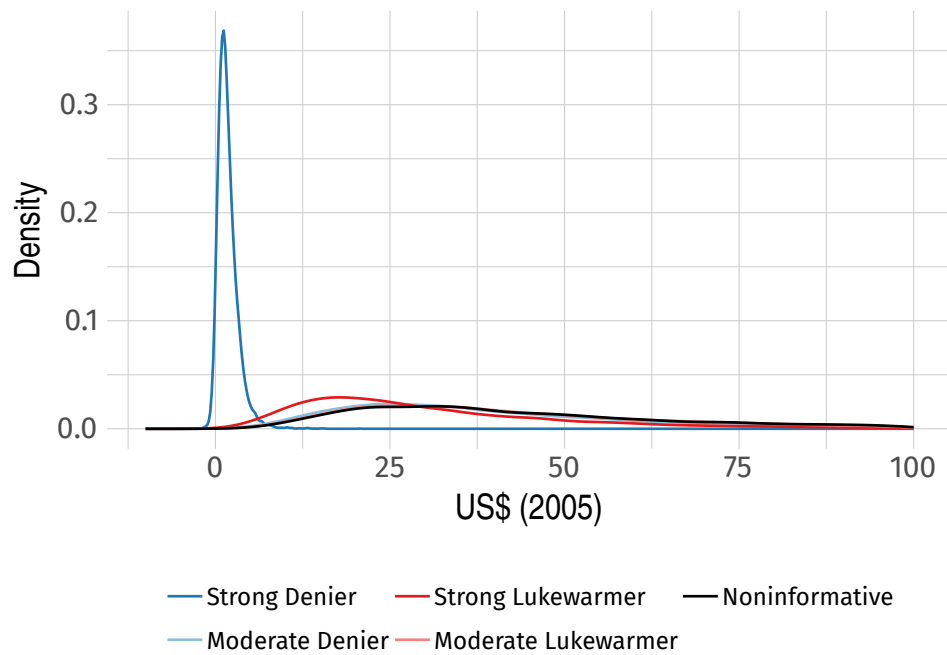


Figure SM1: Social cost of carbon (US\$2005 per tonne)

Note: The results for each agent type are obtained from 10,000 simulation runs of PAGE09 (Hope, 2011a). Posterior TCR distributions serve as key inputs to the model, while the remaining parameters are set to the PAGE09 model defaults. The x-axis is truncated at 100 to aid visual inspection; the uppermost tails of the distributions being well in excess of the range given here. See discussion in Section 4.4.