

Alma Mater Studiorum Università di Bologna
Archivio istituzionale della ricerca

Nondogmatic Climate Policy

This is the final peer-reviewed author's accepted manuscript (postprint) of the following publication:

Published Version:

Nondogmatic Climate Policy / Niko Jaakkola; Antony Millner. - In: JOURNAL OF THE ASSOCIATION OF ENVIRONMENTAL AND RESOURCE ECONOMISTS. - ISSN 2333-5955. - STAMPA. - 9:4(2022), pp. 807-841. [10.1086/719237]

Availability:

This version is available at: <https://hdl.handle.net/11585/896583> since: 2022-10-19

Published:

DOI: <http://doi.org/10.1086/719237>

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>).
When citing, please refer to the published version.

(Article begins on next page)

This is the final peer-reviewed accepted manuscript of:

Jaakkola, N., & Millner, A. (2022). Nondogmatic Climate Policy. *Journal of the Association of Environmental and Resource Economists*, 9(4), 807-841.

The final published version is available online at:

<https://doi.org/10.1086/719237>

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>)

When citing, please refer to the published version.

Non-dogmatic climate policy

Niko Jaakkola^{*1} and Antony Millner^{†2}

¹University of Bologna

²University of California, Santa Barbara & NBER

November 6, 2021

Abstract

Disagreements about normative aspects of social time preferences have led to estimates of the Social Cost of Carbon (SCC) that differ by orders of magnitude. We investigate how disagreements about the SCC change if planners are *non-dogmatic*, i.e., they admit the possibility of a change in their normative views, and internalise the preferences of future selves. Although non-dogmatic planners may disagree about all the contentious aspects of social time preferences, disagreements about the SCC reduce dramatically. Admitting the possibility of a change in views once every 40 years results in a five-fold reduction in the range of recommended SCCs.

Keywords: Social cost of carbon, non-dogmatism, social discounting

^{*}Jaakkola: niko.jaakkola@unibo.it.

[†]Millner: amillner@econ.ucsb.edu. We are grateful to Lint Barrage, Kieran Walsh, seminar participants at UC Berkeley and Uppsala University, and three anonymous referees, for helpful feedback.

1 Introduction

Disagreements about social time preferences are a major source of contention in climate change economics. The Social Cost of Carbon (SCC) – the welfare cost of a ton of carbon dioxide (CO₂) emissions – is perhaps the most important indicator of the optimal intensity of climate policy, but is highly sensitive to parameters of social time preferences that capture e.g. social impatience and aversion to intertemporal consumption inequalities. While one might hope that new data or improved estimation methods could resolve disagreements about their values, estimates of these parameters from market observables are of questionable use due to a variety of factors, including market failures, and the fact that most of those who will be affected by climate policies are not yet born (see Millner & Heal, 2021, for further discussion). Indeed, disagreements about social time preferences have an irreducibly normative character (Dasgupta, 2008; Drupp et al., 2018).

This paper examines how disagreements about the SCC might change if advocates of diverse normative theories of intertemporal social welfare were *non-dogmatic* (Millner, 2020). Non-dogmatic planners favour idiosyncratic theories of intertemporal social welfare, but exhibit some humility; they admit the possibility of a change in their normative views, and internalise the preferences of their future selves. We show that although such planners may disagree on all the contested aspects of social time preferences, disagreements about the SCC would decrease dramatically if they were even mildly non-dogmatic. Even if each planner admitted the possibility of a change in their views only once every 40 years on average, the range of recommended SCC values shrinks by almost a factor of five. Thus,

even a small amount of humility about normative judgements can help to reduce seemingly intractable disagreements about the appropriate level of climate policy.

The theory of non-dogmatic social time preferences was developed in Millner (2020). That paper motivated non-dogmatism as a normative principle in its own right, arguing that planners should exhibit a degree of humility when forming their normative judgments. It then showed that although non-dogmatic planners may disagree about every free parameter of their intertemporal welfare functions, they all agree on the long-run social discount rate, i.e., the rate of decline of the social value of marginal payoffs at long maturities. Since current CO₂ emissions cause persistent warming that lasts for significantly more than one hundred years (Ricke & Caldeira, 2014), intuition might suggest that non-dogmatism will reduce disagreement about climate policy variables like the SCC. However, it is not at all obvious how much reduction in disagreement can be expected, and it turns out that even the direction of the change is not certain.¹ While Millner (2020) used a calibrated version of his model to demonstrate a substantial reduction in disagreements about long-run social discount rates, his analysis focussed on evaluating exogenous, marginal, policies. By contrast, this paper focusses on endogenous, optimal, policies. Since SCCs are calculated along optimal consumption paths they are complex functions of the distribution of non-dogmatic planners' preferences, and of the variation in the stochastic trajectory of the

¹The intuition that non-dogmatism will reduce disagreements on the SCC often holds, but is not correct in general. This is a consequence of the fact that there is a many-to-one mapping between conventional (i.e., dogmatic) social time preferences and a given value of the SCC. Consider two dogmatic planners with different preferences who happen to agree on the SCC. If we then require these planners to be a little non-dogmatic, this will introduce novel variation across their social discount rates at shorter maturities. This will drive their *a priori* identical SCC values apart – non-dogmatism *increases* disagreement in this case. Assessing the effect of non-dogmatism on the SCC is thus a fundamentally empirical question that requires a detailed calibrated model.

economy associated with that distribution. A new modelling approach is thus required in order to assess the consequences of non-dogmatism for endogenous policy choices, and disagreements about the SCC.

Operationalising the formalism developed in Millner (2020) in a fully optimising climate-economy model raises several challenges. First, the formalism developed in Millner (2020) must be generalised to a model of preferences over state-contingent consumption plans; this is a non-trivial extension, as we discuss below. Second, since non-dogmatic planners' preferences may change over time one might expect them to be time inconsistent, potentially complicating the computation of equilibrium policies. We show however that the Millner (2020) model admits preference changes *with* time consistency, and thus focus our attention on the time consistent version of the model. Third, since climate-economy models usually have a number of endogenous state variables, and future preferences are uncertain in the non-dogmatic paradigm, curse of dimensionality problems familiar from numerical stochastic-dynamic optimisation raise their head. Fourth, since we need to solve the model repeatedly from the perspective of many different non-dogmatic planners, the computational demands multiply. We show that these last two computational difficulties can be overcome with judicious simplifications of the problem.

Related literature

Disagreements about social time preferences are at the heart of the extensive debate about optimal climate policy that followed the publication of the Stern Review on the

Economics of Climate Change (Stern, 2008; Nordhaus, 2007; Dasgupta, 2008; Weitzman, 2007). The extent of these disagreements has been documented in a recent survey of economists (Drupp et al., 2018). It is well known that the SCC is highly sensitive to these disagreements (see e.g. Nordhaus, 2008; Heal & Millner, 2014; Anderson et al., 2014, and Figure 2 below).

Alternative approaches to resolving disagreements about social time preferences in the context of climate policy include Dietz & Matei (2016); Heal & Millner (2014). The former paper focusses on an incomplete dominance relation (which does not admit calculations of the SCC) in a non-optimising model, while the latter studies utilitarian aggregation of time preferences, also in a non-optimising model. By contrast, the formalism we develop here delivers a set of complete social preference relations, and we use this model to investigate the effects of non-dogmatism on the distribution of SCCs in a fully optimising climate-economy model.

The notion of ‘non-dogmatic’ social time preferences builds on and reinterprets a prior literature in which altruistic agents internalise others’ preferences (see e.g. Ray (1987); Saez-Marti & Weibull (2005); Galperti & Strulovici (2017), and the discussion in Millner (2020)). Alternative approaches to dealing with disagreements about intertemporal social preferences are explored in Gollier & Zeckhauser (2005); Feng & Ke (2018); Chambers & Echenique (2018); Millner & Heal (2018).² While these papers examine a variety of aggrega-

²At a high level of abstraction, the analysis of Weitzman (1998, 2001) can be seen as a precursor to some of this work. It has some technical features in common with papers that study utilitarian aggregation rules (Weitzman considers linear aggregation across exogenous exponentially declining discount factors), but lacks a formal welfare analytic foundation, and delivers different results. See Freeman & Groom (2015); Millner & Heal (2021) for further discussion.

tion techniques, non-dogmatism makes no attempt to aggregate preferences; disagreements are left intact, but each planner is required to exhibit some humility when forming his/her normative judgements. We show that this is enough to dramatically reduce disagreements about the SCC.

2 Non-dogmatic social time preferences

We begin with a description of non-dogmatic social time preferences and their properties. Our presentation builds on the analysis in Millner (2020), but extends it to the context of the present paper. While Millner (2020) considers preferences over exogenous deterministic consumption paths, we consider preferences over state-contingent consumption plans, which may be endogenous. This additional complexity is necessary, as we aim to study optimal paths in a stochastic-dynamic economy in this paper. We show that under mild conditions the main result in Millner (2020), i.e., that all non-dogmatic planners agree on long-run consumption discount rates, can be extended to this setting. While it has a similar flavour, the result we prove in this section is a non-trivial extension of that in Millner (2020). Incorporating state-dependent consumption requires us to use different mathematical methods, and leads to a different formula for the consensus value of long-run discount payoffs.³

The primitives of the model are a set of N normatively plausible social time preferences, indexed by $i = 1, \dots, N$. The key novelty in our approach is that at each time τ , a planner

³See footnote 8 for a discussion of the differences between the results in these two papers.

who currently adheres to preference i admits the possibility that she may adopt preference j at time $\tau + 1$. Moreover, current planners anticipate this possibility, and incorporate the preferences of future selves into their current preferences in a way we will make precise below. Our model thus allows for the possibility of *normative humility* – at any given time planners still advocate a single normative paradigm unequivocally, but they recognise the possibility of a future change of heart, and their current preferences are constructed with that possibility in mind. This approach allows planners to retain sovereignty over their ethical judgements, while recognising the potential legitimacy of alternative normative paradigms.

Let $w_{ij} \geq 0$ denote the probability that a planner switches from social preference i to j in the next time period, where we assume that $\sum_{j=1}^N w_{ij} = 1$ for all i . The $N \times N$ matrix of switching probabilities $\mathbf{w} = (w_{ij})$ defines a time homogeneous Markov chain on the set of preferences. When there is no possibility for confusion we use \mathbf{w} to refer to the transition matrix, and to the Markov chain itself. The physical state of the economy (e.g., capital, CO₂) at time τ is given by a vector of state variables S_τ , and this state is assumed to evolve according to

$$S_{\tau+1} = F(S_\tau, c_\tau) \tag{1}$$

where F is an arbitrary continuous function, and c_τ is the value of a choice variable at time τ . For expositional simplicity we will think of c_τ as a one dimensional consumption measure; our presentation is easily extended to the case where c_τ is a vector of choice variables, at the expense of more cumbersome notation. As the model in (1) and the

preference switching process \mathbf{w} are both Markovian, and the preferences we will study are history independent, the decision-relevant history of events at time τ is captured by the ordered pair (S_τ, i_τ) , where $i_\tau \in \{1, \dots, N\}$ tells us which of the N preferences the planner holds at τ . A state-contingent plan $\mathbf{c} = c(S, i)$ is a function with domain equal to the set of all possible states $\{(S, i)\}$ that takes values in some choice set, which may depend on the physical state S .

Denote the preferences of a planner who adheres to theory i in physical state S by $\succeq_{(S,i)}$; these preferences rank all state-contingent plans \mathbf{c} . We study the following representation of planner preferences:

$$\mathbf{c} \succeq_{(S,i)} \mathbf{c}' \iff V^{(S,i)}(\mathbf{c}) \geq V^{(S,i)}(\mathbf{c}'),$$

where

$$V^{(S,i)}(\mathbf{c}) = U^i(c(S, i), S) + \beta^i \sum_{j=1}^N w_{ij} V^{(F(S, c(S, i)), j)}(\mathbf{c}) \quad (2)$$

and $\beta^i \in (0, 1)$. When $w_{ij} > 0$ for all i, j , planners with preferences (2) will be said to be *non-dogmatic*.⁴ We allow the utility functions $U^i(c, S)$ to depend directly on the state S , which could include e.g. population size. Given a fixed value of the physical states S_τ at the current time τ , the set $\{V^{(S_\tau, i)}(\mathbf{c})\}_{i=1, \dots, N}$ is the plurality of values that our hypothesised non-dogmatic planners assign to a plan \mathbf{c} . Note that we use the word ‘preferences’ to refer to the ranking of plans induced by $V^{(S_\tau, i)}$ for a *particular* value of i . The set of such

⁴This condition is stronger than is necessary for the result in Theorem 1 to hold, but is normatively intuitive. See Millner (2020) for further discussion of this point.

rankings is referred to as the ‘set of non-dogmatic preferences’.

The interpretation of the preference representation in (2) is as follows: a planner who advocates normative theory i in physical state S favours her own idiosyncratic theory of intertemporal social welfare $V^{(S,i)}$, associated with the utility function $U^i(\cdot)$, the discount factor β^i , and the weights w_{ij} . However, this planner admits the possibility that her normative views may change in the next period; she believes that her future self at time $\tau + 1$ will advocate theory j with probability w_{ij} . The planner is non-dogmatic – she *internalises* the possible preferences of the self at time $\tau + 1$ into her current preferences, i.e., current welfare depends directly on the welfare measure that a future self may advocate. Finally, non-dogmatism is persistent: planners are always non-dogmatic – they never rule out the possibility of a future switch to another plausible normative theory of social welfare. Preference internalisation and persistence together yield a recursive preference system in which current preferences depend on future preferences, each of which is in turn recursively defined.

Let (x, \mathbf{c}) denote a plan that chooses x in the current state, followed by the continuation of plan \mathbf{c} in future states. We say that a set of non-dogmatic preferences is *time consistent* if for all states (S, i) , all current choices x , and all plans \mathbf{c}, \mathbf{c}' ,

$$(x, \mathbf{c}) \succeq_{(S,i)} (x, \mathbf{c}') \iff \forall j \in \{1, \dots, N\} \mathbf{c} \succeq_{(F(S,x),j)} \mathbf{c}'. \quad (3)$$

It is straightforward to verify by inspection that, although the model in (2) allows for the possibility that utility functions and discount factors may change over time, preferences

over state-contingent plans are nevertheless time consistent. This follows since current values $V^{(S_\tau, i_\tau)}$ are increasing functions of future values $V^{(S_{\tau+1}, i_{\tau+1})}$. Millner (2020) considers a more general model of non-dogmatic preferences, where values at τ may depend not just on future values at $\tau + 1$, but also on values adopted at $\tau + 2, \tau + 3, \dots$. In Appendix A we show that the model in (2) is the only such model that is time consistent. Time consistency is an attractive normative property of preferences, and also significantly simplifies the analysis of equilibrium policies. In addition, the model in (2) has the desirable feature that it reduces to a set of standard exponential discounted utilitarian time preferences when $w_{ij} = 0$ for $j \neq i$. Thus (2) is a natural generalisation of the social preferences that are commonly used in models of optimal climate policy, including the DICE model that will be the focus of our numerical work. For these reasons, we focus on the time consistent model in this paper.

Given a consumption plan $\mathbf{c} = c(S, i)$, the Markov chain \mathbf{w} on the set of preferences induces a Markov process on the full state space. The transition function for this process is

$$T(S', j|S, i) = w_{ij}\delta(S' - F(S, c(S, i))) \quad (4)$$

where $\delta(\cdot)$ is the Dirac delta function,⁵ and $\sum_j \int T(S', j|S, i)dS' = 1$ for all i . We say that a pair (\mathbf{c}, \mathbf{w}) is *ergodic* if the Markov process it induces in (4) has a unique globally

⁵In one dimension the Dirac delta function is defined by $\delta(x) = 0$ for $x \neq 0$, and $\int_{-\infty}^{\infty} \delta(x)dx = 1$, with the obvious generalisation to higher dimensions. We use the Dirac delta function, rather than simply defining the transition function piecewise, since in the $t \rightarrow \infty$ limit we will need to integrate over S .

asymptotically stable stationary distribution, which we denote by $p_\infty(S, i)$. When this is the case the long-run distribution of states that the Markov process (4) visits is independent of the initial state. We will discuss when this occurs in equilibrium below. When it exists, the expected value of a function $h(S, i)$ with respect to the stationary distribution on state space will be denoted by $E_\infty h(S, i)$.⁶

Given an initial state (S_τ, i_τ) , define a *trajectory* of the dynamical system in (4) as any sequence $\{(S_{\tau+t}, i_{\tau+t})\}_{t=0, \dots, \infty}$, where for each time $\tau + t$, $(S_{\tau+t+1}, i_{\tau+t+1})$ is in the support of the transition function $T(S', i | S_{\tau+t}, i_{\tau+t})$ in (4). We will denote a generic trajectory by χ . It is convenient for what follows to write preferences (2) in terms of expectations over trajectories. The value of a trajectory $\chi = \{(S_{\tau+t}, i_{\tau+t})\}_{t=0, \dots, \infty}$ is defined recursively as

$$V_\chi^{(S_\tau, i_\tau)}(\mathbf{c}) = U^{i_\tau}(c(S_\tau, i_\tau), S_\tau) + \beta^{i_\tau} V_\chi^{(S_{\tau+1}, i_{\tau+1})}(\mathbf{c}). \quad (5)$$

This is the value planner i_τ would obtain if she knew for sure that the trajectory of the economy, and future preferences, would be χ . Clearly we have

$$V^{(S_\tau, i_\tau)}(\mathbf{c}) = E_{\chi | (S_\tau, i_\tau)} V_\chi^{(S_\tau, i_\tau)}(\mathbf{c})$$

where the expectation $E_{\chi | (S_\tau, i_\tau)}$ is taken with respect to the measure on the space of tra-

⁶Note that the variables in (4) may be measured in efficiency units. For notational simplicity we don't draw an explicit distinction between the consumption variable that enters the state equation (which may be in efficiency units) and the consumption variable that enters the utility function (i.e., per capita consumption). Our main results below apply to all cases by allowing for an arbitrary growth process for per capita consumption in the stationary distribution.

jectories induced by the transition probability law (4), given the initial condition (S_τ, i_τ) .⁷

Denote the consumption variable at maturity t along trajectory χ by c_t^χ . We define the consumption discount rate at maturity t along χ as

$$r(t|\chi) = \left(\frac{\frac{\partial V_\chi^{(S_\tau, i_\tau)}}{\partial c_t^\chi} \Big|_\chi}{\frac{\partial V_\chi^{(S_\tau, i_\tau)}}{\partial c_0^\chi} \Big|_\chi} \right)^{-\frac{1}{t}} - 1. \quad (6)$$

Consider a marginal project π with consumption payoffs $\pi(S', j)$ in state (S', j) , and let π_t^χ be this project's payoffs at maturity t along a trajectory χ . The project is welfare improving according to the planner at time τ if and only if

$$\mathbb{E}_{\chi|(S_\tau, i_\tau)} \sum_{t=0}^{\infty} (1 + r(t|\chi))^{-t} \pi_t^\chi > 0. \quad (7)$$

We denote the expected present value of this project at maturity t , according to a planner who currently advocates theory i in state S , by

$$\Delta_\pi(t|S, i) = \mathbb{E}_{\chi|(S, i)} (1 + r(t|\chi))^{-t} \pi_t^\chi$$

where the expectation $\mathbb{E}_{\chi|(S, i)}$ is over all trajectories that emanate from the initial state (S, i) .

We are interested in the dependence of $\Delta_\pi(t|S, i)$ on the initial state (S, i) at large maturities, i.e., as $t \rightarrow \infty$. To simplify the exposition we make two additional assumptions.

⁷See Section 8.2 of Stokey et al. (1989) for a detailed discussion of how this measure is constructed.

First, we specialise to the case where social utility functions are classical utilitarian and iso-elastic in consumption – we maintain these assumptions in our numerical work below. Denoting population size by P , and assuming that utility $U^i(c, S)$ depends on the vector of physical states S only through P , this assumption amounts to:

$$U^i(c, P) = P \times \begin{cases} \frac{c^{1-\eta^i}}{1-\eta^i} & \eta^i \geq 0, \eta^i \neq 1 \\ \ln c & \eta^i = 1. \end{cases} \quad (8)$$

Second, we assume that project payoffs π depend on the physical states S , but *not* on the preferences j of the planner, i.e., $\pi = \pi(S)$. For all practical purposes this assumption is without loss of generality, since cost-benefit analysis assumes that projects are marginal, and hence exogenous to planners' actions. For the sake of completeness we examine the case where project payoffs depend on j in Appendix B.

Denote population and 1-period pure time discount factors at maturity t along a trajectory χ by P_t^χ , and β_t^χ respectively. We also define

$$\begin{aligned} \lambda_t^\chi &= \frac{P_t^\chi}{P_{t-1}^\chi} - 1, \\ g_t^\chi &= \frac{c_t^\chi}{c_{t-1}^\chi} - 1, \\ \rho_t^\chi &= (\beta_{t-1}^\chi)^{-1} - 1 \end{aligned}$$

for $t \geq 1$. λ_t^χ is the 1-period growth rate of population, g_t^χ the 1-period growth rate of consumption, and ρ_t^χ the 1-period pure time discount rate at maturity t along trajectory

χ . Generic values of these quantities in a sequential pair of states $\{(S, i), (F(S, c(S, i)), j)\}$ will be denoted by $\tilde{\lambda}$, \tilde{g} , and $\tilde{\rho}$ respectively.

The main result of this section is as follows:

Theorem 1. *Suppose that preferences are given by (2) and (8) and that the Markov process (4) induced by the pair (\mathbf{c}, \mathbf{w}) is ergodic. In addition, assume that there exists $\epsilon > 0$ such that $c_t^x > \epsilon$ for all trajectories χ that emanate from an initial state (S, i) . For a fixed value of η^j , define*

$$\hat{r}(j) = \exp \left(\mathbb{E}_\infty \log \left[\frac{(1 + \tilde{\rho})(1 + \tilde{g})^{\eta^j}}{1 + \tilde{\lambda}} \right] \right) - 1, \quad (9)$$

and let

$$\hat{r} = \min_j \hat{r}(j). \quad (10)$$

Then for any marginal project $\boldsymbol{\pi}$ with bounded payoffs $\pi(S)$ in state S ,

$$\lim_{t \rightarrow \infty} \frac{\Delta_{\boldsymbol{\pi}}(t|S, i)}{(1 + \hat{r})^{-t} \mathbb{E}_\infty \pi(S')} = 1. \quad (11)$$

Proof. See Appendix B. □

Let's unpack this result. The limiting formula in (11) says that, regardless of the initial state S and the planner's preference i , the long-run present value of a marginal project is approximately

$$(1 + \hat{r})^{-t} \mathbb{E}_\infty \pi(S'),$$

where this approximation becomes exact as $t \rightarrow \infty$. The expectation E_∞ is taken with respect to the stationary distribution on the physical state space, i.e. $\sum_{j'=1}^N p_\infty(S', j')$, which crucially does not depend on the initial state (S, i) . Thus, cost-benefit analysis of long-run payoffs does *not* depend on which of the N social preferences a planner advocates in the current period. Whatever their normative proclivities, all non-dogmatic planners agree on the value of long-run payoffs.

To understand the expression for the ‘consensus’ long-run discount rate \hat{r} implied by (9–10), notice that when the random variables $\tilde{g}, \tilde{\lambda}, \tilde{\rho}$ are all small in magnitude this quantity is well approximated by

$$\hat{r} \approx E_\infty \tilde{\rho} - E_\infty \tilde{\lambda} + \min_j \{ \eta^j E_\infty \tilde{g} \}.$$

This is a standard Ramsey formula. The first term is the long-run average pure rate of social time preference. This term simply reflects the average of 1-period pure time discount factors along an arbitrary trajectory of the economy. Since the economy is ergodic (by assumption), time averages across almost all trajectories converge to expectations over state space computed with the stationary distribution. The second term is very similar – it reflects long-run average population growth, which is again computed by taking a stationary expectation over population growth rates. Finally, the third term represents a standard consumption smoothing effect, which in this case depends on the long-run average consumption growth rate and an *extreme* value of η^j . Only the lowest (highest) value of the elasticity of marginal utility is relevant in the $t \rightarrow \infty$ limit, since marginal utility declines

(increases) slowest (fastest) for this value when $E_\infty \tilde{g}$ is positive (negative).⁸

The proof of Theorem 1 proceeds by using the properties of ergodic dynamical systems to show that, regardless of which initial condition (S_0, i) prevails, the distribution of long-run discount rates that arises from the stochastic trajectories of the economy is the same. It is then a relatively straightforward matter to show that this implies that a single discount rate dominates the computation of long-run present values. It will be important for what follows to keep this distinction between long-run discount *rates*, and long-run *present values* clear. Under the conditions of Theorem 1 long-run discount rates (i.e., $r(t|\chi)$ for large t) have a common distribution across planners, but do not converge to a single value. By contrast, long-run present values depend on an expectation over discount *factors*, and may be computed using the single value \hat{r} in (10). In essence, this occurs because long-run present values can be shown to depend on an expected discount factor of the form $E_j(1+\hat{r}(j))^{-t} \sim_{t \rightarrow \infty} (1+\min_j \hat{r}(j))^{-t}$, where the expectation E_j is over the same distribution for the discount rates $\hat{r}(j)$ for all non-dogmatic planners. Appendix B provides formal details of these points.

⁸The formula in (10) is subtly different from that in Millner (2020). In Millner (2020) the consensus long-run pure rate of time preference is related to the largest eigenvalue of a matrix of intertemporal weights, whereas here we have an expectation over a stationary distribution. The reason for this difference is that consumption is not state contingent in Millner (2020), and hence the physical state space is degenerate; all trajectories of the Markov chain \mathbf{w} map to the same future consumption variables. In that setting one must take expectations over *all* trajectories of length t when computing the discount rate associated with consumption at maturity t . This expectation is computed by exponentiating a matrix of intertemporal weights, an operation that is dominated by the largest eigenvalue of the underlying matrix at large maturities. By contrast, in the current model future consumption values along different trajectories are generically unique, since they correspond to different paths on the physical state space. We thus don't need to take expectations over many trajectories that end up in the same place when computing state-contingent discount rates in the current model. Assuming ergodicity, almost all trajectories in our model have the same statistics at large maturities, given by the stationary distribution. Hence the presence of the stationary expectation in (10).

A condition of Theorem 1 is that the Markov process on state space induced by the pair (\mathbf{c}, \mathbf{w}) in (4) is ergodic. When does this occur? The most relevant case to consider for our normative analysis is when the state-contingent plan \mathbf{c} is optimal, given preferences (2) and the physical state equations (1). In this case the conditions for ergodicity have been studied in a classic contribution by Hopenhayn & Prescott (1992), with more recent extensions to non-compact state spaces by Kamihigashi & Stachurski (2014).⁹ When the elements of the transition matrix \mathbf{w} are strictly positive these conditions are quite mild, and are satisfied in many common economic environments, including models of economic growth of the kind we study below. A detailed presentation of these conditions is beyond the scope of the present paper – we refer the reader to the literature for details. Rather than studying ergodicity analytically in the abstract, we’ll examine it numerically in our applied model below.

Equation (7) indicates that the fact that non-dogmatic planners may agree on the present value of long-run projects could have especially significant consequences for disagreements about the merits of policies with long-run effects, climate policy being an archetypal example. We turn to this application of the theory next.

3 Application in a DICE-like model

In order to illustrate the implications of non-dogmatic time preferences for disagreements about the SCC, we investigate their implications in a version of the DICE integrated

⁹See also Stokey et al. (1989); Ljungqvist & Sargent (2004) for textbook discussions of ergodicity and its many applications in economics.

assessment model of climate policy (Nordhaus, 2017). DICE combines a Ramsey-style growth model with a dynamical climate module. The climate module quantifies the CO₂ emissions intensity of economic output, how emissions accumulate in the atmosphere, how the stock of CO₂ changes global average temperatures, and finally how those temperature changes feed back into the economy via climate change damages. DICE is arguably the simplest and most widely used model in climate economics, and has been used in a variety of high-profile policy applications.

The standard version of the DICE model is deterministic. It has 6 endogenous state variables (capital, two temperature variables, and three carbon cycle variables), 2 control variables (savings and the emissions control rate), and relies on several exogenous time series (for TFP, population, non-CO₂ forcings, and mitigation costs). DICE uses a discounted utilitarian social welfare function with an iso-elastic utility function to measure intertemporal social welfare:

$$V_{\tau}^{DICE} = P_{\tau} \frac{c_{\tau}^{1-\eta}}{1-\eta} + \beta V_{\tau+1}^{DICE} = \frac{1}{1-\eta} \sum_{s=0}^{\infty} P_{\tau+s} \beta^s c_{\tau+s}^{1-\eta}, \quad (12)$$

where P_{τ} is population, c_{τ} is per capita consumption, $\eta \geq 0$ is the elasticity of marginal utility, and $\beta = (1 + \rho)^{-1}$, where $\rho > 0$ is the pure rate of social time preference. The parameters ρ and η capture social impatience and aversion to intertemporal consumption inequalities respectively. Our object is to provide a counterfactual analysis of the SCC that preserves much of the standard structure of the DICE model, but allows us to investigate how non-dogmatism might reduce disagreement amongst planners who favour different

values of ρ and η . While there are other normative choices in the model that could be interrogated (e.g. its treatment of population ethics and uncertainty (Fleurbaey et al., 2018)), much of the debate in the literature has focussed on these two parameters. As our goal is to demonstrate the consequences of non-dogmatism in a well-known benchmark model, we make minimal changes to the DICE framework itself.

In order to bring the non-dogmatic preferences (2) into contact with the DICE model in a computationally tractable manner, we make simplifying assumptions on the set of planner preferences we consider, and on how to model the climate system. We describe these simplifications in turn. Full details of the model we use are available in Appendix C

3.1 Calibrating preferences

To implement (2) in a quantitative application we need to calibrate preferences to a sample of plausible normative theories of intertemporal social welfare. We follow Millner (2020), who calibrates the model to data from a survey of economists who have published papers on social discounting (Drupp et al., 2018). The survey elicited 173 respondents’ views on the appropriate values of the parameters ρ and η of the discounted utilitarian intertemporal welfare function in (12). To demonstrate how non-dogmatism reduces disagreements about the SCC, we need to solve non-dogmatic versions of the DICE model for each of these parameter values. That is, we run many versions of the DICE model assuming that a different planner is in complete control of policy in each run; however, each planner is non-dogmatic.

In the version of the model used by Millner (2020) non-dogmatic planners are assumed to have symmetric switching probabilities w_{ij} that take the following form:

$$w_{ij} = \begin{cases} 1 - x & j = i \\ \frac{x}{N-1} & j \neq i \end{cases} \quad (13)$$

Thus, planners stick to their current theory in the next period with probability $1 - x \in [1/N, 1]$. Conditional on switching (which occurs with probability $x \in [0, 1 - 1/N]$), they switch to any of the plausible alternative normative theories with equal probability. The variable x thus measures the ‘degree’ of non-dogmatism. Note that this model does not imply that switching probabilities are uniformly distributed over the space of possible preferences. Regions of (ρ, η) space that are more common in our sample will be switched to more often, on average, than sparsely populated regions. In the dogmatic case where $x = 0$, the preferences in (2) coincide with the standard discounted utilitarian preferences in (12).

A direct implementation of this calibration methodology leads to 173 different intertemporal optimization problems, and 173 different value functions on the space of physical state variables that must be solved for in each time step of the model.¹⁰ Since this problem is prohibitively costly to implement computationally, we simplify the dataset. We use a k -means clustering algorithm to assign each data point in the sample of (ρ, η) pairs to

¹⁰Since we work with a finite horizon model, value functions depend explicitly on time. One could alternatively think of this problem as solving for a single (time dependent) value function that depends on the state (S, i) , rather than N value functions that depend on S , but the numerical difficulties are of course the same in this framing.

one of $M \ll 173$ clusters, in such a way that the within cluster variance is minimised. All points within a cluster are identified with the centroid of that cluster, and a weight $m_i = \frac{\# \text{ points in cluster } i}{N}$ is assigned to cluster i . The switching probabilities between clusters, denoted w_{ij}^C , are then assumed to be proportional to their weight, so that the uniform switching probabilities at the individual level in (13) are reflected at the level of clusters:

$$w_{ij}^C = \begin{cases} 1 - x + xm_i & j = i \\ xm_j & j \neq i \end{cases}. \quad (14)$$

Figure 1 illustrates the dataset, and its partition into 10 clusters. Our main analysis works with these preference clusters, and uses the formula in (14) to define the switching probabilities w_{ij} . In Appendix E we present results for a different specification of the switching probabilities in which planners are more likely to switch to an alternative theory that is ‘close’ to their current preferred theory; the results are qualitatively similar in this case.

3.2 Simplifying the climate dynamics

Our second simplification relates to the way DICE represents climate dynamics, i.e., the relationship between CO₂ emissions and temperature change. As we observed above, DICE uses 5 state variables to represent the climate system. The climate model in DICE builds on an early simple climate model (Schneider & Thompson, 1981), which aimed to represent the lags between emissions and warming that were thought to occur because of the dynamics of heat transfer between the atmosphere and the oceans (Nordhaus & Boyer, 2000). More




Figure 1: Recommended values for the pair of parameters (ρ, η) from the Drupp et al. (2018) survey. Points with the same colour are assigned to the same cluster by a k -means clustering algorithm, and cluster centroids are marked with a cross. Ten clusters are used in this example. Note that cluster sizes are not fully represented in this figure as there may be several data points at the same value of (ρ, η) .

recent work with sophisticated earth system models has however shown that the inertia in the climate system as represented in DICE is overstated, largely because it neglects feedbacks in the climate system, and the saturation of carbon sinks in particular. To a very good approximation over time scales of a century or more, emitting a ton of CO₂ causes an almost immediate,¹¹ permanent, increase in temperature that is independent of initial CO₂ concentrations (Matthews et al., 2009; Ricke & Caldeira, 2014). With the 10 year time step that DICE uses, the relationship between temperature T_τ and emissions e_τ in successive time steps is very well approximated by:

$$T_{\tau+1} - T_\tau = \mu e_\tau$$

for some $\mu > 0$. Summing this relationship, we find

$$T_\tau - T_0 = \mu \sum_{t=0}^{\tau} e_t.$$

The Intergovernmental Panel on Climate Change’s central estimate is that 1 trillion tons of CO₂ emissions will cause approximately 2°C of warming (Collins et al., 2013) – this allows us to calibrate the value of μ . Dietz & Venmans (2019) provide a valuable discussion of the scientific literature that demonstrates the linear relationship between temperature change and cumulative emissions, and of the deficiencies of the climate model in DICE on this

¹¹Peak warming after an emissions pulse occurs after about 10 years in modern earth system models. The falloff in temperature response after the peak is very small over a period of 100 years (Ricke & Caldeira, 2014). Dietz & Venmans (2019) consider a simple two state model that allows one to fit the short-run peak in the temperature response curve, but show that such a model is virtually indistinguishable from a one state model in which warming is immediate.

dimension.

Using this simple ‘cumulative carbon’ model of the climate system allows us to reduce the number of climatic state variables in DICE from five down to one – cumulative CO₂ emissions. Although this is undeniably a simplification of reality, this model is more representative of the state of the art in climate science than the more complex climate model that DICE employs.

3.3 Numerics

Other than the changes to the climate model and social time preferences described above, the version of the DICE model we work with is virtually identical to that in Nordhaus (2017).¹² Given our simplifications, the model depends on the pair of physical state variables $S_\tau = (K_\tau, \text{CO}_{2,\tau})$, where K_τ is the capital stock, and $\text{CO}_{2,\tau}$ is cumulative carbon dioxide emissions. Letting Z_τ denote the pair of controls (i.e., consumption and the emissions control rate) at time τ , our modified DICE model can now be written as a set of equations of motion of the form:

$$S_{\tau+1} = H(S_\tau, Z_\tau, \tau).$$

These equations depend explicitly on time τ through the exogenous time series in the DICE model, although all these time series tend to constant steady state values at large times.

¹²There are two other minor differences. We do not include exogenous non-CO₂ forcings – these have only small effects on the SCC. In addition, unlike the version of DICE in Nordhaus (2017), we constrain the emissions control rate to lie in $[0,1]$.

The model can now be solved using methods from numerical stochastic dynamic programming (see Appendix D for further details). Formally, we find numerical approximations to the M time dependent functions of S , $\{V_\tau(S, i)\}_{i=1, \dots, M}$, defined through

$$V_\tau(S, i) = \max_Z P_\tau U^i(Z) + \frac{1}{1 + \rho_i} \sum_{j=1}^M w_{ij}^C V_{\tau+1}(H(S, Z, \tau), j).$$

The SCC is the welfare cost of an additional ton of CO₂ emissions in consumption units.

According to planner i at time τ in state S_τ , the SCC is given by¹³

$$SCC_\tau(S_\tau, i) = \left| \frac{\partial V_\tau(S, i) / \partial \text{CO}_2|_{S_\tau}}{\partial V_\tau(S, i) / \partial K|_{S_\tau}} \right|. \quad (15)$$

Each run of the model with a different planner ‘in charge’ starts out with the same initial values of the state variables S_0 . In the initial period $\tau = 0$ the SCC is a single number for each planner. However, the trajectories of state variables from $\tau = 1$ onwards are endogenous to planners’ preferences, and stochastic (due to the possibility of a change in preferences). Thus SCCs for $\tau \geq 1$ are random variables, which are computed at different values of S_τ for each planner.

4 Results

Our first main result – Figure 2 below – plots the distribution of SCCs in the initial period of the model as a function of the ‘non-dogmatism’ parameter x in (14). We work with

¹³To facilitate comparison to Nordhaus (2017) the units of our SCC are 2010\$/tCO₂. Multiply by 44/12 to convert to 2010\$/tC.




Figure 2: SCCs in the initial model year (2015) as a function of the degree of non-dogmatism. Marker sizes are proportional to cluster sizes.

the values $x = \{0, 1.25\%, 2.5\%, 5\%, 10\%\}/\text{yr}$, which correspond to a change in normative views once every $\{\infty, 80, 40, 20, 10\}$ years on average. We take the value $x = 2.5\%/ \text{yr}$ as a reasonable baseline specification, and present some results in more detail in this case. Even for this arguably conservative value of x non-dogmatism gives rise to a dramatic reduction in disagreement about the SCC. Dogmatic SCCs ($x = 0$) fall in the range 3-459\$/tCO₂, while non-dogmatic SCCs fall in the range 27-126\$/tCO₂ when $x = 2.5\%/ \text{yr}$, a 4.6-fold reduction in disagreement. For $x = 10\%/ \text{yr}$, the range of SCCs shrinks to 46-96\$/tCO₂, a ninefold reduction in disagreement.

Although a lot of this reduction in disagreement comes from the planner with preference parameters $\rho = 0.4\%/yr$, $\eta = 0.4$, this planner represents the fourth largest cluster. While this planner’s dogmatic SCC value is an outlier, the preference parameters that underlie it are not. Indeed, if we exclude outliers the results become even more dramatic. To demonstrate this, let $\mathcal{I}_{90}(x)$ denote the smallest SCC interval that contains 90% of the sample when the switching rate is x .¹⁴ This interval excludes SCC values associated with outlying preference parameters, which in practice means the very low SCC values that obtain due to very large values for ρ , η , or both (i.e., the yellow, dark blue, black, or light green clusters in Figure 1). Table 1 records the values of $\mathcal{I}_{90}(x)$. Excluding outliers, we find a 7.7-fold reduction in the range of SCCs when $x = 2.5\%/yr$, with substantially bigger reductions in disagreement as x increases. Although our results are not driven by extreme *preferences*, the fact that non-dogmatism has a very substantial effect on extreme *SCCs* illustrates the value of a little normative humility in the context of the climate debate, or conversely, how inflexible adherence to a dogmatic normative paradigm can hold up agreement.

To begin to understand the results in Figure 2, note that while all non-dogmatic planners agree on the social cost of marginal climate damages in the distant future (under the conditions of Theorem 1), the SCC sums discounted marginal damages over all maturities. Non-dogmatic planners’ discount rates still differ at shorter maturities, so they do not fully agree on the SCC, even for high values of the switching rate x .¹⁵

¹⁴In detail, let Ω be an arbitrary subset of cluster indices such that $\sum_{i \in \Omega} m_i \geq 0.9$, and let $SCC^i[x]$ be the initial SCC in cluster i when the non-dogmatism parameter is x . Let $\Omega^* = \operatorname{argmin}_{\Omega} (\max\{SCC^i[x] | i \in \Omega\} - \min\{SCC^i[x] | i \in \Omega\})$. Then we define $\mathcal{I}_{90}(x) = [\min\{SCC^i[x] | i \in \Omega^*\}, \max\{SCC^i[x] | i \in \Omega^*\}]$.

¹⁵As we use a 10 year time step, even with an annual switching rate of 100% a planner’s preferences

Table 1: Reduction in disagreement excluding outlying 10% of sample

x (%/yr)	$\mathcal{I}_{90}(x)$ (2010\$/tCO ₂)	$ \mathcal{I}_{90}(0) / \mathcal{I}_{90}(x) $
0	13-459	1
1.25	61-166	4.2
2.5	68-126	7.7
5	72-105	13.6
10	76-96	22.4

Consider the case $x = 10\%/yr$ in Figure 2; this case represents rapid mixing of preferences over time. SCCs in this case are broadly ordered by the pure rate of time preference (ρ): planners with a low ρ have a high SCC, while planners with a high ρ have lower SCCs. Notice that SCC values for the low ρ planners are clustered close together. These planners put a lot of weight on the long run, and they also agree on how to value long-run marginal payoffs; this gives rise to substantial agreement on the SCC. The high ρ planners are more separated from each other. They regard short-run outcomes as more important, and they also disagree about the value of these short-run outcomes.

Now consider lower values of the switching probability x , e.g. $x = 1.25\%/yr$ or $2.5\%/yr$. The low ρ planners' SCC values still exhibit significant convergence in these cases. This is again because they put a lot of weight on the long run, where they agree on how to value marginal payoffs. In effect, low ρ planners 'see' more of the temporal region where preferences (and hence discount rates) are well mixed.

change only after 10 years. This can also limit our estimates of the convergence of SCCs, especially for high switching rates.

Notice that the ranking of the SCCs of two planners may reverse when x is increased from zero. A planner with a low ρ but a high η may have a high consumption discount rate when they are dogmatic;¹⁶ their corresponding SCC may be well below that of a planner with a higher ρ but lower η . But with even a modest degree of non-dogmatism, a planner with a low ρ puts a lot of weight on distant future times where preferences are well mixed, and where, by Theorem 1, the lowest value of η dominates present value calculations (assuming positive expected consumption growth). Such a planner's SCC may thus increase substantially when x is increased from zero. By contrast, a high ρ and low η planner places less weight on the distant future where preferences are well mixed, and her 'effective' long-run value of η also has less far to fall when x is increased from zero. This results in a more modest adjustment of her SCC. Non-dogmatism also affects planners' expected values of η in the medium run. A modest increase in x from zero leads high η planners to account for the possibility of a switch to a more linear utility function in the medium run. Their response is to save more, including in terms of 'natural capital' (i.e., the CO₂ stock). Thus, planners with above average η (red, pink, light green and dark green) see a sharper increase in their SCC with the introduction of non-dogmatism. Conversely, this 'medium run' effect of non-dogmatism causes the SCC of the planner with the lowest η (cyan) to fall: the possibility of switching implies expected medium run marginal utility falls for this planner.

These different impacts of non-dogmatism on the 'impatience' and 'consumption smooth-

¹⁶Recall that in the dogmatic case the consumption discount rate at maturity t is $r_t \approx \rho + \eta g_t$, where g_t is the average consumption growth rate at maturity t .

ing’ aspects of preferences are further illustrated by the possible non-monotonic effect of an increase in x on a planner’s SCC value, as occurs for the Red planner with $\rho = 0.2\%/yr$, $\eta = 1.8$, for example. This planner has the lowest value of ρ , and a moderately high η . When x is increased from zero to say $1.25\%/yr$, Red’s SCC shoots up due the reduction in her expected value of η in both the medium- and long-run. Since her dogmatic value of ρ is so low, her effective impatience *increases* when x increases from zero, since non-dogmatic pure time preference rates mix in other preference parameters, and converge to a stationary average of all rates of time preference at long maturities. As x is increased further, this mixing happens faster leading to a more substantial increase in this planner’s effective impatience, which acts to lower her SCC. This effect works against the reduction in her effective η , which is the dominant contributor to the initial increase in her SCC. The interaction of these two opposing effects leads to the non-monotonic behaviour we see in Figure 2 for this planner. The fact that non-dogmatism may have these different impacts on impatience and consumption smoothing motives leads to ambiguity in the magnitude of the effect of non-dogmatism on SCC estimates for different planners. We investigate this in more detail below.

Of course, the previous discussion is premised on the conditions of Theorem 1 – in particular ergodicity – being approximately satisfied in our numerical model. To investigate this, as well as the differential impact of non-dogmatism on planners’ impatience and propensity to smooth consumption, it is useful to rewrite our expression (15) for the SCC. The SCC is the current value of a small change in cumulative CO₂ emissions. A small change in cumulative emissions can be thought of as leaving the set of trajectories that

the economy may follow undisturbed, but giving rise to small consumption perturbations along those trajectories. Denote the consumption change that a kick in CO_2 gives rise to at maturity t along a trajectory χ by π_t^χ . Then, as in (7), we can write the SCC according to planner i as:

$$SCC(S_\tau, i) = \mathbb{E}_{\chi|(S_\tau, i)} (1 + r(t|\chi))^{-t} \pi_t^\chi.$$

where the expectation is over all trajectories that emanate from (S_τ, i) .

Because the consumption perturbations π_t^χ are marginal (and hence exogenous to planner's choices), they depend only on the trajectory of physical states $\{S_{\tau+t}\}_{t=0, \dots, \infty}$; they don't vary across future realisations of the planner's preferences. Thus, all the variation in the SCC that we observe across non-dogmatic planners can be attributed to variations in the set of trajectories that the economy follows, and the associated variation in discount rates $r(t|\chi)$ along those trajectories. If our model is (approximately) ergodic, we should see the statistics of these sets of trajectories converging at long maturities, irrespective of the initial conditions. As our discussion of the proof of Theorem 1 indicated, this will manifest in convergence of the *distributions* of discount rates across non-dogmatic planners at large maturities. We can investigate whether this occurs in our model by simulating the trajectory of the economy many times for each of the M initial values of the state (S_0, i) , and plotting statistics of the distribution of discount rates $r(t|\chi)$ as a function of maturity t . Figure 3 illustrates the results of this exercise for $x = 2.5\%/yr$.

The vertical lines in the top panel of Figure 3 are the 90% confidence intervals for the consumption discount rates $r(t|\chi)$, plotted at several maturities for each of the M non-

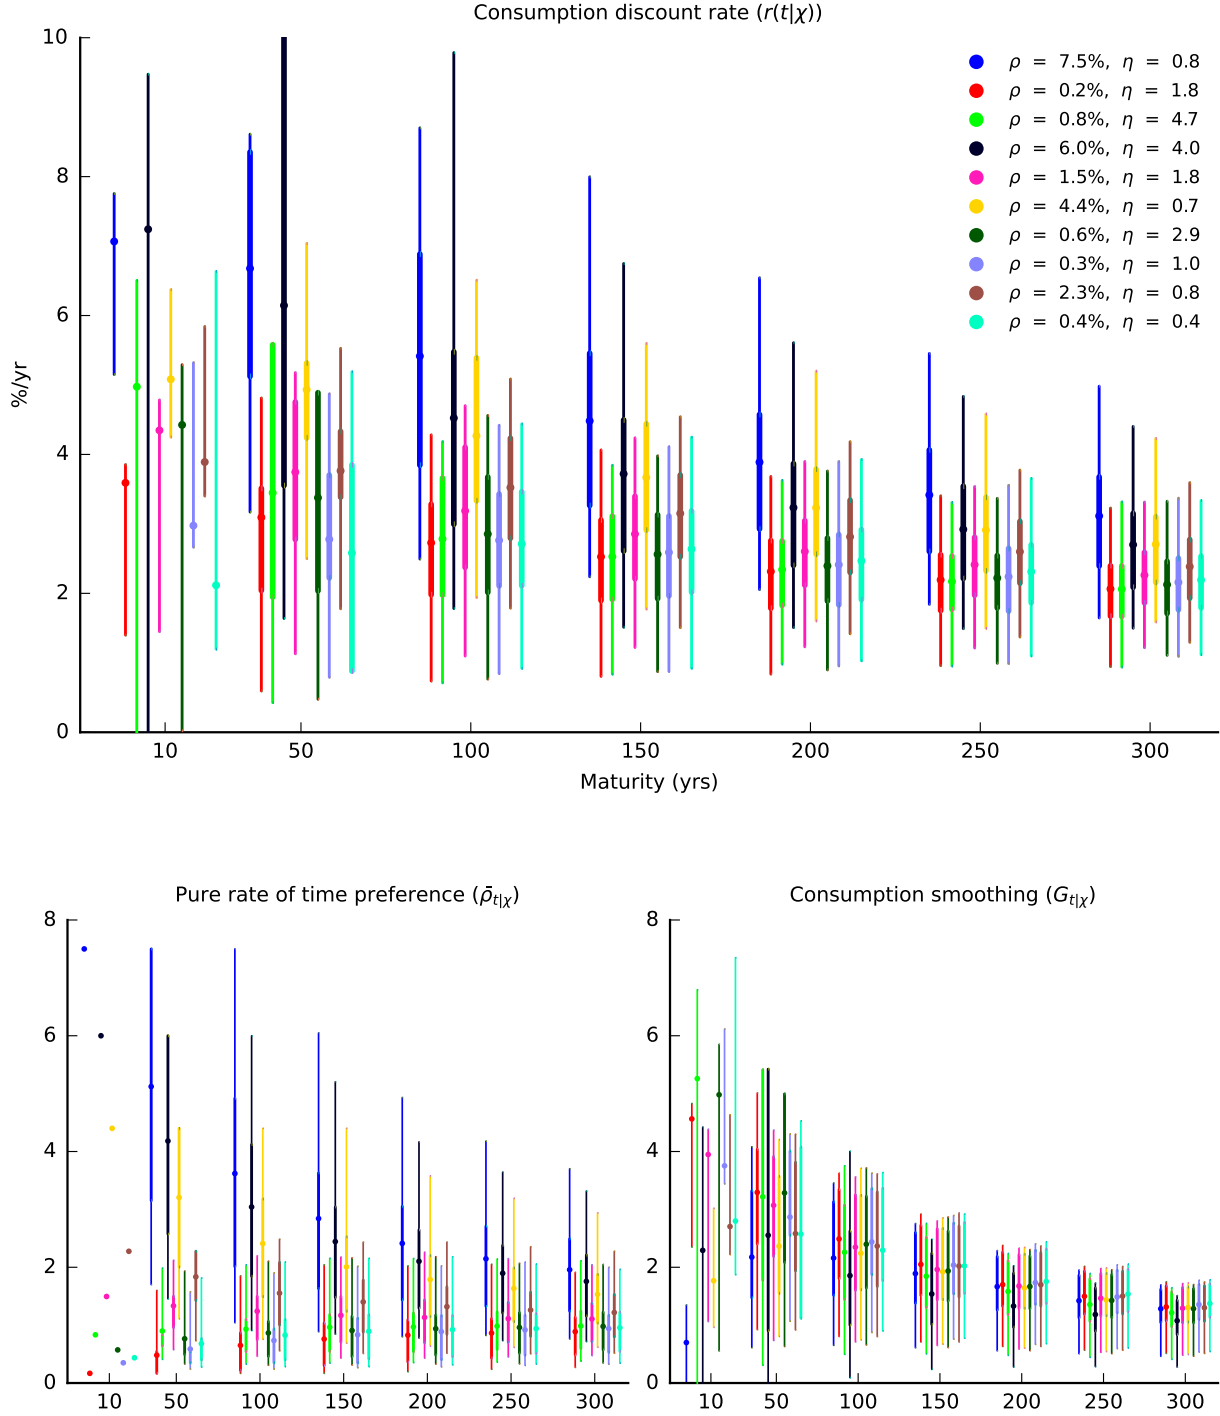


Figure 3: Endogenous distributions of consumption discount rates as a function of maturity for non-dogmatic planners ($x = 2.5\%$). Top: Consumption discount rate $r(t|\chi)$, defined in (6). Bottom left: Pure rate of social time preference $\bar{\rho}_{t|\chi}$, defined in (16). Bottom right: Consumption smoothing term $G_{t|\chi}$, defined in (17). Vertical lines at each maturity denote 5-95% range of discount rates for each planner, with interquartile ranges in bold. Mean values are depicted by solid dots.

dogmatic planners at $\tau = 0$. The figure demonstrates several features of the model. First, it shows that our numerical model is approximately ergodic: the distributions of non-dogmatic planners' consumption discount rates exhibit significant convergence at large maturities. The main factor that prevents our model from being fully ergodic is our use of a finite time horizon; this prevents full asymptotic mixing of trajectories on state space.¹⁷ Nevertheless, because the Markov chain \mathbf{w} mixes preferences and optimal policies over time, we still see substantial convergence of discount rate distributions.

The bottom panels in Figure 3 provide a decomposition of the variability in the discount rates $r(t|\chi)$ into a component due to pure time preference and a component due to consumption smoothing. Equation (20) in Appendix B shows that we can write

$$\begin{aligned} r(t|\chi) &= (1 + \bar{\rho}_{t|\chi})(1 + G_{t|\chi})(1 + \bar{\lambda}_{t|\chi})^{-1} - 1 \\ &\approx \rho_{t|\chi} + G_{t|\chi} - \bar{\lambda}_{t|\chi} \quad \text{when } |\rho_{t|\chi}|, |G_{t|\chi}|, |\bar{\lambda}_{t|\chi}| \ll 1, \end{aligned}$$

where we define

$$\bar{\rho}_{t|\chi} = \left(\prod_{l=0}^{t-1} \beta^{i_l^x} \right)^{-\frac{1}{t}} - 1, \quad (16)$$

$$G_{t|\chi} = \left[\frac{(c_{t|\chi})^{-\eta^x}}{c_0^{-\eta^i}} \right]^{-\frac{1}{t}} - 1, \quad (17)$$

$$\bar{\lambda}_{t|\chi} = \left(\frac{P_{t|\chi}}{P_0} \right)^{\frac{1}{t}}. \quad (18)$$

¹⁷We use a finite horizon to ensure computational tractability. A second (related) factor that curtails convergence is that because cumulative emissions don't depreciate, our simple climate model has a long memory of 'transient' past emissions. A third factor is the presence of exogenous time series for e.g. population and TFP growth in DICE, which only approach steady state values at large maturities.

The quantity $\bar{\rho}_{t|\chi}$ is the maturity t pure rate of time preference along trajectory χ . To interpret $G_{t|\chi}$ choose units so that $c_0 = 1$,¹⁸ and write $c_{t|\chi} = (1 + g_{t|\chi})^t$. Then we have

$$G_{t|\chi} = (1 + g_{t|\chi})^{\eta_t^{\chi}} - 1 \approx \eta_t^{\chi} g_{t|\chi} \text{ when } |g_{t|\chi}| \ll 1.$$

Thus $G_{t|\chi}$ is the standard consumption smoothing term in the Ramsey formula for the consumption discount rate; it reflects aversion to intertemporal consumption inequalities along a trajectory χ . Finally $\bar{\lambda}_{t|\chi}$ is the maturity t population growth rate on trajectory χ . Since population is exogenous in our model, this term does not vary with the trajectory χ , and is common to all non-dogmatic planners. Population growth rates can thus be safely neglected when explaining the variation in discount rates across planners in Figure 3.

The bottom left (right) panel of Figure 3 plots the distributions of $\bar{\rho}_{t|\chi}$ ($G_{t|\chi}$) at several maturities, for each of the non-dogmatic planners. Notice that these two components of discount rates are not ordered uniformly across planners (i.e., high values of $\bar{\rho}$ do not necessarily correspond to high values of G), and they also behave differently as functions of maturity. In particular, there is faster convergence to a common long-run distribution in the consumption smoothing term ($G_{t|\chi}$) than in the pure time preference term ($\bar{\rho}_{t|\chi}$).¹⁹ This confirms our discussion above; since non-dogmatic planners exhibit substantive agreement on the distribution of the consumption smoothing term $G_{t|\chi}$ even at relatively short

¹⁸This is without loss of generality.

¹⁹This is consistent with the findings of Millner (2020). Although this finding is contingent on empirical details (and in particular on the endogenous consumption growth rates that prevail in our model), the online appendix of Millner (2020) provides a qualitative discussion of why this might occur in a simplified setting.

maturities (e.g. 50 years), much of the divergence in their SCCs is explained by residual disagreements about pure time preference. Nevertheless, the bottom left panel of Figure 3 shows that planners with low values of ρ exhibit much more agreement on the long-run distribution of non-dogmatic PRSTPs than their high ρ counterparts, and they also assign more weight to long maturities where agreement on discount rates is strongest. The outliers with high values of ρ (in blue, black, and yellow in the figure) exhibit more persistent disagreement, which is still tempered at large maturities – even these planners largely agree on the distribution of $G_{t|\chi}$ for large t . However, these planners place less weight on long maturities when computing the SCC, and so their SCC values differ more even with some non-dogmatism.

Figure 3 also helps to illustrate how reversals in the order of planners' SCCs can occur as x increases from zero. Consider the Pink planner with $(\rho, \eta) = (1.5\%/yr, 1.8)$ and the Brown planner with $(\rho, \eta) = (2.3\%/yr, 0.8)$. Figure 2 shows that when $x = 0$ Pink has a lower SCC than Brown, due to her larger value of η .²⁰ But when $x = 1.25\%/yr$, the order of their SCCs reverses. A careful examination of Figure 3 demonstrates how this can occur. At short maturities (i.e., $t = 10$), where a planner's dogmatic views still dominate her discount rates, the mean of Pink's consumption discount rate distribution is *above* Brown's. But for $t \geq 50$ Pink's discount rate distribution drops below Brown's, leading to a higher SCC. The bottom panel of Figure 3 shows that this occurs because these planners' values of ρ and η are not ranked uniformly, and their non-dogmatic counterparts $\bar{\rho}$ and G

²⁰For a representative consumption growth rate $g \approx 2\%/yr$, Pink's discount rate is $1.5 + 2 \times 1.8 = 5.1\%/yr$, while Brown's is $2.3 + 2 \times 0.8 = 3.9\%/yr$.

have different rates of convergence to their long-run distributions.

Finally, while our discussion until now has focussed on SCCs in the initial model year, in Figure 4 we illustrate that reductions in disagreement due to non-dogmatism may be larger in the future. Since the trajectory of the economy in our model is stochastic, SCCs in all except the first period are random variables. To provide an illustration of possible future reductions in the SCC that is comparable to our analysis in Figure 2, we thus specify a fixed reference trajectory for the economy. We take this to be the optimal path according to a dogmatic planner whose preference parameters coincide with those used in Nordhaus (2017). We compute SCCs at future times t , assuming that the interim evolution of the economy until t is given by the reference trajectory. Figure 4 depicts the results of this analysis, showing that the effects of non-dogmatism are even larger at later times along the reference trajectory. This is due to the fact that marginal damages are larger in the future (due to a larger economy and higher climate damages); this amplifies the importance of discount rate disagreements for present values, and hence the SCC.

5 Conclusion

This paper presented a normative model of climate policy in which devotees of diverse theories of intertemporal social welfare are non-dogmatic – they admit the possibility of a change in their normative views, and internalise the preferences of possible future selves. The model requires planners (or economists) to exhibit a little humility about their preferred framework for evaluating climate policies. Despite this, the model still allows an-




Figure 4: SCC values along a fixed trajectory for the state variables given by the optimal path according to a dogmatic planner with DICE preferences ($\rho = 1.5\%$, $\eta = 1.45$). Filled dots correspond to $x = 0$, open circles to $x = 2.5\%$, and stars to $x = 10\%$. The dashed line marks the DICE SCC.

analysts the freedom to advocate their preferred normative theory unequivocally. Although disagreements about contested welfare parameters – the pure rate of social time preference and elasticity of marginal utility – remain intact, planners who embrace non-dogmatism will exhibit significantly less disagreement about the SCC. Disagreements reduce by almost a factor of five if planners admit the possibility of a change in views once every forty years, and substantially more if they are more non-dogmatic. While non-dogmatism cannot deliver a universally acknowledged ‘best’ value of the SCC, it can focus policy evaluation on a much narrower range of values, which nevertheless reflects a plurality of normative views. This can hopefully help to move the debate on optimal climate policy past ethics, and on

to practicalities.

References

- B. Anderson, et al. (2014). ‘Uncertainty in Climate Change Modeling: Can Global Sensitivity Analysis Be of Help?’. *Risk Analysis* **34**(2):271–293.
- C. P. Chambers & F. Echenique (2018). ‘On Multiple Discount Rates’. *Econometrica* **86**(4):1325–1346.
- M. Collins, et al. (2013). ‘Long-term climate change: Projections, Commitments and Irreversibility’. In *Climate Change 2013 - The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, chap. 12, pp. 1029–1136. Cambridge University Press.
- I. Cornfeld, et al. (1982). *Ergodic Theory*. Springer.
- P. Dasgupta (2008). ‘Discounting climate change’. *Journal of Risk and Uncertainty* **37**(2):141–169.
- S. Dietz & N. A. Matei (2016). ‘Spaces for agreement: A theory of time-stochastic dominance and an application to climate change’. *Journal of the Association of Environmental and Resource Economists* **3**(1):85–130.
- S. Dietz & F. Venmans (2019). ‘Cumulative carbon emissions and economic policy: In

- search of general principles'. *Journal of Environmental Economics and Management* **96**:108–129.
- M. Drupp, et al. (2018). 'Discounting Disentangled'. *American Economic Journal: Economic Policy* **10**(4):109–134.
- T. Feng & S. Ke (2018). 'Social Discounting and Intergenerational Pareto'. *Econometrica* **86**(5):1537–1567.
- M. Fleurbaey, et al. (2018). 'The Social Cost of Carbon: Valuing Inequality, Risk, and Population for Climate Policy'. *The Monist* **102**(1):84–109.
- M. C. Freeman & B. Groom (2015). 'Positively Gamma Discounting: Combining the Opinions of Experts on the Social Discount Rate'. *The Economic Journal* **125**(585):1015–1024.
- S. Galperti & B. Strulovici (2017). 'A Theory of Intergenerational Altruism'. *Econometrica* **85**(4):1175–1218.
- C. Gollier & R. Zeckhauser (2005). 'Aggregation of Heterogeneous Time Preferences'. *Journal of Political Economy* **113**(4):878–896.
- G. Hardy, et al. (1934). *Inequalities*. Cambridge University Press.
- G. M. Heal & A. Millner (2014). 'Agreeing to disagree on climate policy'. *Proceedings of the National Academy of Sciences* **111**(10):3695–3698.

- H. Hopenhayn & E. C. Prescott (1992). ‘Stochastic monotonicity and stationary distributions for dynamic economies’. *Econometrica* **60**(6):1387–1406.
- T. H. Johnsen & J. B. Donaldson (1985). ‘The Structure of Intertemporal Preferences under Uncertainty and Time Consistent Plans’. *Econometrica* **53**(6):1451–1458.
- K. L. Judd (1998). *Numerical Methods in Economics*. The MIT Press.
- T. Kamihigashi & J. Stachurski (2014). ‘Stochastic stability in monotone economies’. *Theoretical Economics* **9**:383–407.
- L. Ljungqvist & T. J. Sargent (2004). *Recursive Macroeconomic Theory*. MIT Press, Cambridge, Mass, 2nd edn.
- H. D. Matthews, et al. (2009). ‘The proportionality of global warming to cumulative carbon emissions’. *Nature* **459**:829–832.
- A. Millner (2020). ‘Nondogmatic Social Discounting’. *American Economic Review* **110**(3):760–775.
- A. Millner & G. Heal (2018). ‘Discounting by committee’. *Journal of Public Economics* **167**:91–104.
- A. Millner & G. Heal (2021). ‘Choosing the future: Markets, Ethics, and Rapprochement in Social Discounting’. *NBER Working Paper*, <https://tinyurl.com/9s6cyfdj>.
- W. Nordhaus & J. Boyer (2000). *Warming the World: Economic Models of Climate Change*. MITPress, Cambridge, MA, USA.

- W. D. Nordhaus (2007). ‘A Review of the Stern Review on the Economics of Climate Change’. *Journal of Economic Literature* **45**(3):686–702.
- W. D. Nordhaus (2008). *A question of balance*. Yale University Press.
- W. D. Nordhaus (2017). ‘Revisiting the social cost of carbon’. *Proceedings of the National Academy of Sciences* **114**(7):1518–1523.
- D. Ray (1987). ‘Nonpaternalistic intergenerational altruism’. *Journal of Economic Theory* **41**(1):112–132.
- K. L. Ricke & K. Caldeira (2014). ‘Maximum warming occurs about one decade after a carbon dioxide emission’. *Environmental Research Letters* **9**(12):124002.
- M. Saez-Marti & J. W. Weibull (2005). ‘Discounting and altruism to future decision-makers’. *Journal of Economic Theory* **122**(2):254–266.
- S. H. Schneider & S. L. Thompson (1981). ‘Atmospheric CO₂ and climate: Importance of the transient response’. *Journal of Geophysical Research Oceans* **86**(C4):3135–3147.
- N. Stern (2008). ‘The Economics of Climate Change’. *American Economic Review* **98**(2):1–37.
- N. L. Stokey, et al. (1989). *Recursive methods in economic dynamics*. Harvard University Press.
- M. L. Weitzman (1998). ‘Why the Far-Distant Future Should Be Discounted at Its Lowest Possible Rate,’. *Journal of Environmental Economics and Management* **36**(3):201 – 208.

M. L. Weitzman (2001). ‘Gamma Discounting’. *The American Economic Review* **91**(1):260–271.

M. L. Weitzman (2007). ‘A Review of The Stern Review on the Economics of Climate Change’. *Journal of Economic Literature* **45**(3):703–724.

A Time consistency

Since non-dogmatic planners' preferred discount factors may change over time, intuition might suggest that their preferences are necessarily time inconsistent. However, this is not the case. We illustrate this in the context of the more general model of non-dogmatic social time preferences in Millner (2020). Millner (2020) considers preferences over deterministic consumption streams; we extend these to preferences over state-contingent consumption plans. It is perhaps easiest to make the link between preferences over deterministic and state-contingent plans by using trajectory-specific value functions similar to those in (5) as an intermediate step. Given a deterministic trajectory $\chi = \{(S_\tau, i_\tau)\}_{\tau=0, \dots, \infty}$, define

$$W_\chi^{(S_\tau, i_\tau)}(\mathbf{c}) = U^{i_\tau}(c(S_\tau, i_\tau), S_\tau) + \sum_{t=1}^{\infty} \beta_t^{i_\tau} W_\chi^{(S_{\tau+t}, i_{\tau+t})}(\mathbf{c}),$$

where now we have an infinite sequence of discount factors $\beta_t^i > 0$ on future wellbeings at all future times $\tau + t$, for each value of i . We assume that $\max_i \sum_{t=1}^{\infty} \beta_t^i < 1$; this ensures that preferences are complete on the set of bounded utility streams and increasing in all utilities. Preferences over state contingent plans are then given by

$$\begin{aligned} W^{(S_\tau, i_\tau)}(\mathbf{c}) &= \mathbf{E}_\chi W_\chi^{(S_\tau, i_\tau)}(\mathbf{c}) \\ &= U^{i_\tau}(c(S_\tau, i_\tau), S_\tau) + \sum_{t=0}^{\infty} \beta_t^{i_\tau} \mathbf{E}_{(S_{\tau+t}, i_{\tau+t})} W^{(S_{\tau+t}, i_{\tau+t})}(\mathbf{c}), \end{aligned} \quad (19)$$

where $\mathbf{E}_{(S_{\tau+t}, i_{\tau+t})}$ denotes an expectation over future values of the state $(S_{\tau+t}, i_{\tau+t})$. The model in (19) is a natural generalisation of that in Millner (2020). The major difference

is that in Millner (2020) consumption does not depend on the state (S, i) ; in that paper we have $c(S, i) = c(t)$, i.e., consumption is an exogenous function of maturity t . The expectation $\mathbb{E}_{(S_{\tau+t}, i_{\tau+t})}$ reduces to $\mathbb{E}_{i_{\tau+t}}$ in that case, since we do not need to marginalise over the physical states $S_{\tau+t}$, although valuations of the exogenous consumption stream still depend on the planner's future preference $i_{\tau+t}$. As Millner (2020) observes, the preferences (19) are the unique non-dogmatic preferences (i.e., preferences that internalise the values of future selves) that are both time invariant (i.e., the utility functions U^i and sequence of discount factors β_t^i do not depend on calendar time τ), and additively separable in time and states of the world.

When do the preferences (19) satisfy the definition of time consistency in (3)?

Proposition 1. *The preferences in (19) are time consistent if and only if $\beta_s^i = 0$ for all $s \geq 2$, $i = 1 \dots N$.*

Proof. Johnsen & Donaldson (1985) show that history independent and consequentialist²¹ preferences over state-contingent consumption plans \mathbf{c} are time consistent iff they can be represented by

$$V_\tau^i(x, \mathbf{c}) = F_\tau^i(x, \{V_{\tau+1}^j(\mathbf{c})\} | j \in \mathcal{I}_{\tau+1})$$

where x is current consumption, \mathbf{c} is an arbitrary state-contingent consumption plan, and

$\mathcal{I}_{\tau+1}$ indexes a set of events (preference changes in our setting) that occur at the begin-

²¹Preferences are consequentialist if they only depend on reachable nodes of a decision tree, i.e., they do not depend on unrealised events or plans not carried out. They are history independent if preferences at time τ do not depend *explicitly* on actions at times before τ . The preferences in (2) are consequentialist, since current value $W^{(S_\tau, i_\tau)}$ depends only on current consumption and future consumption in states that are reachable with non-zero probability from (S_τ, i_τ) ; they do not depend on past unrealised plans or events. Preferences are also history independent – they only depend on history indirectly through the dependence of the consumption plan \mathbf{c} on the state variables S .

ning of period $\tau + 1$. Applying this constraint on preferences to (2) yields the result, by inspection. \square

Thus, although non-dogmatic planners' utility functions and discount factors change over time, their preferences over state-contingent plans need not be time inconsistent. This is a direct consequence of preference internalisation: if current planners internalise the preferences of future selves in the next time period only, their preferences satisfy a recurrence relation that is reminiscent of a stochastic Bellman equation; they are thus time consistent.

B Proof of Theorem 1

The proof of this result builds on the defining property of ergodic dynamical systems – Birkhoff's Ergodic Theorem. Given a state-contingent plan \mathbf{c} , transition matrix \mathbf{w} , and an initial state (S_τ, i_τ) , define a trajectory in state space as the sequence $\chi = \{(S_{\tau+t}, i_{\tau+t})\}_{t=0, \dots, \infty}$, where for each time $\tau + t + 1$, $(S_{\tau+t+1}, i_{\tau+t+1})$ is in the support of the transition function $T(S', i | S_{\tau+t}, i_{\tau+t})$ in (4). Birkhoff's theorem can be stated as follows:

Theorem 2 (Birkhoff's Ergodic Theorem). *Suppose that the Markov process (4) associated with (\mathbf{c}, \mathbf{w}) is ergodic, and denote its stationary distribution by $p_\infty(S, i)$. Then for any function $h(S, i)$ on state space, and any trajectory $\chi = \{(S_{\tau+t}, i_{\tau+t})\}_{t=0, \dots, \infty}$, we have*

$$\lim_{t \rightarrow \infty} \frac{1}{t} \sum_{t=0}^{\infty} h(S_{\tau+t}, i_{\tau+t}) = \mathbb{E}_\infty h(S, i)$$

almost surely, where E_∞ denotes an expectation taken with respect to $p_\infty(S, i)$.

In words, ergodicity implies that time averages across trajectories are equivalent to expectations across state space taken with the stationary distribution, for almost all trajectories. See e.g. Cornfeld et al. (1982) for a proof of this theorem.

Given a trajectory χ that starts at (S_τ, i_τ) , denote the t -th term of this sequence by $\chi_t = (S_t^\chi, i_t^\chi)$. In addition denote the 1-period population and consumption growth rates at time $\tau + t$ along this trajectory by λ_t^χ and g_t^χ respectively. Then on this trajectory we can write population and consumption at time $\tau + t$ as

$$P_t^\chi = \prod_{n=1}^t (1 + \lambda_n^\chi) P_\tau$$

$$c_t^\chi = \prod_{m=1}^t (1 + g_m^\chi) c_\tau$$

where P_τ and c_τ are population and consumption in the initial state (S_τ, i_τ) . Similarly, the discount factor that preference $V^{(S_\tau, i_\tau)}$ in (2) places on state $(S_{\tau+t}, i_{\tau+t})$ along trajectory χ is

$$\delta_t^\chi \equiv \prod_{l=0}^{t-1} \beta^{i_l^\chi},$$

and the probability of being on trajectory χ at time $\tau + t$ is

$$z_t^\chi \equiv \prod_{l=0}^{t-1} w_{i_l^\chi, i_{l+1}^\chi}.$$

Using these definitions, the non-dogmatic preferences in (2) can be written in terms of

sums over trajectories as follows:

$$V^{(S_\tau, i_\tau)} = \sum_{t=0}^{\infty} \sum_{\chi} z_t^\chi P_t^\chi \delta_t^\chi U^{i_t^\chi}(c_t^\chi),$$

where all trajectories χ in this sum start at (S_τ, i_τ) .

Consider a trajectory χ with $\chi_t = (S, j)$, i.e., (S, j) is the state χ reaches at time $\tau + t$.

The consumption discount rate that is applied to this state, given χ , is

$$\begin{aligned} r^\chi(t|S, j) &= \left(\frac{P_t^\chi \delta_t^\chi (U^j)'(c(S, j))}{P_\tau c_\tau^{-\eta^{i_\tau^\chi}}} \right)^{-\frac{1}{t}} - 1 \\ &= \left(\frac{[\Pi_{n=1}^t (1 + \lambda_n^\chi) P_\tau] [\Pi_{l=0}^{t-1} \beta^{i_{\tau+l}^\chi}] (\Pi_{m=1}^t (1 + g_m^\chi) c_\tau)^{-\eta^j}}{P_\tau c_\tau^{-\eta^{i_0^\chi}}} \right)^{-\frac{1}{t}} - 1. \end{aligned} \quad (20)$$

We are interested in $\lim_{t \rightarrow \infty} r^\chi(t|S_{\tau+t}, j)$. We use Birkhoff's theorem to evaluate this limit as follows:

$$\begin{aligned} \log [r^\chi(t|S, j) + 1] &= -\frac{1}{t} \left(\sum_{n=1}^t \log(1 + \lambda_n^\chi) + \sum_{l=0}^{t-1} \log \beta^{i_{\tau+l}^\chi} - \sum_{m=1}^t \eta^j \log(1 + g_m^\chi) \right) \\ &\quad - \frac{1}{t} (\log P_\tau - \eta^j \log c_\tau - \log P_\tau + \eta^i \log c_\tau) \\ &\rightarrow_{t \rightarrow \infty} -\mathbf{E}_\infty \left[\log(1 + \tilde{\lambda}) + \log(\tilde{\beta}) - \eta^j \log(1 + \tilde{g}) \right] \end{aligned}$$

almost surely. In words, if the state (S, j) is sufficiently far in the future, the discount rate on that state does not depend on the initial state (S_τ, i_τ) , on the future value of the physical state S , or on the trajectory χ that led to that state, almost surely. In particular, long-run discount rates are independent of the planner's current preferred welfare measure

i_τ . Given this result, and writing $\tilde{\beta} = (1 + \tilde{\rho})^{-1}$ we define the long-run discount rate on a distant future state (S, j) as

$$\hat{r}(j) \equiv \exp \left(\mathbb{E}_\infty \left[\log \frac{1 + \tilde{\rho}}{1 + \tilde{\lambda}} + \eta^j \log(1 + \tilde{g}) \right] \right) - 1.$$

Now let us consider the ‘total’ discount rate on a state (S, j) that occurs at time $\tau + t$, which we denote by $d(t|S, j)$. This is related to the probability weighted sum of discount factors over all trajectories that have $\chi_t = (S, j)$, i.e.

$$(1 + d(t|S, j))^{-t} = \sum_{\chi|\chi_t=(S,j)} z_t^\chi (1 + r^\chi(t|S, j))^{-t}$$

Since $r^\chi(t|S, j)$ is independent of χ when $t \rightarrow \infty$, we can factor it out of this expression when t is large. By ergodicity, the quantity $\sum_{\chi|\chi_t=(S,j)} z_t^\chi$ must approach the stationary density $p_\infty(S, j)dS$ as $t \rightarrow \infty$, almost surely.²² Thus we have

$$(1 + d(t|S, j))^{-t} \sim_{t \rightarrow \infty} (1 + \hat{r}(j))^{-t} p_\infty(S, j)dS$$

almost surely.

Consider cost benefit analysis of a project with long-run payoffs $\pi(S, j)$ in a distant state future (S, j) . If marginal rates of substitution and project payoffs are bounded in all states (as we assumed in the statement of Theorem 1 when we required consumption to be bounded above zero), we may neglect measure zero trajectories when computing a

²²This holds provided that (S, j) is in the support of $p_\infty(S, j)$.

project's present value. Thus, according to a planner who currently advocates theory i , the present value of these payoffs are

$$\int_S \sum_j (1 + d(t|S, j))^{-t} \pi(S, j) dS \rightarrow_{t \rightarrow \infty} \int_S \sum_j (1 + \hat{r}(j))^{-t} \pi(S, j) p_\infty(S, j) dS \quad (21)$$

Since the discount rates $\hat{r}(j)$ and stationary probability distribution $p_\infty(S, j)$ do not depend on the initial state (S_τ, i_τ) , all planners agree on the present value of these long-run payoffs.

To make contact between this result and the statement of the result in Theorem 1, assume now that project payoffs depend on the physical state S but not the planner's preferences j , i.e., $\boldsymbol{\pi} = \pi(S)$. In this case we can marginalise over the physical states S in (21) to write the right hand side of this expression as

$$\mathbb{E}_{\infty, j} (1 + \hat{r}(j))^{-t} \mathbb{E}_{\infty, S} \pi(S)$$

where $\mathbb{E}_{\infty, j}$ denotes an expectation over j taken with respect to the stationary marginal distribution $\int_S p_\infty(S, j) dS$, and $\mathbb{E}_{\infty, S}$ denotes an expectation over S taken with respect to the stationary marginal distribution $\sum_j p_\infty(S, j)$.

Now it is a standard result (see e.g. Hardy et al., 1934) that

$$\lim_{t \rightarrow \infty} (\mathbb{E}_{\infty, j} (1 + \hat{r}(j))^{-t})^{-\frac{1}{t}} = \min_j \{1 + \hat{r}(j)\}$$

and hence in this case we can write long-run project present value as

$$(1 + \hat{r})^{-t} \mathbf{E}_{\infty, S} \pi(S)$$

where

$$\hat{r} \equiv \min_j \hat{r}(j) = \exp \left(\min_j \left\{ \mathbf{E}_{\infty} \left[\log \frac{(1 + \tilde{\rho})(1 + \tilde{g})^{\eta_j}}{1 + \tilde{\lambda}} \right] \right\} \right) - 1.$$

C Details of the climate-economy model

We here detail the exact numerical model we use in the paper. We closely duplicate the structure of DICE-2016R September 2016 (downloaded from William Nordhaus's website).

We adjust the time period to $\Delta = 10$ years, instead of 5. Some of the growth rates associated with parameters have been computed or adjusted to reflect this. We also simplify the carbon cycle and climate module, so that surface temperature deviation in our model depends only on cumulative carbon emissions. Finally, given the simplified climate module in which we do not explicitly consider radiative forcing, we also omit the non-CO₂ forcings.

Population. Population P_t (in billions) at time t depends on a population growth parameter ν^P and the ultimate population size P^{∞} , through:

$$P_{t+1} = P_t (P^{\infty} / P_t)^{\nu^P},$$

$$\nu^P = 1 - .97164^{\Delta},$$

$$P_1 = 7.403, \quad P^{\infty} = 11.5.$$

TFP. Total factor productivity (TFP) at time t , denoted by A_t , grows at a declining rate. The sequence of growth rates depends on the parameter vector ν_t^A :

$$\begin{aligned} A_{t+1} &= A_t / (1 - \nu_t^A), \\ \nu_t^A &= \nu_1^A \exp(-.005 \Delta(t - 1)), \\ A_1 &= 5.115, \quad \nu_1^A = (1.0148)^\Delta - 1. \end{aligned}$$

Climate. Temperature is given by

$$T_t = 2CO_{2,t},$$

in which $CO_{2,t}$ is the stock of cumulative carbon emissions at the beginning of period t . The evolution of cumulative emissions depends on annual energy-related and land-use emissions, respectively E_t and E_t^{land} :

$$CO_{2,t+1} = CO_{2,t} + \Delta(E_t + E_t^{\text{land}})/1000,$$

in which the division by 1000 translates emissions into GtC/year.

Emissions. Energy-related annual carbon emissions are given by

$$E_t = \sigma_t(1 - \mu_t)F_t,$$

where μ_t is the emissions control rate, and F_t is gross output (defined below). The carbon

intensity of GDP σ_t grows according to

$$\begin{aligned}\sigma_{t+1} &= \sigma_t \exp(\nu_t^\sigma \Delta), \\ \nu_{t+1}^\sigma &= \nu_t^\sigma (1 - .001)^\Delta, \\ \sigma_1 &= \frac{35.85}{105.5(1 - .03)} \frac{12}{44}, \quad \nu_1^\sigma = -.0152.\end{aligned}$$

Land-use annual emissions are given exogenously by

$$\begin{aligned}E_t^{\text{land}} &= E_1^{\text{land}} (1 - .115)^{\frac{\Delta}{5}(t-1)}, \\ E_1^{\text{land}} &= 2.6(12/44).\end{aligned}$$

Output. The production function is Cobb-Douglas:

$$F_t(K, P) = A_t K^{.3} P^{.7}.$$

where K is capital and P is population.

Net output. Output net of abatement cost and damages is

$$Y_t(K, P, \mu, T) = (1 - \Omega(T) - \Xi(\mu)) F_t(K, P)$$

with climate damages (as fraction of GDP) $\Omega(T)$ determined by

$$\Omega(T) = .00236 T^2,$$

and abatement costs (as fraction of GDP) $\Xi(\mu)$ dependent on the rate of emission abatement μ :

$$\Xi(\mu) = \theta_{1,t} \mu^{2.6}$$

. The parameter $\theta_{1,t}$ is given by

$$\theta_{1,t} = (b_t/1000)\sigma_t/2.6.$$

The sequence b_t represents the time dependent marginal cost of 100% abatement, and is given by

$$b_t = b_1(1 - .005051)^{(t-1)\Delta}, \quad b_1 = 550(44/12).$$

Multiplication by 44/12 converts the units from per tonnes of CO₂ into per tonnes of carbon.

Utility. Utility functions are isoelastic in per-capita consumption c :

$$U(c) = \frac{(\nu^u c)^{1-\eta}}{1-\eta},$$

with the scaling parameter $\nu^u = 0.1$.

Capital. Capital accumulation follows

$$K_{t+1} = (1 - \delta)K_t + \Delta(Y_t - c_t P_t),$$

with $c_t P_t$ being aggregate consumption, and the depreciation factor being

$$\delta = 1 - (1 - .10)^\Delta.$$

The initial capital stock is $K_1 = 223$ (trillions of dollars), and the initial cumulative emissions $CO_{2,1} = .5$ (teratons of carbon).

D Details of the numerical solution

We solve the model using numerical stochastic dynamic programming. The model's time step is 10 years, the initial model year is 2015, and the terminal period is taken to be $T = 50$, corresponding to the year 2515, at which point the economy comes to an end.²³ The final period value function for each type i is given by the utility of consuming all output plus any remaining undepreciated capital. By approximating the value functions at any $t \leq T$, we can solve the problem for each type at time $t - 1$. We iterate until $t = 0$.²⁴

Value functions are approximated by Chebyshev collocation. That is, we choose a degree of approximation (N_K, N_{CO_2}) and approximate each value function $V^{i,t}$ by $\tilde{V}^{i,t}(K, CO_2) \equiv$

²³We could use an alternative terminal condition, for example that the economy is forced to reach a steady state at this point. Given the DICE assumptions about TFP growth, and conservatively using the preference parameter combination $\min_i \rho_i, \min_i \eta_i$, the effects due to the choice of terminal condition are small.

²⁴Since our model is not ratio-scale invariant in consumption (i.e., rescaling consumption values does not leave preferences unchanged), we must pick consumption units for our simulations that give rise to reasonable wealth effects. We don't want to choose units that give rise to artificial large differences in marginal utilities across types in the initial period when all planners are in an equal *a priori* position (i.e., they all start from the same initial values of the state variables). In future periods, where the values of state variables are determined endogenously by modelled events, divergences in marginal utility across planners are economically meaningful – they capture wealth effects across types that are consequences of the model structure. We thus pick units so that all planners' marginal utilities are approximately equal to 1 *a priori*.

$\sum_{j=0}^{N_K-1} \sum_{k=0}^{N_{CO_2}-1} z_{jk}^{i,t} \zeta_j(K) \zeta_k(CO_2)$, where $z_{jk}^{i,t}$ is the set of $N_K N_{CO_2}$ coefficients for type i at time t , and $\zeta_j(\cdot)$ is the Chebyshev polynomial of degree j . The value function is approximated in $[\underline{K}, \overline{K}] \times [\underline{CO_2}, \overline{CO_2}]$ (with the state intervals normalised to $[-1, 1]$ for the Chebyshev approximation itself). We compute the set of Chebyshev nodes in this region, solve the optimisation problem of type i at time t at each node, and choose the coefficients such that the approximation holds exactly at these nodes (Judd, 1998).

We choose $N_K = 24$, $N_{CO_2} = 10$. The problem is set up in Julia, using JuMP and the solver IPOPT. We use tight tolerances and verify that the relative errors in the Bellman equation are small (typically of the order of 10^{-4}), despite the existence of a kink where the abatement rate reaches unity (the maximal degree of abatement). The algorithm is stable as long as the state space region is chosen large enough.

After solving for the value functions, simulation of time paths (given an initial state) is straightforward: given a realisation of the type sequence, the optimal decision can be calculated for each t . Social costs of carbon are calculated using the formula (15), given our approximations for the value functions.

E Alternative specification of the switching probabilities

In this section we present numerical results for an alternative specification of the switching probabilities in (14). In the body of the paper we assumed that, conditional on switch-

ing preferences, which occurs with probability x , each planner has a constant probability $1/(N - 1)$ of switching to each of the alternative normative theories. Here we adopt a different model where planner i 's probability of switching to theory j is inversely related to a measure $d(i, j)$ of the ‘distance’ between their preferences. Given the preference data in Figure 1, we use the well-known Mahalanobis metric to measure this distance. For planner i , define a vector

$$\mathbf{q}_i = \begin{pmatrix} \rho_i \\ \eta_i \end{pmatrix}$$

Let \mathbf{S} be the covariance matrix of the (ρ, η) data in Figure 1. Then we define

$$d(i, j) = \sqrt{(\mathbf{q}_i - \mathbf{q}_j)^T \mathbf{S}^{-1} (\mathbf{q}_i - \mathbf{q}_j)}. \quad (22)$$

In the case where ρ and η are uncorrelated, this distance measure reduces to a standardised Euclidean distance in the (ρ, η) plane, where squared differences between the ρ (η) component of two vectors are normalised by the variance of ρ (η) in the data. When ρ and η are correlated, $d(i, j)$ accounts for the fact that differences in the ρ components of vectors are statistically related to differences in their η components (and vice versa), and adjusts the standardisation of the Euclidean distance measure accordingly.

We then define a new transition matrix \mathbf{w}^C on the set of planner preferences as follows:

$$w_{ij}^C = \begin{cases} 1 - x + A_i x m_i & j = i \\ A_i \frac{x m_j}{1 + \mu d(i, j)} & j \neq i \end{cases} \quad (23)$$

where $\mu > 0$ is a parameter, and

$$A_i = \left(\sum_j \frac{m_j}{1 + \mu d(i, j)} \right)^{-1}$$

is a normalisation constant chosen so that

$$\sum_j w_{ij}^C = 1.$$

Clearly, the specification in (23) reduces to that in (14) when $\mu = 0$. In (23) with $\mu > 0$ however, conditional on switching preferences, the probability of switching from i to j is downweighted by a factor $(1 + \mu d(i, j))^{-1}$, so that preferences that are further from i are switched to less often. For the present analysis we choose the parameter μ so that

$$\min_{(i,j)} (1 + \mu d(i, j))^{-1} = 0.1$$

In words, the i, j pair that is furthest apart receives one tenth of the weight that a switch from one point in preference cluster i to another point in the same cluster receives.

Figure 5 presents the results of SCC calculations for the alternative specification of the switching matrix in (23). This figure tells a remarkably similar story to that in Figure 2, where we used the uniform switching probabilities in (14). The upshot is that at least in this model admitting *some* possibility of a switch in preferences (with probability x) matters more for SCC calculations than precisely *how* preferences switch when they do. One factor that contributes to this is that because most climate damages occur in the distant future,




Figure 5: SCC estimates for the switching probabilities in (23).

both models of the switching process in (14) and (23) give rise to ‘well-mixed’ discount rate distributions at the relevant maturities, and thus the short-run structure of switching probabilities is not so important for explaining SCC calculations. Theorem 1 also indicates a second relevant factor. Regardless of the details of the switching probabilities, long-run present value calculations for non-dogmatic planners are dominated by the lowest value of η (assuming positive expected consumption growth), which is of course independent of the precise nature of the switching process. As Figure 3 indicates, the consumption smoothing motive is a major determinant of social discount rates and SCCs, so it is significant that the long-maturity behaviour of this component is independent of the details of the switching

process.