General Bayesian Learning in Dynamic Stochastic Models: Estimating the Value of Science Policy

Ivan Rudik,[*] Maxwell Rosenthal,[†] and Derek Lemoine[‡]

February 6, 2018

Preliminary

We integrate climate scientists into an economic model of climate change by calibrating a statistical model for updating beliefs about the climate's sensitivity to greenhouse gas emissions to the actual history of scientific progress. We find that nonconjugate priors are critical for representing the observed dynamics of scientific knowledge. In order to investigate the implications for policy, we extend recursive dynamic programming methods to allow for nonconjugate learning about an uncertain parameter. We find that today's policymaker must set emission policy without the expectation that new information will enable timely revisions to policy. Improving scientific monitoring and climate modeling to enable faster learning would be worth up to $XX dollars.

JEL: C61, D83, H23, Q54, Q58

Keywords: uncertainty, learning, nonconjugate, dynamic programming, recursive, climate, carbon

[*]Charles H. Dyson School of Applied Economics and Management, Cornell University; irudik@cornell.edu
[†]Department of Economics, University of Arizona; rosenthm@email.arizona.edu
[‡]Department of Economics, University of Arizona and NBER; dlemoine@email.arizona.edu

The optimal stringency of policies to control greenhouse gas emissions depends on how sensitive the climate is to greenhouse gas emissions, which is is uncertain and is likely to remain uncertain in the future. In fact, we are not much more certain about the climate's sensitivity to emissions than we were forty years ago, despite massive research effort devoted to precisely this question (Knutti et al., 2017). Many scientists even argue that uncertainty about long-term warming has actually increased over this period as a result of new observations and new understanding of climate change in past geological epochs (e.g., Hansen et al., 2008; Kiehl, 2011; Royer, 2016). The next fifty years are unlikely to generate an unprecedented reduction in uncertainty, which has led some scientists to call for research effort to be devoted elsewhere (e.g., Allen and Frame, 2007). Optimal climate policy should be influenced by this slow rate of learning, because it becomes important to act without the expectation of being able to effectively revise policy over the next decades.

Dynamic economic models of climate change have recognized the potential importance of the sensitivity of the climate to emissions, the potential importance of uncertainty about the sensitivity of the climate to emissions, and the potential importance of learning about this sensitivity in future years.[1] Modelers have carefully calibrated current beliefs about the sensitivity of the climate to emissions to current scientific understanding. They have also carefully calibrated the dynamics of the economy and the climate system to data and to scientific understanding. However, modeling decision-making under uncertainty requires tracking informational states as well as physical states. The dynamics of these informational states are potentially as important as the dynamics of physical states, yet belief dynamics have been left entirely uncalibrated. We show that existing climate-economy models predict reductions in uncertainty that are dramatically inconsistent with both historical and projected scientific progress. We take up the task of calibrating learning in a dynamic economic model of climate change.

We here demonstrate how to calibrate the dynamics of Bayesian beliefs, we develop a computational approach to implementing realistic belief dynamics in a recursive dynamic programming framework, and we apply our new calibrational and computational approach to a prominent question in climate change policy. Standard economic approaches to calibration can be roughly divided into three types: a first type of calibration matches the levels of variables to data (e.g., matching equilibrium employment, savings, or the equity premium); a second type matches the dynamics of variables to data (e.g., targeting a balanced growth path or changes in the equity premium over the business cycle); and a third type matches the distribution of variables to the distribution of data (e.g., dynamic stochastic general equilibrium models might match moments of hours worked and trade models might target the distribution of productivity). We instead match the endogenously generated dynamics of beliefs to the actual record of beliefs. In fact, our learning model is nearly identical to how climate scientists empirically estimate the climate's sensitivity to emissions (e.g., Aldrin

[1]See Lemoine and Rudik (2017) for a review.

et al., 2012; Ring et al., 2012; Schwartz, 2012; Otto et al., 2013; Lewis, 2013; Skeie et al., 2014). We therefore integrate climate scientists into an economic model of climate change.

Our climate change setting is particularly amenable to the calibration of belief dynamics because we have a published record of consensus beliefs over many decades as well as the primary data (observations of carbon dioxide and temperature) that inform these beliefs. This type of calibrational approach could be applicable in other settings with a record of beliefs as implied by, for instance, option prices or credit default swaps.[2] However, our climate change setting poses especially thorny computational challenges when it comes to implementing our calibration. Our learning model requires historical data on three different climate states over several centuries, and tracking these history states expands our state space well beyond the computational limits of previous models.

Our calibrated learning model is so computationally demanding because we had to abandon a standard restriction on beliefs: that the prior distribution must be conjugate to the likelihood function. This restriction is universally imposed in dynamic policymaking environments for the sake of tractability.[3] It ensures that the posterior distribution comes from the same parametric family as the prior distribution. In that case, learning can be implemented by extending the state vector to include the sufficient statistics for the prior distribution, with their transition equations derived from Bayes' rule. For example, a normal prior and normally distributed data combine to ensure that the posterior is also normal, in which case Bayesian learning can be tractably implemented by extending the time t state vector to include the mean and variance of the time t prior distribution. We find that the assumption of conjugacy overly restricts belief dynamics in our climate change application. Conjugate models that implement a reasonable prior from 40 years ago and then update it with subsequent climate data predict that we should now be fairly certain about the climate's sensitivity to emissions. In particular, these models predict that we can now rule out a highly sensitive climate. This prediction is clearly counterfactual. We find that the learning model can match the historical record if we use a lognormal prior, but a lognormal prior implies a nonconjugate model of learning. We expect that nonconjugate learning models will also better describe many other environments of economic interest. For instance, recent work

[2]Recent work has focused on the implications of learning for asset prices but has not calibrated the learning process (Collin-Dufresne et al., 2016). Pastor and Veronesi (2009a) survey work in finance that considers learning. For example, Brennan (1998) studies the allocation of investment to a risky asset under learning about the equity premium, and Brennan and Xia (2001) study the equity premium that emerges from equilibrium investment decisions when agents learn about the expected dividend growth rate.

[3]For instance, Cogley and Sargent (2008) describe how they limit attention to distributions that can be summarized by a finite-dimensional vector of summary statistics and for which learning occurs via conjugate priors. We here relax both standard restrictions. In fact, much work in macroeconomics has not only relied on conjugate priors but has further eased tractability by using "anticipated utility" (see Kreps, 1998). In these frameworks, the decision maker updates beliefs from period to period as a Bayesian would but formulates policies as if current beliefs will never change. The decision maker learns only passively, without considering the informational value of his potential actions. Our decision makers will always formulate policies in full awareness of the possibility of future learning and of how their policy choices affect the rate of learning.

in macrofinance has focused on the possibility of rare disasters (e.g., Rietz, 1988; Barro, 2006). Models of nonconjugate learning could allow for more flexible implementations of rare disasters in an environment with anticipated learning.[4]

We extend frontier computational methods for dynamic economic modeling in order to integrate Bayesian nonconjugate learning into a dynamic decision problem. The challenge is that the posterior is not typically a named, parametric distribution under nonconjugate learning. We therefore cannot summarize the posterior by storing parameters such as the mean and variance and relying on a known density function to assign probability weights. Instead, we must track a history of observations so that we do not lose information by fitting a parametric distribution to the posterior at each timestep. But the complexity of dynamic models grows in the dimension of the state space (a specific form of this effect is known as the "Curse of Dimensionality"), and the dimension of the state space becomes enormous once we begin tracking a history of observations. Further, nonconjugate learning must be implemented by combining the data with an algorithm to approximate the posterior (such as Markov chain Monte Carlo), which can substantially increase the computational time required to update beliefs.[5]

We combine and extend two recent developments in computational economics to render nonconjugate learning tractable.[6] First, the standard practice in computational economic modeling has been to approximate the unknown value function by using a tensor product of nodes in the state space (e.g., Judd, 1998; Miranda and Fackler, 2002). The size of this tensor product grows exponentially with the dimension of the state space. Recently, several papers have used a sparse grid technique pioneered by Smolyak (1963) to approximate value functions and policy functions (e.g. Krueger and Kubler, 2004; Gonzalez and Rojas,

[4]All of the work described in footnote 2 uses conjugate priors, and nearly all of it uses a normal prior in particular. Johnson (2007) explores learning about a new technology's returns-to-scale as a means of generating a bubble in investment, and Pástor and Veronesi (2009b) study how learning about the productivity gain from a new technology can generate a stock price bubble. Both of these papers use a conjugate, normally distributed prior, even though Johnson (2007) explicitly notes that the conjugate prior is not a good fit to his setting because its unbounded support allows the production function to be explosive. In many ways, our work is closer to Veronesi (2004) and Orlik and Veldkamp (2014), where learning can generate fluctuations in uncertainty about fundamental states.

[5]We require Markov chain Monte Carlo techniques instead of simply evaluating the posterior at a set of predetermined (Gauss-Legendre) quadrature nodes because, as we will show later, we must compute an expected posterior over a number of posteriors we recover from an interpolation procedure. Averaging over the density values of the posteriors at the quadrature nodes will not necessarily be equivalent to averaging the entire densities.

[6]To our knowledge, the only other implementations of nonconjugate Bayesian learning in a model of dynamic decision making project each period's posterior into a parametric distribution in the same family as the prior by minimizing the distance between the nonconjugate posterior and the target distribution (Zhou et al., 2010; Springborn and Sanchirico, 2013; Kling et al., 2017; MacLachlan et al., 2017). This "density projection" method avoids the need to track the history of observations. Our approach differs in remaining fully nonparametric: we do not rely on earlier timesteps' approximated distributions when modeling learning at later timesteps.

2009; Winschel and Krätzig, 2010; Rudik, 2016).[7] The size of the Smolyak grid grows only polynomially in the dimension of the state space, which allows us to track the history of observations at regular intervals. But even such a grid would be intractably large if it tracked the history at the annual interval consistent with actual data. We further reduce the computational burden by developing methods that allow for tracking the history only at larger intervals. Specifically, we use a Brownian bridge to approximate learning over the intervals between each tracked observation. The combination of larger intervals and a Smolyak grid helps to keep our problem to a tractable size.

But we still must solve the problem of having to run a costly nonconjugate learning algorithm many times. This is where we take advantage of a second recent development in computational economics: using precomputation to reduce the time cost of the solution. Precomputation recognizes that some costly calculations depend only on the location in the state space rather than on policy choices. In these cases, costly calculations can be done one time "offline" before running the optimization model, with only the results of the calculations called over and over in the optimization routine. Precomputation has been applied to taking expectations (Judd et al., 2017) and computing labor decisions (Maliar and Maliar, 2005). We here extend precomputation to the computation of nonparametric beliefs. This extension allows us to implement nonconjugate learning offline, rather than calling the costly learning algorithm over and over throughout the optimization routine. In our application the benefits of precomputation are immense. If we are to solve several alternative models that maintain the same climate system (e.g. changing consumption preferences, the discount rate, or the function that quantifies the damages from climate change), having a precomputed set of beliefs conservatively reduces computation time by three orders of magnitude.

We find that implementing nonconjugate learning has strong implications for greenhouse gas emission policies. Because our learning dynamics are consistent with the history of real world beliefs, our model projects much less future learning than do other economic models. In cases where the climate is particularly sensitive to emissions, the degree of uncertainty is essentially unchanged until the end of the twenty first century. Today's policymakers must therefore act without the expectation of being able to revise policy if they happen to be in a bad state of the world. We find that uncertainty tends to lead optimal policy to allow more emissions than the perfect-information policy would allow. This effect arises because near-term warming is concave in the sensitivity of the equilibrium climate to emissions, even though long-term warming is linear in this sensitivity.[8] A policymaker who discounts the future at a sufficiently high rate will therefore find that greater uncertainty can reduce expected temperature. Contrary to previous work in the economics of climate change (described in

[7] The Smolyak grid can be seen as an optimal grid in the sense that it minimizes approximation error measured with L_2 or L_∞ norms given a fixed number of collocation nodes (Brumm and Scheidegger, 2016).

[8] We show that in the DICE-2007 climate-economy model, ocean heat sinks respond more strongly when climate sensitivity is greater, which makes the marginal effect of today's emissions on temperature concave in climate sensitivity. This is effect is consistent with results in climate science (e.g., Baker and Roe, 2009).

the next section), we find that the wedge between optimal policy and the perfect-information policy grows over the course of the century. Even if the climate's sensitivity to emissions is consistent with initial expectations, end-of-century carbon dioxide (temperature) is 40 ppm (0.2°C) greater under the optimal policy than under the perfect-information policy. These wedges become much larger if the climate is in fact more sensitive to emissions than currently expected.

Economists have focused on evaluating policies to control emissions and to direct R&D towards low-emission technologies, but the most thorough, longstanding U.S. policies around climate change support scientific monitoring and modeling. The federal government funds satellites that measure carbon dioxide fluxes and temperature, networks of ocean sensors that measure the distribution of heat stored in the ocean, and large-scale modeling enterprises that draw on vast computing power to project future changes in the earth system. Because we generate realistic belief dynamics, we can assess the value of improved scientific monitoring and modeling. Better scientific monitoring and modeling acts like reducing the noise in the year-to-year temperature transition, which allows the policymaker to more confidently back out the climate sensitivity to emissions from a sequence of temperature observations. Because our setting predicts only slow learning, we find that scientific information about climate change is quite valuable. Nonconjugate learning alone increases balanced growth equivalent consumption by only 0.3% relative to never updating the initial prior belief. However, balanced growth equivalent consumption would increase by up to 8.5% if we were to rapidly increase scientific information and thereby resolve climate sensitivity uncertainty very soon. This high value of scientific information argues for increased funding for climate modeling efforts and for atmospheric and oceanic monitoring systems that could reduce the unexplained variance in year-to-year temperature changes and hence allow for faster learning about the climate's sensitivity to emissions.

The next section demonstrates that the standard, conjugate approach to learning about long-run warming is inconsistent with the observed rate of scientific progress. A nonconjugate model of learning performs much better. Section 2 details our new computational approach to implementing nonconjugate learning into a dynamic economic model of policymaking. Section 3 calibrates the belief dynamics implied by our learning model. Section 4 briefly describes the DICE-2007 climate-economy model into which we integrate our nonconjugate model of learning. Section 5 presents results. The final section concludes. The appendix contains additional methodological details and results.

1 The Binding Constraint of Conjugate Learning in Climate Economics

Dynamic economic models have generally been restricted to studying learning in settings where the prior distribution is "conjugate" to the likelihood function. In this special case, the posterior is known to have the same type of distribution as the prior, so that the same types of parameters that identify the prior distribution also identify the posterior distribution. These informational parameters become state variables in the dynamic economic model, with transition equations from time t to $t+1$ derived from combining Bayes' rule with time t observations. We here demonstrate that the standard constraint of conjugacy in fact binds in climate change applications: imposing conjugacy prevents the modeler from capturing realistic belief dynamics.

We analyze the case of learning about the sensitivity of the climate to greenhouse gas emissions. First, consider the real world dynamics of learning, which are the dynamics that an economic model of climate change should match. A key metric for the strength of climate change is "climate sensitivity," or the equilibrium response of global mean surface temperature to doubling the concentration of carbon dioxide (CO_2) from its preindustrial level.[9] In a benchmark early report, Charney et al. (1979) estimated climate sensitivity "to be near 3°C with a probable error of ±1.5°C."[10] Since then, computational climate models have become vastly more complex and realistic, satellite data have become available, paleoclimatologists have learned more about past swings in the planet's climate, and we have obtained several decades of new temperature observations in a regime with a clearly changing climate. One might therefore expect that the variance of our distribution for climate sensitivity has shrunk. This is not the case. The Intergovernmental Panel on Climate Change (IPCC) periodically summarizes the state of scientific knowledge about anthropogenic climate change. In their First Assessment Report, the IPCC estimated that climate sensitivity "lies between +1.5°C and +4.5°C" (Mitchell et al., 1990). In their recent Fifth Assessment Report, the IPCC judged that climate sensitivity "is likely in the range 1.5°C to 4.5°C", "extremely unlikely" to be less than 1°C, and "very unlikely" to be greater than 6°C (Collins et al., 2013). Uncertainty since 1979 has endured despite climate sensitivity being the major metric of interest in climate science, with substantial scientific effort devoted to estimating its value (Knutti et al., 2017). This lack of scientific progress has led to theoretical arguments that the nonlinear nature of the climate system may inherently constrain our ability to learn about climate sensitivity (e.g., Roe and Baker, 2007) and has induced calls to redirect scientific research effort towards tasks that are more promising than seeking constraints on climate sensitivity (e.g., Allen and Frame, 2007).

Not only did the likely range of climate sensitivity remain constant from 1979 to 2013, but many would argue that uncertainty about climate sensitivity actually increased over

[9] "Equilibrium" here refers to the steady state, so climate sensitivity measures the change from one steady state to another. The traditional definition allows the ocean and the atmosphere to come to equilibrium but prevents some types of processes from responding to temperature. For instance, the carbon cycle will respond to temperature in reality and thereby change the concentration of CO_2 and also temperature, but the traditional definition of climate sensitivity abstracts from this process. For overviews of climate sensitivity, see Andronova et al. (2007), Knutti and Hegerl (2008), and Knutti et al. (2017), among others.

[10] About 80 years earlier, Arrhenius (1896) estimated climate sensitivity to be in the range of 5°C–6°C. Ten years later, Arrhenius (1906) honed in even closer to modern beliefs with a best estimate of 4°C.

this interval. First, many scientists have recently argued that paleoclimatic data and recent observations of earth system elements such as land ice sheets suggest that values much greater than 4.5°C are not so unlikely after all (e.g., Hansen et al., 2008; Lunt et al., 2010; Pagani et al., 2010; Kiehl, 2011; Royer, 2016). Second, the IPCC's 2007 Fourth Assessment Report actually offered a slightly different range for climate sensitivity, judging it "likely to lie in the range 2°C to 4.5°C" and to be "very likely larger than 1.5°C" (Meehl et al., 2007). However, new empirical research since that report, much like what climate scientists in our model perform, made the possibility of a climate sensitivity below 2°C seem more plausible, so that the lower bound of the "likely" interval dropped back to 1.5°C in the 2013 Fifth Assessment Report.

We propose that a reasonable economic model of learning about climate change should be consistent with the history of scientific progress.[11] If models of decision-making under uncertainty cannot hindcast how beliefs responded to past information, then we may have little confidence in their ability to forecast how beliefs will respond to new information. In particular, a reasonable learning model should not demonstrate substantial changes from the 1979 prior to the 2013 posterior when simulated with data from the intervening years, and a reasonable model should allow for the possibility that new observations increase the variance of beliefs.

Previous economic models of learning about climate change meet neither requirement. First, previous models learn too fast when simulated over the past decades. Figure 1 plots the consequences of simulating the learning models of Kelly and Kolstad (1999), Kelly and Tan (2015), and Hwang et al. (2017) with annual temperature observations from 1979 to 2013. Next to each set of authors is the mean and standard deviation of the model's posterior distribution. All three models start with a prior that is similar to the type of prior found throughout the climate science literature and often taken to be consistent with Charney et al. (1979) (light blue).[12] All three models suggest that we should have learned a substantial amount about climate sensitivity by now: the standard deviation in each model is considerably smaller than the standard deviation of the 1979 prior and thus also smaller than the standard deviation of what today's scientific literature takes to also be a reasonable

[11] Our proposal extends recent calls for the next generation of integrated assessment models to have climate modules that can accurately hindcast the climate's history (National Academies of Sciences, 2017). It also extends recent work testing the ability of integrated assessment models' economic modules to hindcast economic growth over the past century (Millner and McDermott, 2016). We here emphasize that hindcasting experiments should focus on belief dynamics as well as the dynamics of the physical climate and the economy.

[12] Kelly and Kolstad (1999) does not fit the climate science literature's prior as well as do the others because they constrain climate sensitivity to be normally distributed. Among published papers, we omit only Leach (2007), which extends Kelly and Kolstad (1999) to allow for uncertain serial correlation that slows learning a bit, and Lemoine and Rudik (2017), which learns decadally instead of annually and thus is not directly comparable to the others. In unpublished work, Jensen and Traeger (2016) also consider learning about climate sensitivity, and we find that their model implies learning speeds comparable to the published models plotted in Figure 1. Lemoine and Rudik (2017) review this literature's methods and results.

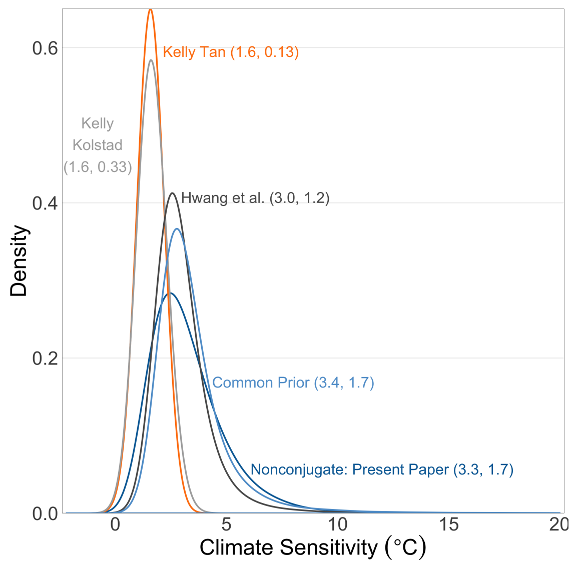

Figure 1: A common representation of the year 1979 prior on climate sensitivity (light blue), and the year 2013 posterior implied by combining subsequent annual observations of surface and ocean temperature with the present paper's learning model (dark blue) as well as the learning models in Kelly and Kolstad (1999) (light gray), Kelly and Tan (2015) (orange), and Hwang et al. (2017) (dark gray). Beside each set of authors is the mean and standard deviation.

posterior. Whereas some climate scientists have called for abandoning the futile search for constraints on climate sensitivity, these learning models instead judge that we should already know that a climate sensitivity above 5°C (or even 4°C) is very unlikely.

Previous models' bias towards overly rapid learning has dramatic consequences for their projections of future knowledge. For instance, when the true climate sensitivity is near 3°C in model simulations, Kelly and Tan (2015) find that the policymaker will effectively rule out values much above 4°C within only 10 more years. Learning is slower in Hwang et al. (2017), but the policymaker can still expect to effectively rule out values much above 4°C within 100 years and the posterior's coefficient of variation is cut by half in 50 years. The learning models in Kelly and Tan (2015) and Hwang et al. (2017) imply that learning will be dramatically faster over the next 40 years than it was over the last 40 years and imply that scientists' fears about being inherently unable to rule out high values of climate sensitivity are misplaced. These types of results work to limit the importance of uncertainty about climate sensitivity for policy.

Second, all three of these models a priori eliminate the possibility that new observations will increase the variance of beliefs. All three models use a "normal-normal" learning model, in which both the prior and the posterior are normally distributed and the likelihood function arises from a normally distributed error. This is one of the most common conjugate prior models used throughout economics, in part because normally distributed data generating processes arise naturally when error terms are additive and normally distributed. In the present case, the error term is an additive weather shock in the temperature transition equation, and it is reasonable to assume that such a shock would be normally distributed. Yet not only do these "normal-normal" models learn counterfactually fast, but they also severely restrict even the types of learning that are feasible. In a normal-normal learning model, the variance of the prior evolves as the fraction of the data's subjective variance that is due to the data rather than the prior. The variance of the posterior therefore declines monotonically, and the monotonicity of the variance's evolution is independent of the observed data: what really matters for the variance of one's posterior are the number of observations accumulated and the variance of the error term in the data generating process. It seems reasonable that observations that are very inconsistent with one's current beliefs should make one less certain about climate sensitivity, but the conjugate learning transitions a priori rule out this effect. This restriction is one reason why all of the previous models learn so fast in Figure 1. And this restriction is contrary to many scientists' fears about the future. Indeed, Hannart et al. (2013) argue that it is very possible that uncertainty about climate sensitivity could increase in the future precisely because the distribution of climate sensitivity is non-normal.[13]

We have found that relaxing the restriction to conjugate learning is critical for replicating the dynamics of scientific knowledge. Section A.1 provides additional evidence that nonconjugacy is the key difference between our learning dynamics and previous, conjugate learning dynamics, and it also shows how conjugacy implies counterfactual climate sensitivity beliefs going back to 1880 that are far outside the range of actual documented climate sensitivity beliefs. As we will describe in Section 3, we develop a setting that learns from the post-1880 observational record of temperature and CO_2 and endogenously matches the present state of scientific knowledge. Figure 1 shows that our learning model (in dark blue) does not suggest that the data since 1979 should have strongly constrained beliefs about climate sensitivity. In fact, largely consistent with scientific progress over the past few decades, the variance of our posterior distribution actually *increased* over this timeframe while the mean was effectively unchanged.

Nonconjugate learning is critical to our calibration. However, nonconjugate learning has not been used in dynamic economic models because it is thought to be intractable. We next describe how we make nonconjugate learning tractable and thereby allow for an assessment of climate change policy under a realistic model of belief dynamics.

[13] In Kelly and Tan (2015) and Hwang et al. (2017), the prior is over a "feedback factor" that determines climate sensitivity rather than over climate sensitivity directly, which allows them to mimic the non-normal shape of conventional distributions for climate sensitivity (see Roe and Baker, 2007; Roe, 2009; Lemoine, 2010). Observations that shift the mean towards a higher feedback factor might increase the variance of the distribution for climate sensitivity even as they reduce the variance of the distribution for the feedback factor. However, the variance of the distribution for the feedback factor falls so quickly in these two models that this potential effect is of little practical relevance.

2 Dynamic Programming with Nonconjugate Learning

We have seen that dynamic economic models' standard restriction to conjugate priors has prevented economic models of climate change from reflecting the actual dynamics of scientific knowledge. A model of nonconjugate learning performs much better. However, the restriction to conjugate priors is used in large part because nonconjugate learning is thought to be intractable when combined with a dynamic economic model of decision-making.

We here develop and implement a method for integrating nonconjugate learning with dynamic programming. In dynamic programming, the computational difficulty of the problem grows (often exponentially) with the dimensionality of the state space, a problem commonly referred to as the "Curse of Dimensionality." When using conjugate priors, modeling learning requires expanding the state space to include informational states that define the prior distribution (e.g., the mean and variance when the prior is a normal distribution). However, when learning is nonconjugate, the posterior is not guaranteed to be in the same distributional family as the prior, so we do not generally have sufficient statistics for the posterior. Modeling nonconjugate learning requires expanding the state space to keep track of the entire history of observations. This dramatic expansion of the state space requires frontier computational methods.

Consider solving the following Bellman equation:

$$
\begin{aligned}
V_t(S_t, H_t) = \max_{z_t} \Big\{ & u_t(z_t, S_t) \\
& + \beta \int \int V_{t+1}(S_{t+1}(z_t, S_t; \theta, \epsilon_{t+1}), H_{t+1}(S_t, H_t))\, p(\theta | S_t, H_t)\, p(\epsilon_{t+1})\, \mathrm{d}\theta\, \mathrm{d}\epsilon_{t+1} \Big\}.
\end{aligned}
$$

Here z_t is a vector of controls, $u_t(\cdot, \cdot)$ is the per-period payoff, and ϵ_{t+1} is a shock with a known distribution. S_t is a vector of payoff-relevant states, and H_t is a vector that contains the entire history of states that provide information about the unknown parameter θ. This history is used to develop the posterior $p(\theta|S_t, H_t)$, which determines the expectation of the continuation value $V_{t+1}(\cdot, \cdot)$. A standard way of solving the Bellman equation (and one of the only methods that works in the context of climate-economy integrated assessment models) is to use value function iteration: approximate the value function V_{t+1} over the state space (S_{t+1}, H_{t+1}), solve the time t Bellman equation on a grid of collocation nodes in the state space, use the results to approximate the value function V_t over (S_t, H_t), and iterate until the approximant converges (in an infinite horizon setting) or until we reach time 0 (in a finite horizon setting).[14]

Nonconjugate learning introduces two problems. First, the history vector H_t expands as t grows, which increases both the dimensionality of the state space and the number of colloca-

[14]Benchmark climate-economy integrated assessment models lack a steady state, have more states than controls (and so cannot be solved by Euler equation methods), and have kinked policy functions. Value function iteration has therefore been the primary solution technique for recursive implementations of these models. See Lemoine and Rudik (2017) for further discussion of conventional computational approaches.

tion nodes where we must maximize the Bellman equation. To address this problem, we construct our collocation grid using Smolyak's sparse grid method (Smolyak, 1963; Judd et al., 2014) rather than using the more conventional full tensor product grid. Smolyak's method uses a subset of the nodes generated from standard tensor product approaches while maintaining favorable properties for approximation accuracy (Barthelmann et al., 2000; Brumm and Scheidegger, 2016). The number of nodes on a Smolyak grid grows only polynomially, rather than exponentially, in the dimensionality of the state space. However, as t becomes very large (climate-economy models use a horizon that extends for centuries), the lengthy histories pose problems even for a Smolyak grid. We therefore also develop a method that requires tracking only a subset of the history vector. We use Brownian bridges to interpolate the omitted portion of the history vector while maintaining the stochasticity of temperature data. Our policymaker therefore learns from the full set of data even as our history vector tracks only a subset of the data.[15]

The second problem posed by nonconjugate learning concerns the computation of the posterior $p(\theta|S_t, H_t)$. This posterior must be computed at each maximization step, which means that it must be computed at each of the collocation nodes. With conjugate priors, the posterior is governed by closed-form transition equations for the parameters defining the prior distribution. Computing the posterior looks like computing the future value of any standard state variable, and the conventional assumption of normally distributed variables lends itself to efficient Gauss-Hermite quadrature rules for integration over the continuation value. In contrast, when priors are nonconjugate, the posterior is typically constructed by sampling from its distribution through techniques like Markov chain Monte Carlo. We here take advantage of the fact that the posterior depends only on the time t state variables, not on the time t controls. We construct the posterior "offline" for every node combination prior to solving the Bellman equation even once, and each maximization step then draws on the precomputed information about the posterior generated by its node in (S_t, H_t) space.

We now step through the algorithm, as outlined for our climate-economy application in Figure 2. The appendix provides further details and pseudocode.

2.1 Collocation

The solution algorithm begins by constructing a collocation grid in the model's state space, which we will use to approximate the policymaker's value function (e.g., Judd, 1998; Miranda and Fackler, 2002). In our climate-economy application, the state space is comprised of the states contained in DICE-2007 (Nordhaus, 2008) as well as the histories of states that provide information about uncertain parameters. The standard states in DICE are capital, atmospheric CO_2, upper ocean CO_2, lower ocean CO_2, surface temperature, and lower ocean

[15] Even recent advances in adaptive sparse grids, such as in Brumm and Scheidegger (2016), would be overwhelmed by the size of the state space. A 300-year problem would result in a state space with nearly 1000 dimensions.

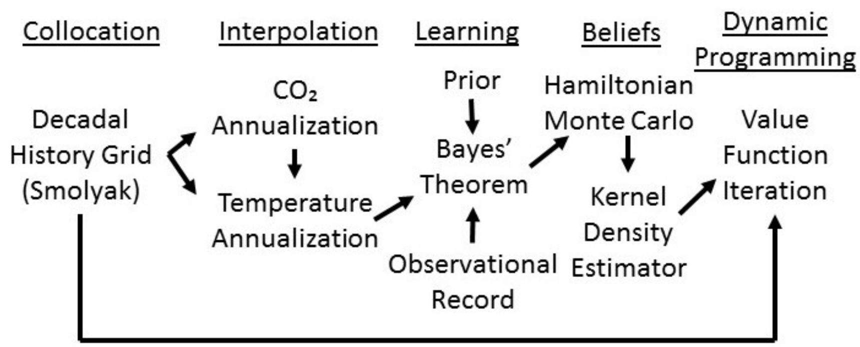

Figure 2: Schematic of the method for incorporating nonconjugate learning.

temperature. In order to undertake nonconjugate learning, we also include vectors of historical observations of atmospheric CO_2, surface temperature, and lower ocean temperature. We measure time in years, with the model's start year defined as time 0. This state space would have dimension $3t + 6$.[16] However, as discussed in the "interpolation" step below, we track the state variables only at every tenth timestep (i.e., only every decade). Our state space will therefore actually have dimension $3(t/10) + 6$. We construct our grid using a more efficient version of Smolyak's method (Judd et al., 2014), as discussed above. Nonetheless, our grid will still require on the order of 10^5 nodes for some decades of the problem.[17]

2.2 Interpolation

The interpolation step explains how we use annual data in our learning calculations despite storing only 10-year timesteps in our history vectors. This step is critical, because storing annual timesteps is not feasible when t becomes large, but using lower resolution histories to compute the posterior would underestimate the amount of data actually available. Consider a given collocation node in the state space (S_t, H_t), and define S_t^H as the vector of contemporaneous states (e.g., surface temperature, lower ocean temperature, and atmospheric CO_2) that will enter the history in future periods ($S_t^H \subseteq S_t$). The history vector H_t contains state vectors $S_{t-10}^H, S_{t-20}^H, \ldots, S_0^H$.[18] The learning algorithm needs to know what happened between S_{t-20} and S_{t-10}: if we do not give it that data, it will act is if the intra-decadal

[16] As discussed in Section 3, our learning model uses observational data from the past century in addition to histories generated after model time 0. However, these observational data are fixed from the perspective of our policymaker at all times $t = 0, 1, 2, ...$, so we do not need to treat them as states.

[17] As we go further in the future, we need to store longer histories and therefore have more states and collocation nodes.

[18] Note that the combination of state values defined by the collocation node can be entirely nonphysical, in that the combination of values generated by our collocation grid may never actually arise under the model's dynamics. For instance, the history vector defined by a given node may tell us that CO_2 increased monotonically while temperature decreased monotonically over some time period, even though the model's

values were never observed and underestimate the quantity of information actually available at t.

We develop these intradecadal data by constructing a Brownian bridge from, for instance, S_{t-20} to S_{t-10} with a drift function determined by the model's transition equations. This bridge defines a normal distribution for temperature at each time $s \in \{t-20, ..., t-10\}$, conditioned on the drift function and the need to start at S_{t-20} and end at S_{t-10}. Each Brownian bridge produces noise which, although necessary to interpolate stochastically evolving temperature between decades, we do not want propagating into our value function. As discussed below, we marginalize out the noise from the Brownian bridge by taking N =1,000 independent and identically distributed samples of the entire Brownian bridge. We then provide this random sample of climate histories to the learning algorithm, compute a posterior for each of the 1,000 unique histories, and average over the resulting posteriors. We thereby obtain the expected posterior distribution conditional on observing the thinned history stored in H_t.[19]

In our particular climate change application, learning will depend on the history of CO_2 and temperature. CO_2 evolves slowly and deterministically. We therefore linearly and deterministically interpolate CO_2 within a decade. The left panel of Figure 3 displays this interpolation for the optimal CO_2 trajectory for the deterministic DICE-2007 model. The linear interpolation fits the trajectory well, without noticeable kinks.

Interpolating surface and lower ocean temperatures is the trickier, and more critical, step. This is where we rely on the Brownian bridge. We define the drift of each of these two stochastic processes by using downscaled versions of the DICE model's physically-based transition equations (see appendix). The Brownian bridge ensures that the simulated temperature 10 years from, for instance, a given draw of S_{t-20} matches the temperature reported in S_{t-10}, not the temperature implied by the model's deterministic physics. The volatility of the Brownian bridge is directly estimated in our downscaled DICE transition equations with data from the historical climate record, as described in Section 3. The right panel of Figure 3 displays a Brownian bridge interpolation of surface temperature along the DICE-2007 optimal trajectory. The annually interpolated points pass through the decadal end points even as they evolve stochastically within each decade.

(deterministic) physics would not allow such a history. One could imagine further trimming the history vector by imposing restrictions based on model dynamics. This technique would bear some relation to adaptive grid methods (see Maliar and Maliar, 2015, 2014). We leave this to future work.

[19]The precise posterior obtained from a full history could differ from this calculated posterior depending on the precise values of the missing history states. One should therefore beware of trimming the history vector too much, because the variance of the interpolated states and thus the variance of the realized posterior would increase in the number of states that need to be interpolated between the tracked states. Also, note that our method is not equivalent to using expected values in the interpolation (i.e., to fixing the Brownian bridge's shocks to zero). That approach would overestimate the rate of learning because the policymaker would always see the least noisy time series that could actually occur.

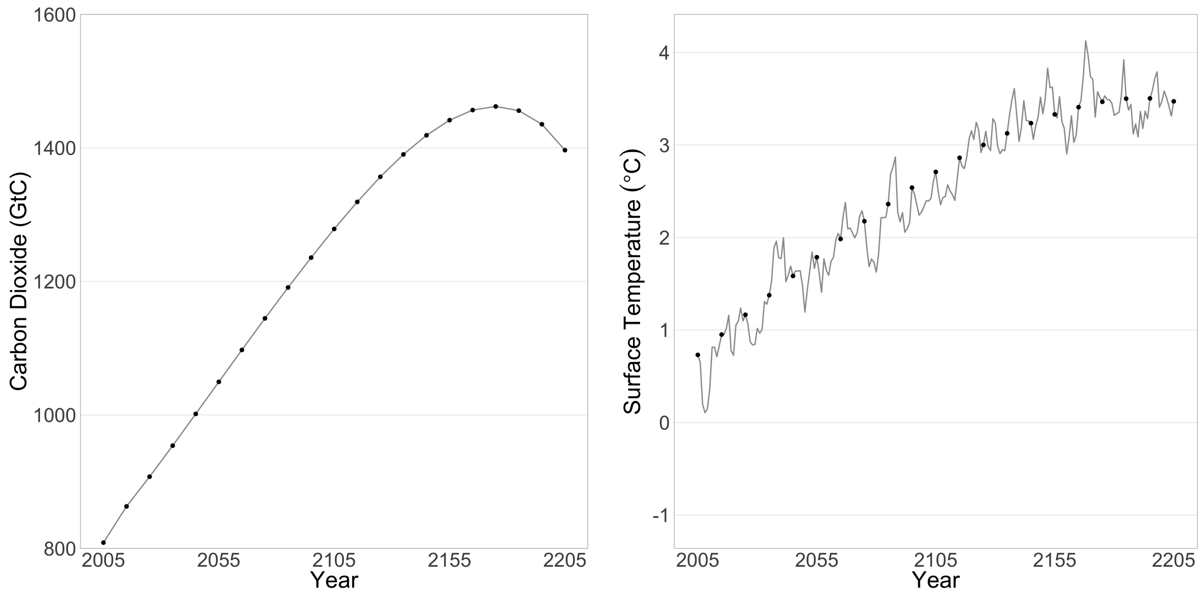

Figure 3: The annually interpolated CO_2 trajectory (left) and one sample of an annually interpolated surface temperature trajectory (right) along the DICE-2007 optimal trajectory. The decadal DICE-2007 trajectory is depicted with the black circles.

2.3 Learning

Our policymaker is a Bayesian. She formulates the posterior distribution for her unknown parameter θ by combining her prior over θ with the current observations in S_t and the history of observations in H_t using Bayes' rule:

$$p(\theta|S_t, H_t) = \frac{p(S_t, H_t|\theta)\, p(\theta)}{\int p(S_t, H_t|\theta) p(\theta) \mathrm{d}\theta},$$

where $p(\theta)$ is her time 0 prior about θ. In our climate change application, she learns about the sensitivity of the planet's climate to CO_2. At each collocation node (S_t, H_t), she updates her beliefs by combining her prior over this sensitivity with climate data from the twentieth century (see Section 3), with the data stored in the history vectors, and with the most recent observations of CO_2 and temperature.

2.4 Beliefs

Dynamic programming applications typically restrict attention to conjugate priors, in which case the posterior can be derived analytically from Bayes' rule. We consider nonconjugate priors. As detailed in the appendix, we use Hamiltonian Monte Carlo (HMC), a method closely related to Markov chain Monte Carlo, to obtain a sample from the N posterior distributions implied by the N different Brownian bridge interpolations for each decadal history. This method is too computationally costly to be implemented within a value function iteration routine.

Instead, we implement the HMC routine "offline," before attempting to solve the dynamic programming problem. We are able to "precompute" the posterior at every combination of states because this posterior depends only on the combination of states reached, not on the policies that will be adopted within the dynamic programming calculations. In our climate change application, we (stochastically) interpolate the history vectors at each of the approximately 100,000 nodes on our grid 1,000 times each using the methods outlined in Section 2.2. We then run the HMC algorithm to generate 50,000 samples for each of the 1,000 histories at each node and we then combine the samples to approximate the nonparametric posterior used for decision-making at the node.[20] In total, we use 50,000 $\times$ 1,000 = 50,000,000 samples to approximate the potential posteriors at each node. We distribute this precomputation across 1,400–4,800 cores at the University of Arizona's high performance computing facility.[21]

In order to utilize the posterior sample when solving our Bellman equation, we use a kernel density estimator to nonparametrically estimate the average posterior density of the 1,000 Brownian bridges.[22] The averaging of the posteriors effectively marginalizes out the bridge's shocks. The appendix describes our Markov chain sampling method and kernel density estimation in more detail.

Why does precomputation matter? Computing one set of beliefs takes approximately 10 days or 300,000–1,000,000 cores-hours depending on how many cores we are able to use. If we wish to solve different kinds of models, as long as we do not vary the structure of the model that directly generates the learning data we only need to perform the learning component of the algorithm once. In our setting this allows us to vary any part of the model that is not directly related to the temperature transitions. Being able to precompute beliefs matters because the rest of the algorithm is computationally cheap: solving the value function only takes 300–400 core-hours. Precomputing beliefs reduces the cost of future alternative model runs by 99.9%.

[20] The number of nodes declines as we step backwards in time through the value function iteration algorithm because we will have a shorter history and thus fewer states as we get closer to the initial time.

[21] How many processors we are able to use depends on requested usage by other researchers.

[22] To average the posterior we combine the posterior samples from the 1,000 different Brownian bridges and then we apply the kernel density estimator with a Gaussian kernel. The bandwidth we use is the one that minimizes the mean squared error of the posterior sample mean and the posterior mean implied by using the quadrature rule. We opt for a Gaussian kernel instead of an Epanechnikov kernel since it provides lower mean squared error in our application.

2.5 Dynamic programming

The previous steps result in a precomputed posterior distribution at every collocation node in the state space. The final step is to solve the Bellman equation. We use value function iteration (Judd, 1992, 1998; Miranda and Fackler, 2002), where we approximate the value function via collocation on a sparse grid with the Smolyak basis (Judd et al., 2014). This step is standard, albeit with a posterior density defined by the kernel density estimator described above. We begin at the terminal time T and solve the Bellman equation at every node. We fit a time T value function approximant to the results. We then step back to the previous time period and repeat the procedure, using the time T value function approximant to define the continuation value.[23] We proceed until reaching time 0. The appendix provides the full model equations and reports computational details. At the conclusion of this procedure, we have a value function approximant at every date, and we can use that approximant to solve for optimal policy and to simulate optimally controlled trajectories.

3 Calibrating the Learning Model

We now describe the calibration of our learning model, which we will integrate into the benchmark DICE-2007 climate-economy integrated assessment model (Nordhaus, 2008). This climate-economy model uses a two-box climate system, in which surface temperature and ocean temperature interact over time. Our learning model combines the transition equations for surface and ocean temperature from an annualized version of DICE-2007 with observations of surface temperature, ocean temperature, and "forcing" (a measure of the heat directly trapped by greenhouse gases like CO_2) since 1880, as plotted in Figure 4.[24] We put the temperature record on the same scale as the Brownian bridge interpolation in Figure 3 for comparison. The appendix provides the formal description. We here note salient details.

First, the ocean temperature series begins in 1955, not 1880. Rather than throw away 75 years of surface temperature and forcing data by truncating the observational record to 1955, we estimate the missing ocean temperature data along with climate sensitivity (see also Aldrin et al., 2012).[25] Our ability to easily estimate multiple uncertain parameters at once is a strength of our nonconjugate approach to learning, as compared to the standard conjugate approach to learning. These missing ocean temperature data are identified, first, by variation in contemporary surface temperature that is not explained by variation in contemporary forcing and, second, by the ocean temperature state's transition equations

[23] Since policymaking only occurs every 10 years, each time period for the value function is a decade even though the policymaker is still using the annual data obtained from the Brownian bridge step.

[24] Large dips in radative forcing are generally caused by reflective particulates being ejected by volcanic eruptions such as Krakatoa in 1883 or Pinatubo in 1991.

[25] Uncertainty about ocean temperature is a key part of actual scientific uncertainty about climate change.

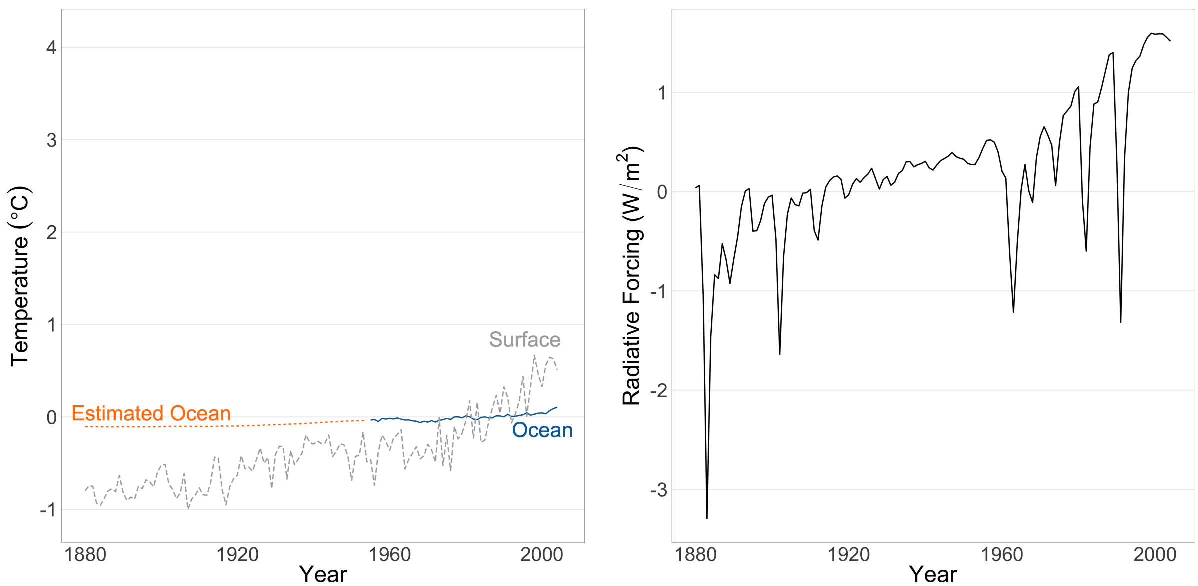

Figure 4: Observations of surface temperature (gray), ocean temperature (blue, 0–700m depth, vertically averaged), and forcing (due to greenhouse gases such as CO_2) from 1880–2005. Surface temperature is from the NOAA National Climatic Data Center, ocean temperature is from the NOAA National Oceanographic Data Center, and forcing is from the NASA Goddard Institute for Space Studies. The plotted ocean temperature for 1880–1955 (orange) is the mean estimate.

that link the missing years to the 1955 measurement. For each missing ocean temperature data point, we use an improper (uniform) prior over the real numbers. We estimate fairly tight confidence intervals around each data point. We will therefore fix each missing ocean temperature data point at its estimated mean for the dynamic programming implementation.

Second, we must also estimate a scale adjustment to the post-1955 ocean temperature series. Both ocean temperature and surface temperature are measured as anomalies with respect to some baseline. The climate system's transition equations depend on the two measurements being on the same scale, but the absolute scale for the ocean temperature series is not available and the reported anomaly is not necessarily based on the same reference years as is the reported surface temperature anomaly. In order to integrate these data series, we follow Aldrin et al. (2012) in estimating the adjustment factor for reported ocean temperature from the data. This adjustment is identified from the post-1955 relationship between reported ocean and surface temperature anomalies. We follow Aldrin et al. (2012) in using an improper (uniform) prior over the real numbers. Uncertainty about this adjustment is another component of actual scientific uncertainty about the state of the ocean.

Third, we also estimate the standard deviation of the disturbance term in the transition equation for surface temperature and the standard deviation of the disturbance term in the transition equation for ocean temperature. The disturbance term in the surface temperature transition equation is sometimes treated as reflecting "weather" shocks. We estimate a tight confidence interval around each standard deviation from the historical data, even though we use an improper (uniform) prior over the positive real numbers. We will therefore fix each standard deviation at its estimated mean for the dynamic programming implementation.

Finally, we calibrate the 1880 prior over climate sensitivity so that the year 2005 posterior matches the current state of scientific knowledge. The climate sensitivity parameter is identified by changes in surface temperature that are residual to changes in ocean temperature and are correlated with changes in forcing. The standard deviation of the surface temperature transition is identified by the remaining variation in the surface temperature series. We find that a uniform prior yields posterior distributions that are on the high end of the current state of scientific knowledge. Further, theoretical derivations of climate sensitivity suggest that a uniform prior does not reflect the knowledge available before looking at the data (e.g., Roe and Baker, 2007).[26] In contrast, we find that a lognormal prior over climate sensitivity generates a posterior with an acceptable fit to current scientific understanding.[27] We calibrate the prior so as to minimize the Kullback-Leibler distance between our year 2005 posterior and the distribution of climate sensitivity derived in the recent review by Knutti et al. (2017).[28] We then fit a lognormal distribution to the 2005 nonparametric density that minimized the Kullback-Leibler divergence and use this lognormal posterior as our prior for the dynamic programming implementation.[29] We obtain a lognormal density in 2005 that has location parameter 1.28 and scale parameter 0.602, and our 1880 prior has a location parameter of 1.73 and a scale parameter of 0.604. Arrhenius (1896) published the first recorded estimate of the climate sensitivity shortly after our observational record begins. This estimate is consistent with our calibrated 1880 prior.

Figure 5 plots the evolution of the climate sensitivity's posterior distribution from our learning model since 1880. We see that the last 120 years' observations have generated

[26] A uniform prior is not uninformative but implies a certain type of knowledge about the climate system: that climate sensitivity values as high as 20°C are as likely as those around $3-4$°C. Actual theoretical understanding is in fact different from this implied knowledge.

[27] Consistent with Aldrin et al. (2012), we truncate the distribution of climate sensitivity at 20°C. We also truncate the lower end of the distribution at 1.2°C under the standard assumption that the total climate feedback is nonnegative.

[28] Specifically, we target the observationally constrained distribution in their Figure 5b. We take 100,000 draws from this target distribution and drop all values of climate sensitivity that are less than 0°C or greater than 20°C. We then use a kernel smoother to fit a density over the same Gauss-Legendre quadrature nodes used in the dynamic programming distribution. We use a differential evolution algorithm to search for the 1880 prior that minimizes the Kullback-Leibler distance between the target density at these nodes and the density of the posterior distribution at these same nodes.

[29] Fitting a lognormal to the year 2005 posterior avoids having to re-estimate the model on the entire historical series at every node in the dynamic programming routine described in Section 2, which saves a substantial amount of computational time.

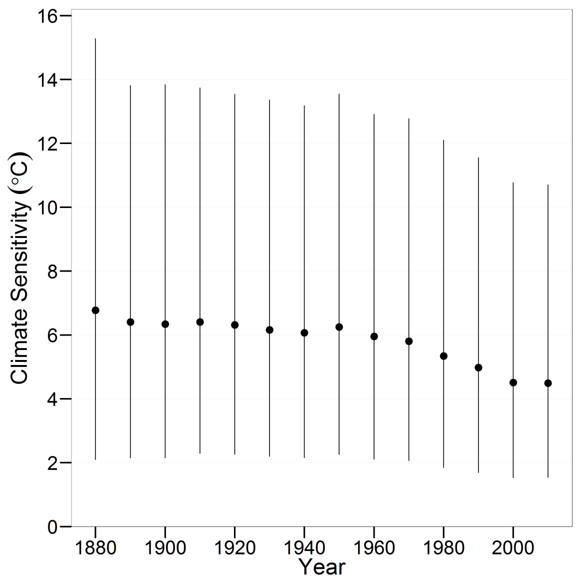

Figure 5: The evolution of the distribution for climate sensitivity as observations accrue. The point is the mean of the distribution and the line indicates the 90% confidence interval.

nontrivial information about climate sensitivity, but we also see that the learning process has been slow.

We implement the learning model using Hamiltonian Monte Carlo, a higher-performance analogue of Markov chain Monte Carlo.[30] Figure 6 shows our original prior in 1880, the 2005 target distribution from Knutti et al. (2017), and the lognormal distribution in 2005 that minimized the Kullback-Leibler divergence. The lognormal distribution is a reasonable approximation to the target distribution.[31]

4 DICE Climate-Economy Model

We integrate our learning model into the benchmark DICE climate-economy model (Nordhaus, 1992, 2008, 2013). As illustrated in Figure 7, DICE is a Ramsey growth model coupled

[30]Hamiltonian Monte Carlo is helpful in high-dimensional spaces because the majority of the volume has very low probability density even as the area of high density (near the mode) occupies very little volume. Markov chain Monte Carlo's random walk algorithm tends to explore regions that are not highly relevant, but Hamiltonian Monte Carlo restricts exploration to the relevant region of the state space. See Neal (2011) and Betancourt (2017) for more on Hamiltonian Monte Carlo.

[31]We roll the uncertainty from the uncertain ocean parameters into the year 2005 prior for climate sensitivity in the dynamic programming application. Our approach is acceptable because these parameter estimates are stable by 2005.

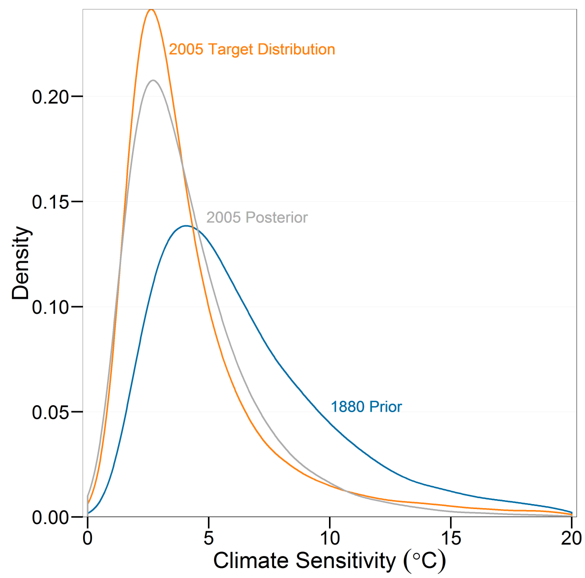

Figure 6: The year 2005 target distribution (orange), the optimized lognormal 2005 posterior (gray), and the lognormal 1880 prior that minimizes the Kullback-Leibler divergence between the lognormal posterior and the target distribution (blue).

to a climate module where the policymaker's objective is to maximize the present value of the flow of expected utility. We use a 600-year horizon. In each policymaking period (here a decade), a final good is produced from a Cobb-Douglas production function with capital and labor inputs and with exogenously improving technology. The policymaker can use the final good in three ways: for consumption, for abatement of industrial CO_2 emissions, or for investment that increases the future capital stock. Consuming output increases the current period's flow utility, where the utility function has the usual power (constant relative risk aversion) form.

The policymaker can devote some output to abatement, which reduces the net quantity of CO_2 emissions that enters the atmosphere. Each period, CO_2 can move between the atmosphere and the land/ocean CO_2 reservoir and between the land/ocean CO_2 reservoir and the deep ocean CO_2 reservoir. Additional atmospheric CO_2 traps more heat at the earth's surface and raises surface temperature. Climate sensitivity, our uncertain parameter of interest, governs how additional atmospheric CO_2 translates into surface warming in equilibrium. Heat from the surface can also move into the ocean and from the ocean back to the surface, depending on their relative temperatures. Increasing surface temperature causes damage to the economy, and it is this damage that may motivate a policymaker to divert output away from consumption and savings and towards abatement.

Each period, the policymaker makes decisions based on the current state of the world,

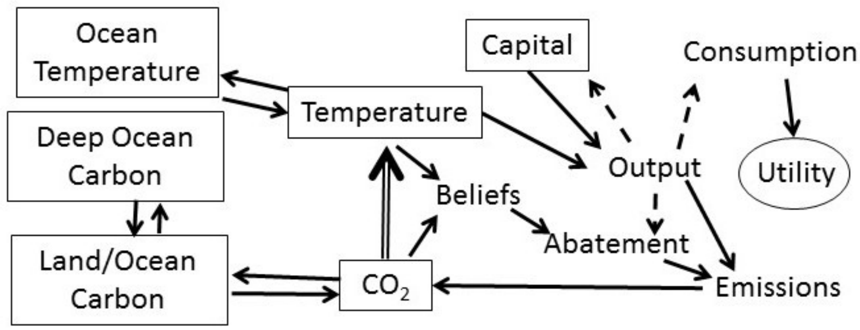

Figure 7: Schematic of the extended DICE climate-economy model. States of the model are contained in boxes.

which consists of the observable states (boxes in Figure 7) and her beliefs about the true value of the climate sensitivity. The policymaker's beliefs in each period are formed by combining current and past observations of atmospheric CO_2, surface temperature, and ocean temperature with her initial prior in the standard Bayes' formula.

5 Results

We now present the results of analyzing climate change policy under realistic learning about climate sensitivity. First, we describe the evolution of the policymaker's beliefs along the optimal path for different true values of the climate sensitivity. Second, we describe how the optimal carbon tax evolves under these different climate sensitivities.

Figure 8 shows the evolution of the policymaker's beliefs along the optimal emission path when the true climate sensitivity is either 2°C, 3°C, 4°C, or 5°C.[32] The points give the mean of the policymaker's belief and the bars depict the 90% confidence interval. We will discuss the lower left panel's gray dots below. In the first period (year 2005), the policymaker's expected climate sensitivity is 4.3°C and the standard deviation of her beliefs is 2.8°C.[33] The top left panel shows the evolution of beliefs for a true climate sensitivity of 2°C. The 90% confidence interval declines relatively slowly over the next century, with the standard deviation about 60% smaller by the end of the century. At that time, the policymaker's expected climate sensitivity is 2.5°C, still a bit larger than the true value.

We see that the policymaker learns more slowly when the climate is actually more sensitive to emissions, an effect that is consistent with climate scientists' projections (e.g., Allen

[32]Note that the optimal emission path depends on the true climate sensitivity because emissions at any future time t depend on temperature and beliefs at that time, both of which respond to changes in the true climate sensitivity.

[33]Recall that this initial belief is obtained from the historical record, as described in Section 3.

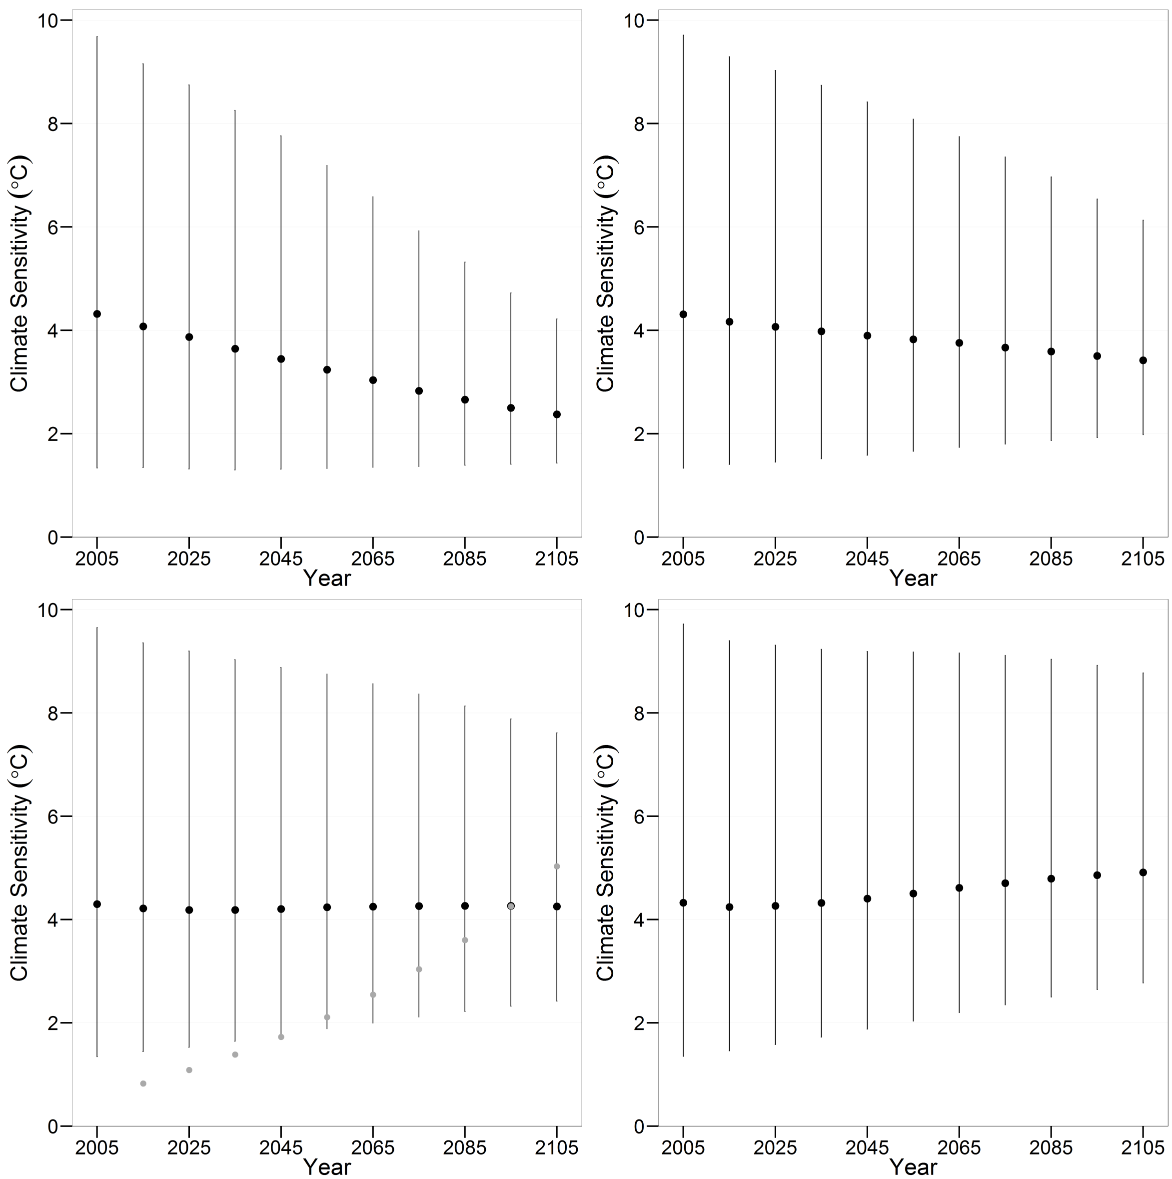

Figure 8: Evolution of the mean (points) and 90% confidence interval (bars) of climate sensitivity beliefs along the optimal emission trajectory when the true climate sensitivity is 2 (top left), 3 (top right), 4 (bottom left), or 5 (bottom right) degrees Celsius. We sample 1000 Brownian bridges to compute the expected posterior and evaluate decadal temperature shocks at their mean. The gray dots in the lower left panel indicate the single value of climate sensitivity that would generate that year's expected temperature along the emission trajectory that is optimal when the true climate sensitivity happens to be 4°C.

and Frame, 2007). When the true climate sensitivity is around 3°C, the year 2105 policymaker's expectation of climate sensitivity is 3.5°C, but the standard deviation of her beliefs is only around 45% smaller than it was at the beginning of this century. Learning is even slower for high values of climate sensitivity. When facing a true climate sensitivity of 4°C, the policymaker's end-of-century expectation remains around 4.3°C, but her standard deviation has now fallen by only 33% over the century. And if the true climate sensitivity is 5°C, then the policymaker does not start becoming much more certain about climate sensitivity until near the end of the first century. At 2105 she still has a modestly incorrect expectation that the climate sensitivity is 4.9°C and her standard deviation has fallen by only 25%, mainly because of ruling out low climate sensitivity values but not high ones.

These results teach us two important lessons, both of which are consistent with climate science. First, the policymaker remains quite uncertain about the true climate sensitivity in all cases. Our nonconjugate policymaker expects to learn much less over the next 100 years than standard conjugate climate-economy models' policymakers would learn even over 40 years (see Figure 1). The nonconjugate learning dynamics are more consistent with the scientific literature than are the learning dynamics implied by previous economic models of climate change. Second, learning is slower when the true climate sensitivity is higher. This is the opposite of what occurs in standard conjugate learning models. There, a greater true climate sensitivity actually accelerates learning because it produces higher temperatures, which increase the climate sensitivity signal relative to the noise in the temperature transition (and this signal-to-noise ratio is all that matters in a normal-normal conjugate learning model). This fast learning in the face of undesirable climate sensitivity outcomes reduces the importance of uncertainty because future policymakers can adjust policy once they learn that they live in a bad state of the world. Our nonconjugate learning model eliminates this unrealistic effect.

Now consider optimal emission policy. Figure 9 displays the optimal emission tax trajectories for a policymaker who is a nonconjugate learner and for a counterfactual policymaker who has perfect information. The panels correspond to climates where the true climate sensitivity is either 2°C, 3°C, 4°C, or 5°C. In all cases the learning policymaker sets the same initial carbon tax of \$7/tCO_2 because initial information does not vary by scenario.

If the true climate sensitivity is 2°C, then the initial emission tax with perfect information is only \$5/tCO_2, or 30% smaller than the optimal emission tax under uncertainty. Since the climate is relatively insensitive to CO_2, additional CO_2 does not cause much damage and the perfect-information tax rises slowly over time to only \$31/tCO_2 in 2105. The learning policymaker ends up following a flatter tax trajectory, reaching only \$26/tCO_2 in 2105. When the true climate sensitivity is 3°C, the perfectly informed policymaker actually selects a greater initial emission tax than does the learning policymaker. Further, the slow resolution of uncertainty means that the learning policymaker's emission tax ends up substantially smaller than the perfectly informed policymaker's emission tax at the end of the century. Because the learning policymaker's emission tax does not soon adjust towards the perfect-

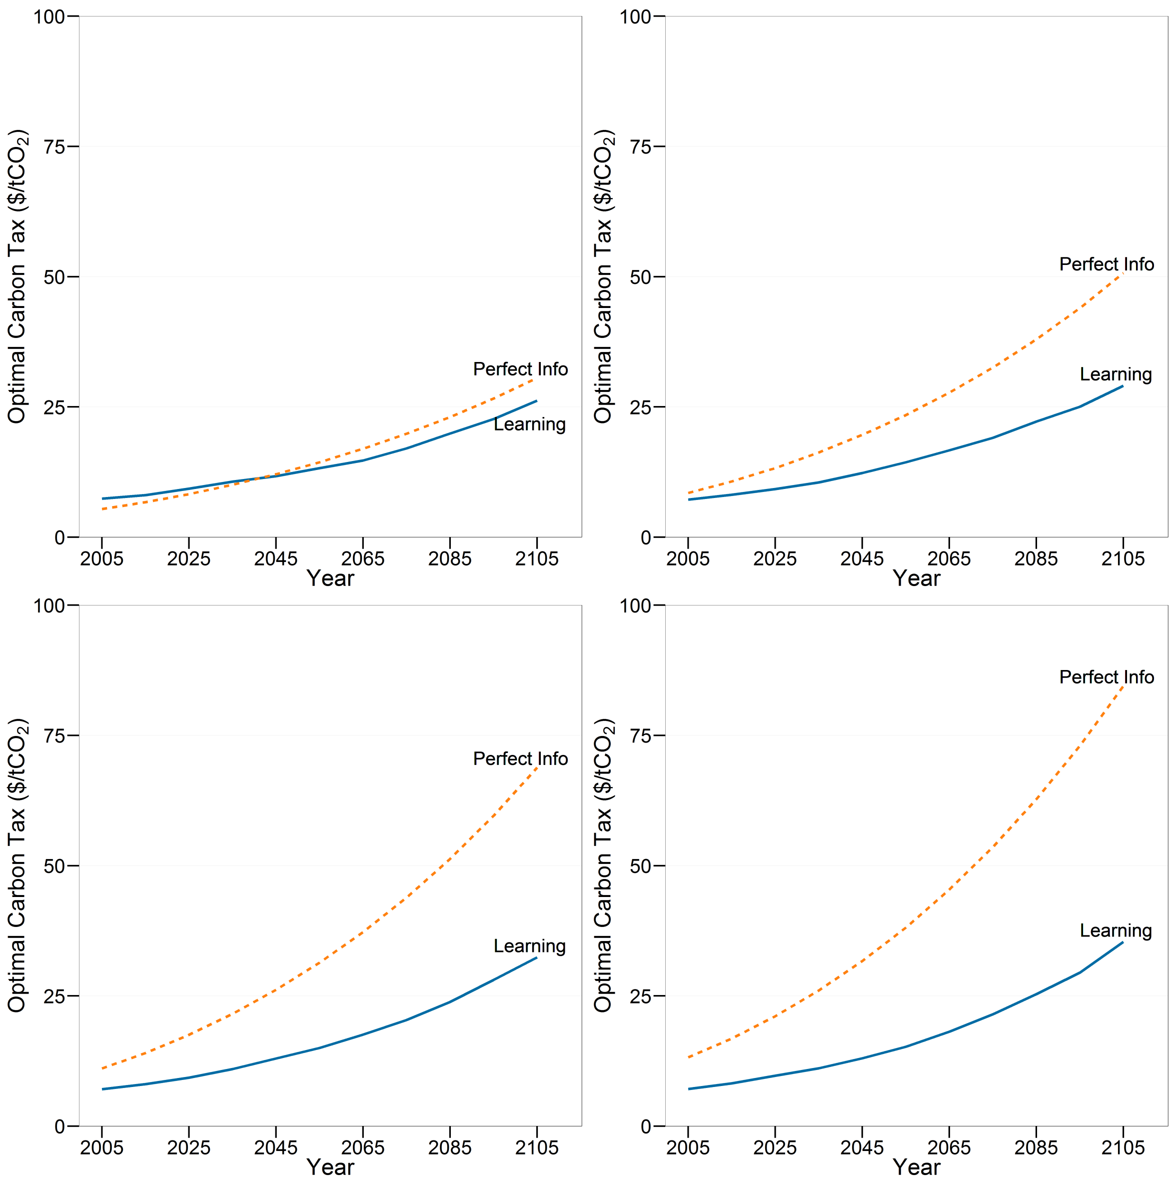

Figure 9: Evolution of the efficient carbon tax ($\$/tCO_2$) when the policymaker has perfect information (dashed orange) and when the policymaker is uncertain but learns about the climate sensitivity (solid blue). The panels correspond to a true climate sensitivity of 2°C (top left), 3°C (top right), 4°C (bottom left), and 5°C (bottom right). We sample 1000 Brownian bridges to compute the expected posterior, and we evaluate decadal temperature shocks at their mean.

information values, temperature (plotted in the appendix) reaches a level that is 0.1°C greater in 2105 than the level that it would reach under the perfectly informed policy.

Why does uncertainty reduce the initial carbon tax relative to a world in which we know that climate sensitivity would be less than its expected value?[34] Equilibrium temperature is linear in climate sensitivity, but climate sensitivity affects nearer-term temperatures nonlinearly. In particular, as climate sensitivity increases in DICE-2007, oceans absorb more heat early on, so that near-term warming is concave in climate sensitivity.[35] The lower left panel of Figure 8 illustrates this effect. Each gray dot represents the unique climate sensitivity value that, given the the emission trajectory that is optimal when the true climate sensitivity happens to be 4°C, would yield the expected temperature in that year based on the year 2005 posterior. We may think of these dots as representing the certainty-equivalent climate sensitivity when outcomes are measured in terms of temperature (i.e., before accounting for nonlinear damages and risk aversion). With year 2005 information, the expected climate sensitivity is around 4.3°C, but we see that the certainty-equivalent climate sensitivity is substantially smaller over much of the next century. Uncertainty therefore reduces the certainty-equivalent climate sensitivity because of the strong concavity of near-term temperature in climate sensitivity. And for a policymaker who discounts the future, this near-term reduction in certainty-equivalent climate sensitivity can be more important than the increase in expected long-term warming.[36]

[34]In their review, Lemoine and Rudik (2017) describe uncertainty about climate sensitivity as having either nearly no effect or a very small positive effect in previous settings.

[35]This effect is in part why it is so difficult to constrain climate sensitivity from the observational record and is consistent with scientific modeling (e.g., Baker and Roe, 2009).

[36]Following Calel et al. (2015) and ignoring the deep ocean, a continuous-time version of DICE's temperature dynamics is $\dot{T}(t) = [F(t) - [F_2/s]T(t)]/C$, where $F(t)$ is time t forcing, F_2 is forcing from doubled CO_2, s is climate sensitivity, and C is the effective heat capacity of the climate system. Solving the differential equation, we have:

$$T(t) = e^{-\frac{F_2}{s\,C}(t-t_0)}T(t_0) + \frac{1}{C}\int_{t_0}^{t} e^{-\frac{F_2}{s\,C}(t-i)}F(i)\,\mathrm{d}i.$$

Let time t_0 emissions be e_0. Then

$$\frac{\mathrm{d}T(t)}{\mathrm{d}e_0} = \frac{1}{C}\int_{t_0}^{t} e^{-\frac{F_2}{s\,C}(t-i)}\frac{\mathrm{d}F(i)}{\mathrm{d}e_0}\,\mathrm{d}i > 0.$$

Twice-differentiate with respect to s:

$$\frac{\mathrm{d}^3T(t)}{\mathrm{d}e_0\,\mathrm{d}s^2} = \frac{F_2}{s^3C^2}\left[-2\int_{t_0}^{t} e^{-\frac{F_2}{s\,C}(t-i)}\frac{\mathrm{d}F(i)}{\mathrm{d}e_0}(t-i)\,\mathrm{d}i + \frac{F_2}{sC}\int_{t_0}^{t} e^{-\frac{F_2}{s\,C}(t-i)}\frac{\mathrm{d}F(i)}{\mathrm{d}e_0}(t-i)^2\,\mathrm{d}i\right].$$

The first term is negative and the second is positive. Note that $\mathrm{d}F(i)/\mathrm{d}e_0 \geq 0$ declines with i because time 0 emissions' effect on later atmospheric CO_2 decays over time. Also note that DICE uses a value for F_2 of around 3.8 and, according to calculations in Calel et al. (2015), a value for C of around 1.8, so $F_2/sC < 2$ for $s > 1$ and $F_2/sC < 1$ for $s > 2.1$. Thus, for all relevant values of s, the first integral is multiplied by a

The optimal tax trajectories diverge even more strongly when the true climate sensitivity is greater than 3°C. In these cases, the perfectly informed policymaker would set an even greater initial emission tax and increase it even more strongly over time. The learning policymaker does eventually raise her tax to greater levels than in the cases with a smaller true climate sensitivity, but because she learns especially slowly when climate sensitivity is high, the divergence between her policy and the perfectly informed policy becomes especially severe: the learning policymaker's year 2105 emission tax is only \$33/tCO_2 (\$36/tCO_2) when the true climate sensitivity is 4°C (5°C), but the perfectly informed policymaker's year 2105 emission tax would be \$69/tCO_2 (\$85/tCO_2). As a result, end-of-century emissions are 32% (46%) greater for the learning policymaker and end-of-century temperature is 0.2°C (0.3°C) greater.

We have seen that our nonconjugate learning model predicts that we will learn something about climate sensitivity over the next century but that progress will be slow. We now calculate the expected value of the information gained from learning relative to not updating our initial beliefs. The balanced growth equivalent (BGE) gain (Mirrlees and Stern, 1972) is the constant relative difference in consumption between two counterfactual consumption trajectories that grow at the same constant rate and also yield the exact same welfare as two models we are interested in (in our case, learning and not learning). We find that the BGE gain from nonconjugate learning relative to never updating beliefs from the initial prior is 0.3%. The greatest possible gain from learning would arise if the policymaker could expect to learn the true value of climate sensitivity right before she begins policymaking. This BGE tells us the expected welfare gain from perfect information. We find that expecting to obtain perfect information yields a BGE gain of 8.5% relative to a world without learning. This gain is nearly thirty times larger than the gain from nonconjugate learning. Better scientific monitoring and modeling could be hugely valuable if it would speed learning towards the perfect information case. In ongoing work, we are calculating the welfare gains from reducing the unexplained volatility of temperature observations, as would be achieved by directing research toward improving ocean and surface monitoring systems or improving the accuracy of climate models.

6 Conclusions

TBA

larger term, and often by a much larger term. Because the second integral becomes especially small for s large, concavity becomes more important when more weight is placed on large values of climate sensitivity. Alternatively, using the standard discrete-time DICE temperature transition, it is easy to show that the marginal effect of emissions on temperature two periods ahead (climate sensitivity and emissions do not interact only one period ahead) is strictly concave in climate sensitivity.

References

Aldrin, Magne, Marit Holden, Peter Guttorp, Ragnhild Bieltvedt Skeie, Gunnar Myhre, and Terje Koren Berntsen (2012) "Bayesian estimation of climate sensitivity based on a simple climate model fitted to observations of hemispheric temperatures and global ocean heat content," *Environmetrics*, Vol. 23, No. 3, pp. 253–271.

Allen, Myles R. and David J. Frame (2007) "Call off the quest," *Science*, Vol. 318, No. 5850, pp. 582–583.

Andronova, Natalia, Michael Schlesinger, Suraje Dessai, Mike Hulme, and Bin Li (2007) "The concept of climate sensitivity: history and development," in Michael E. Schlesinger, Haroon S. Kheshgi, Joel Smith, Francisco C. de la Chesnaye, John M. Reilly, Tom Wilson, and Charles Kolstad eds. *Human-Induced Climate Change: An interdisciplinary assessment*, New York: Cambridge University Press, pp. 5–17.

Arrhenius, SA (1896) "Nature's heat usage," *Nordisk Tidskrift*, Vol. 14, pp. 121–130.

Arrhenius, Svante (1906) *Världarnas utveckling*: H. Geber.

Baker, Marcia B. and Gerard H. Roe (2009) "The shape of things to come: Why is climate change so predictable?" *Journal of Climate*, Vol. 22, No. 17, pp. 4574–4589.

Barro, Robert J. (2006) "Rare disasters and asset markets in the twentieth century," *The Quarterly Journal of Economics*, Vol. 121, No. 3, pp. 823–866.

Barthelmann, Volker, Erich Novak, and Klaus Ritter (2000) "High dimensional polynomial interpolation on sparse grids," *Advances in Computational Mathematics*, Vol. 12, No. 4, pp. 273–288.

Betancourt, Michael (2017) "A conceptual introduction to Hamiltonian Monte Carlo," *arXiv:1701.02434 [stat]*, arXiv: 1701.02434.

Brennan, M. J. (1998) "The role of learning in dynamic portfolio decisions," *Review of Finance*, Vol. 1, No. 3, pp. 295–306.

Brennan, Michael J. and Yihong Xia (2001) "Stock price volatility and equity premium," *Journal of Monetary Economics*, Vol. 47, No. 2, pp. 249–283.

Brumm, Johannes and Simon Scheidegger (2016) "Using adaptive sparse grids to solve high-dimensional dynamic models," working Paper.

Cai, Yongyang, Kenneth L. Judd, and Thomas S. Lontzek (2015) "The social cost of carbon with economic and climate risks," *arXiv:1504.06909 [q-fin]*, arXiv: 1504.06909.

Calel, Raphael, David A. Stainforth, and Simon Dietz (2015) "Tall tales and fat tails: the science and economics of extreme warming," *Climatic Change*, Vol. 132, No. 1, pp. 127–141.

Charney, Jule G., Akio Arakawa, D. James Baker, Bert Bolin, Robert E. Dickinson, Richard M. Goody, Cecil E. Leith, Henry M. Stommel, and Carl I. Wunsch (1979) "Carbon Dioxide and Climate: A Scientific Assessment,"Technical report, National Academy of Sciences, Washington, D.C.

Cogley, Timothy and Thomas J. Sargent (2008) "Anticipated utility and rational expectations as approximations of bayesian decision making," *International Economic Review*, Vol. 49, No. 1, pp. 185–221.

Collin-Dufresne, Pierre, Michael Johannes, and Lars A. Lochstoer (2016) "Parameter learning in general equilibrium: The asset pricing implications," *American Economic Review*, Vol. 106, No. 3, pp. 664–698.

Collins, M., R. Knutti, J. Arblaster, J.-L. Dufresne, T. Fichefet, P. Friedlingstein, X. Gao, W. J. Gutowski, T. Johns, G. Krinner, M. Shongwe, C. Tebaldi, A. J. Weaver, and M. Wehner (2013) "Long-term climate change: Projections, commitments and irreversibility," in T. F. Stocker, D. Qin, G.-K. Plattner, M. Tignor, S. K. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex, and P. M. Midgley eds. *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, Cambridge, United Kingdom and New York, NY, USA: Cambridge University Press.

Gonzalez, A Gavilan and JA Rojas (2009) "Solving Portfolio Problems with the Smolyak-Parameterized Expectations Algorithm."

Hannart, Alexis, Michael Ghil, Jean-Louis Dufresne, and Philippe Naveau (2013) "Disconcerting learning on climate sensitivity and the uncertain future of uncertainty," *Climatic Change*, Vol. 119, No. 3-4, pp. 585–601.

Hansen, James, Makiko Sato, Pushker Kharecha, David Beerling, Robert Berner, Valerie Masson-Delmotte, Mark Pagani, Maureen Raymo, Dana L. Royer, and James C. Zachos (2008) "Target atmospheric CO2: Where should humanity aim?" *The Open Atmospheric Science Journal*, Vol. 2, pp. 217–231.

Hoffman, Matthew D and Andrew Gelman (2014) "The No-U-turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo.," *Journal of Machine Learning Research*, Vol. 15, No. 1, pp. 1593–1623.

Hwang, In Chang, Frdric Reyns, and Richard S. J. Tol (2017) "The effect of learning on climate policy under fat-tailed uncertainty," *Resource and Energy Economics*, Vol. 48, pp. 1–18.

Jensen, Svenn and Christian Traeger (2016) "Pricing climate risk," working paper.

Johnson, Timothy C. (2007) "Optimal learning and new technology bubbles," *Journal of Monetary Economics*, Vol. 54, No. 8, pp. 2486–2511.

Judd, Kenneth L (1992) "Projection methods for solving aggregate growth models," *Journal of Economic Theory*, Vol. 58, No. 2, pp. 410–452.

Judd, Kenneth L. (1998) *Numerical Methods in Economics*, Cambridge, Mass.: MIT Press.

Judd, Kenneth L., Lilia Maliar, Serguei Maliar, and Inna Tsener (2017) "How to solve dynamic stochastic models computing expectations just once," *Quantitative Economics*, Vol. 8, No. 3, pp. 851–893.

Judd, Kenneth L., Lilia Maliar, Serguei Maliar, and Rafael Valero (2014) "Smolyak method for solving dynamic economic models: Lagrange interpolation, anisotropic grid and adaptive domain," *Journal of Economic Dynamics and Control*, Vol. 44, pp. 92–123.

Kelly, David L. and Charles D. Kolstad (1999) "Bayesian learning, growth, and pollution," *Journal of Economic Dynamics and Control*, Vol. 23, No. 4, pp. 491–518.

Kelly, David L. and Zhuo Tan (2015) "Learning and climate feedbacks: Optimal climate insurance and fat tails," *Journal of Environmental Economics and Management*, Vol. 72, pp. 98–122.

Kiehl, Jeffrey (2011) "Lessons from Earth's past," *Science*, Vol. 331, No. 6014, pp. 158 –159.

Kling, David M., James N. Sanchirico, and Paul L. Fackler (2017) "Optimal monitoring and control under state uncertainty: Application to lionfish management," *Journal of Environmental Economics and Management*, Vol. 84, No. Supplement C, pp. 223–245.

Knutti, Reto and Gabriele C. Hegerl (2008) "The equilibrium sensitivity of the Earth's temperature to radiation changes," *Nature Geoscience*, Vol. 1, No. 11, pp. 735–743.

Knutti, Reto, Maria A. A. Rugenstein, and Gabriele C. Hegerl (2017) "Beyond equilibrium climate sensitivity," *Nature Geoscience*, No. September.

Kreps, David (1998) "Anticipated utility and dynamic choice," in Donald P. Jacobs, Ehud Kalai, and Morton Kamien eds. *Frontiers of Research in Economic Theory*, Cambridge, England: Cambridge University Press.

Krueger, Dirk and Felix Kubler (2004) "Computing equilibrium in OLG models with stochastic production," *Journal of Economic Dynamics and Control*, Vol. 28, No. 7, pp. 1411–1436.

Leach, Andrew J. (2007) "The climate change learning curve," *Journal of Economic Dynamics and Control*, Vol. 31, No. 5, pp. 1728–1752.

Lemoine, Derek M. (2010) "Climate sensitivity distributions depend on the possibility that models share biases," *Journal of Climate*, Vol. 23, No. 16, pp. 4395–4415.

Lemoine, Derek and Ivan Rudik (2017) "Managing climate change under uncertainty: Recursive integrated assessment at an inflection point," *Annual Review of Resource Economics*, Vol. 9, p. forthcoming.

Lewis, Nicholas (2013) "An objective Bayesian improved approach for applying optimal fingerprint techniques to estimate climate sensitivity," *Journal of Climate*, Vol. 26, No. 19, pp. 7414–7429.

Lunt, Daniel J., Alan M. Haywood, Gavin A. Schmidt, Ulrich Salzmann, Paul J. Valdes, and Harry J. Dowsett (2010) "Earth system sensitivity inferred from Pliocene modelling and data," *Nature Geoscience*, Vol. 3, No. 1, pp. 60–64.

MacLachlan, Matthew J., Michael R. Springborn, and Paul L. Fackler (2017) "Learning about a moving target in resource management: Optimal Bayesian disease control," *American Journal of Agricultural Economics*, Vol. 99, No. 1, pp. 140–162.

Maliar, Lilia and Serguei Maliar (2005) "Parameterized expectations algorithm: how to solve for labor easily," *Computational Economics*, Vol. 25, No. 3, pp. 269–274.

——— (2014) "Numerical methods for large-scale dynamic economic models," in Karl Schmedders and Kenneth L. Judd eds. *Handbook of Computational Economics*, Vol. 3: Elsevier, pp. 325–477, doi: 10.1016/B978-0-444-52980-0.00007-4.

——— (2015) "Merging simulation and projection approaches to solve high-dimensional problems with an application to a new Keynesian model," *Quantitative Economics*, Vol. 6, No. 1, pp. 1–47.

Meehl, G.A., T. F. Stocker, W. D. Collins, P. Friedlingstein, A.T. Gaye, J. M. Gregory, A. Kitoh, R. Knutti, J. M. Murphy, A. Noda, S. C. B. Raper, I.G. Watterson, A. J. Weaver, and Z.-C. Zhao (2007) "Global climate projections," in Susan Solomon, D. Qin, M. Manning, Z. Chen, M. Marquis, K.B. Averyt, M. Tignor, and H.L. Miller eds. *Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*, Cambridge, United Kingdom and New York, NY, USA: Cambridge University Press, pp. 748–845.

Millner, Antony and Thomas K. J. McDermott (2016) "Model confirmation in climate economics," *Proceedings of the National Academy of Sciences*, Vol. 113, No. 31, pp. 8675–8680.

Miranda, Mario J. and Paul L. Fackler (2002) *Applied Computational Economics and Finance*, Cambridge, Massachusetts: MIT Press.

Mirrlees, J. A and N. H Stern (1972) "Fairly good plans," *Journal of Economic Theory*, Vol. 4, No. 2, pp. 268–288.

Mitchell, J. F. B., S. Manabe, V. Meleshko, and T. Tokioka (1990) "Equilibrium climate change—and its implications for the future," in J. T. Houghton, G. J. Jenkins, and J. J. Ephraums eds. *Climate Change: The IPCC Scientific Assessment*, Cambridge: Cambridge University Press, pp. 131–172.

Neal, Radford M. (2011) "MCMC using Hamiltonian dynamics," in Steve Brooks, Andrew Gelman, Galin L. Jones, and Xiao-Li Meng eds. *Handbook of Markov Chain Monte Carlo*: Chapman & Hall/CRC, pp. 113–162, arXiv: 1206.1901.

Nordhaus, William (2013) "Integrated economic and climate modeling," in Peter B. Dixon and Dale W. Jorgenson eds. *Handbook of Computable General Equilibrium Modeling*, Vol. 1: North Holland, Elsevier B.V. pp. 1069–1131.

Nordhaus, William D. (1992) "An optimal transition path for controlling greenhouse gases," *Science*, Vol. 258, No. 5086, pp. 1315–1319.

——— (2008) *A Question of Balance: Weighing the Options on Global Warming Policies*, New Haven: Yale University Press.

Orlik, Anna and Laura Veldkamp (2014) "Understanding uncertainty shocks and the role of black swans," Working Paper 20445, National Bureau of Economic Research.

Otto, Alexander, Friederike EL Otto, MR Allen, O Boucher, J Church, G Hegerl, PM Forster, NP Gillett, J Gregory, GC Johnson et al. (2013) "Energy budget constraints on climate response," *Nature Geoscience*, Vol. 6, No. 6, pp. 415–416.

Pagani, Mark, Zhonghui Liu, Jonathan LaRiviere, and Ana Christina Ravelo (2010) "High Earth-system climate sensitivity determined from Pliocene carbon dioxide concentrations," *Nature Geoscience*, Vol. 3, No. 1, pp. 27–30.

Pastor, Lubos and Pietro Veronesi (2009a) "Learning in financial markets," *Annual Review of Financial Economics*, Vol. 1, No. 1, pp. 361–381.

Pástor, Luboš and Pietro Veronesi (2009b) "Technological revolutions and stock prices," *The American Economic Review*, Vol. 99, No. 4, pp. 1451–1483.

Rietz, Thomas A. (1988) "The equity risk premium: A solution," *Journal of Monetary Economics*, Vol. 22, No. 1, pp. 117–131.

Ring, Michael J, Daniela Lindner, Emily F Cross, and Michael E Schlesinger (2012) "Causes of the global warming observed since the 19th century," *Atmospheric and Climate Sciences*, Vol. 2, No. 04, p. 401.

Roe, Gerard H. (2009) "Feedbacks, timescales, and seeing red," *Annual Review of Earth and Planetary Sciences*, Vol. 37, No. 1, pp. 93–115.

Roe, Gerard H. and Marcia B. Baker (2007) "Why is climate sensitivity so unpredictable?" *Science*, Vol. 318, No. 5850, pp. 629–632.

Royer, Dana L. (2016) "Climate sensitivity in the geologic past," *Annual Review of Earth and Planetary Sciences*, Vol. 44, No. 1, pp. 277–293.

Rudik, Ivan (2016) "Optimal climate policy when damages are unknown," working Paper, Iowa State University.

Schwartz, Stephen E (2012) "Determination of Earths transient and equilibrium climate sensitivities from observations over the twentieth century: strong dependence on assumed forcing," *Surveys in geophysics*, Vol. 33, No. 3-4, pp. 745–777.

National Academies of Sciences, Engineering (2017) *Valuing Climate Damages: Updating Estimation of the Social Cost of Carbon Dioxide*.

Skeie, R. B., T. Berntsen, M. Aldrin, M. Holden, and G. Myhre (2014) "A lower and more constrained estimate of climate sensitivity using updated observations and detailed radiative forcing time series," *Earth Syst. Dynam.*, Vol. 5, No. 1, pp. 139–175.

Smolyak, S. (1963) "Quadrature and interpolation formulas for tensor products of certain classes of functions," *Soviet Mathematics, Doklady*, Vol. 4, pp. 240–243.

Springborn, Michael and James N. Sanchirico (2013) "A density projection approach for non-trivial information dynamics: Adaptive management of stochastic natural resources," *Journal of Environmental Economics and Management*, Vol. 66, No. 3, pp. 609–624.

Traeger, Christian (2014) "A 4-stated DICE: Quantitatively addressing uncertainty effects in climate change," *Environmental and Resource Economics*, Vol. 59, No. 1, pp. 1–37.

Veronesi, Pietro (2004) "The peso problem hypothesis and stock market returns," *Journal of Economic Dynamics and Control*, Vol. 28, No. 4, pp. 707–725.

Winschel, Viktor and Markus Krätzig (2010) "Solving, estimating, and selecting nonlinear dynamic models without the curse of dimensionality," *Econometrica*, Vol. 78, No. 2, pp. 803–821.

Zhou, Enlu, M.C. Fu, and S.I. Marcus (2010) "Solving continuous-state POMDPs via density projection," *IEEE Transactions on Automatic Control*, Vol. 55, No. 5, pp. 1101–1116.

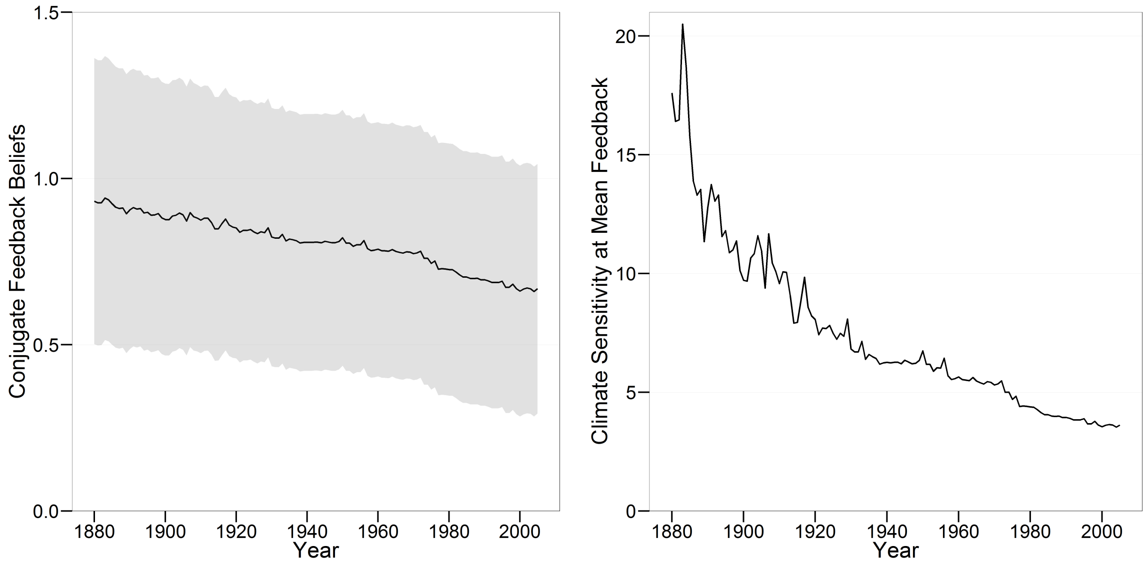

(a) Conjugate Feedback Learning in Reverse (b) The Implied Climate Sensitivity

Figure A-1: Left: The mean and 90% confidence interval of conjugate feedback beliefs when starting from the standard 2005 Roe and Baker (2007) prior and simulating the learning rules in reverse using the observational record from 2005–1880. Right: The implied belief about the climate sensitivity at the mean feedback factor.

A Appendix

The appendix contains additional details on how conjugate learning is a binding constraint to matching real world informational dynamics, plots of climate variables, more details on the Brownian bridge, the Hamiltonian Monte Carlo algorithm, how we approximate posterior beliefs, pseudocode for the entire algorithm, the DICE model, and an error analysis of our results.

A.1 Additional Details on the Binding Constraint of Conjugate Learning

Previous models' bias towards overly rapid learning has strong implications for what we must have believed in the past. The left panel of Figure A-1 shows the trajectory of beliefs about the feedback factor, a transformation of the climate sensitivity and the direct object of learning in conjugate models, that arises when taking the common 2005 prior calibration used in the literature and then simulating conjugate learning *in reverse* using the actual observational record of temperature and CO_2. This is the unique belief trajectory that would generate economic models' initialized beliefs from Bayesian learning on observational

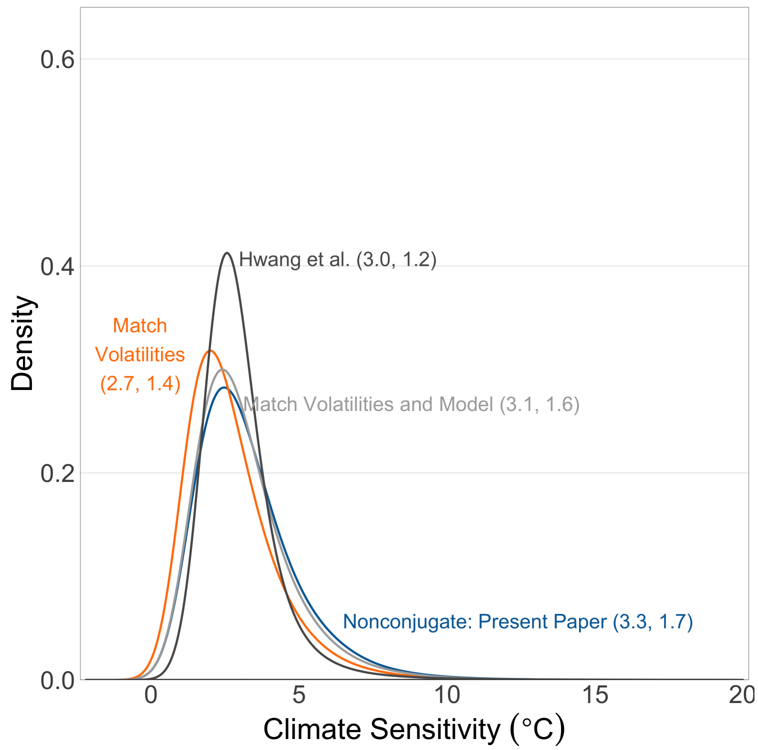

Figure A-2: Changing features of our learning model one by one to match that of Hwang et al. (2017), in order to verify why learning is slower in our setting. Beside each distribution is the mean and standard deviation. The scale of the y-axis matches Figure 1 for comparison.

data.[37] To arrive at year 2005 beliefs in these conjugate learning models, scientists would have expected the feedback to be about 0.9 in 1880. The right panel shows that at the mean feedback, this implies a climate sensitivity of 17°C, far above even the original guess of Arrhenius (1896).

Is nonconjugacy driving the major differences between our posterior and previous models, or is the main driver some other difference between the models? Figure A-2 plots the posteriors obtained by step-by-step matching our learning model to the Hwang et al. (2017) model, which was the most accurate out of the conjugate models. First, consider adapting our model to match the volatility of surface and ocean temperature in Hwang et al. (2017).[38]

[37]Note that the Roe and Baker (2007) climate sensitivity-feedback relationship is given by $\text{cs} = \frac{1.2}{1-\text{fb}}$, so a feedback factor of 1 implies runaway climate change and values higher than 1 mathematically map into negative climate sensitivity values. The common practice in economic modeling is to update beliefs using the untruncated distribution and then truncate the distribution below 1 when calculating expectations about future temperature.

[38]We match the volatilities as a separate step because we directly estimate them from the instrumental

The mean and standard deviation of the posterior both become smaller, but learning is still slower than the slowest conjugate model. Now consider also matching their model structure, so that we have the same transition equation for temperature. Matching the model structure and volatilities of Hwang et al. (2017) tends result in even smaller differences in the posterior compared to the common prior and the nonconjugate posterior. The only remaining difference between our setting and that of Hwang et al. (2017) is nonconjugate learning. Once we replace that with conjugacy, the posterior narrows significantly. Nonconjugacy is indeed the key change that drives slow learning in our setting.

A.2 Additional Figures

Figures A-3 through A-5 plot the trajectories of emissions, CO_2 concentrations, and surface temperature for a policymaker who is a nonconjugate learner and for a counterfactual policymaker who has perfect information. For all true climate sensitivities, the learning policymaker's emission trajectories follow a similar path that increases over the first century. Learning does allow the policymaker to adjust emissions by the end of the century, with emissions under a true climate sensitivity of 2°C reaching near 14.6 GtC per year but reaching only 13.3 GtC per year when the true climate sensitivity is 5°C. But learning is too slow for emissions to reach anything like their levels under the perfectly informed policy. With perfect information, higher climate sensitivities lead to drastically greater emissions cuts. At a climate sensitivity of 4°C (5°C), emissions peak in 2095 (2085) at 10.4 GtC per year (9.4 GtC per year).

Figure A-4 shows how CO_2 concentrations correspondingly evolve. The similar emission trajectories of the learning policymaker across the different climate sensitivities lead to similar CO_2 concentrations. At the end of the century, CO_2 is between 625 ppm and 636 ppm. When the true climate sensitivity is 2°C, the CO_2 trajectory under perfect information is very similar and reaches 632 ppm. However under the higher climate sensitivities, the perfectly informed policymaker's drastic emission reductions lead to substantially lower CO_2 concentrations of 605 ppm, 585 ppm, and 570 ppm for climate sensitivities of 3°C, 4°C, and 5°C. Differences in CO_2 translate into temperature only slowly, but Figure A-5 shows that temperature would already tend to be substantially lower by the end of the century under the perfectly informed policy.

A.3 Brownian Bridge Interpolation

Our collocation grid contains sets of decadal histories of temperature and CO_2, but we want to learn using annual data. We generate annual data by stochastically interpolating temperature between the decadal history observations using a Brownian bridge with drift. Let t index years. A Brownian bridge, $B(t)$ is simply a Brownian motion, $W(t)$, but with the

record using our model, and our values imply more volatility in temperature, which would slow learning.

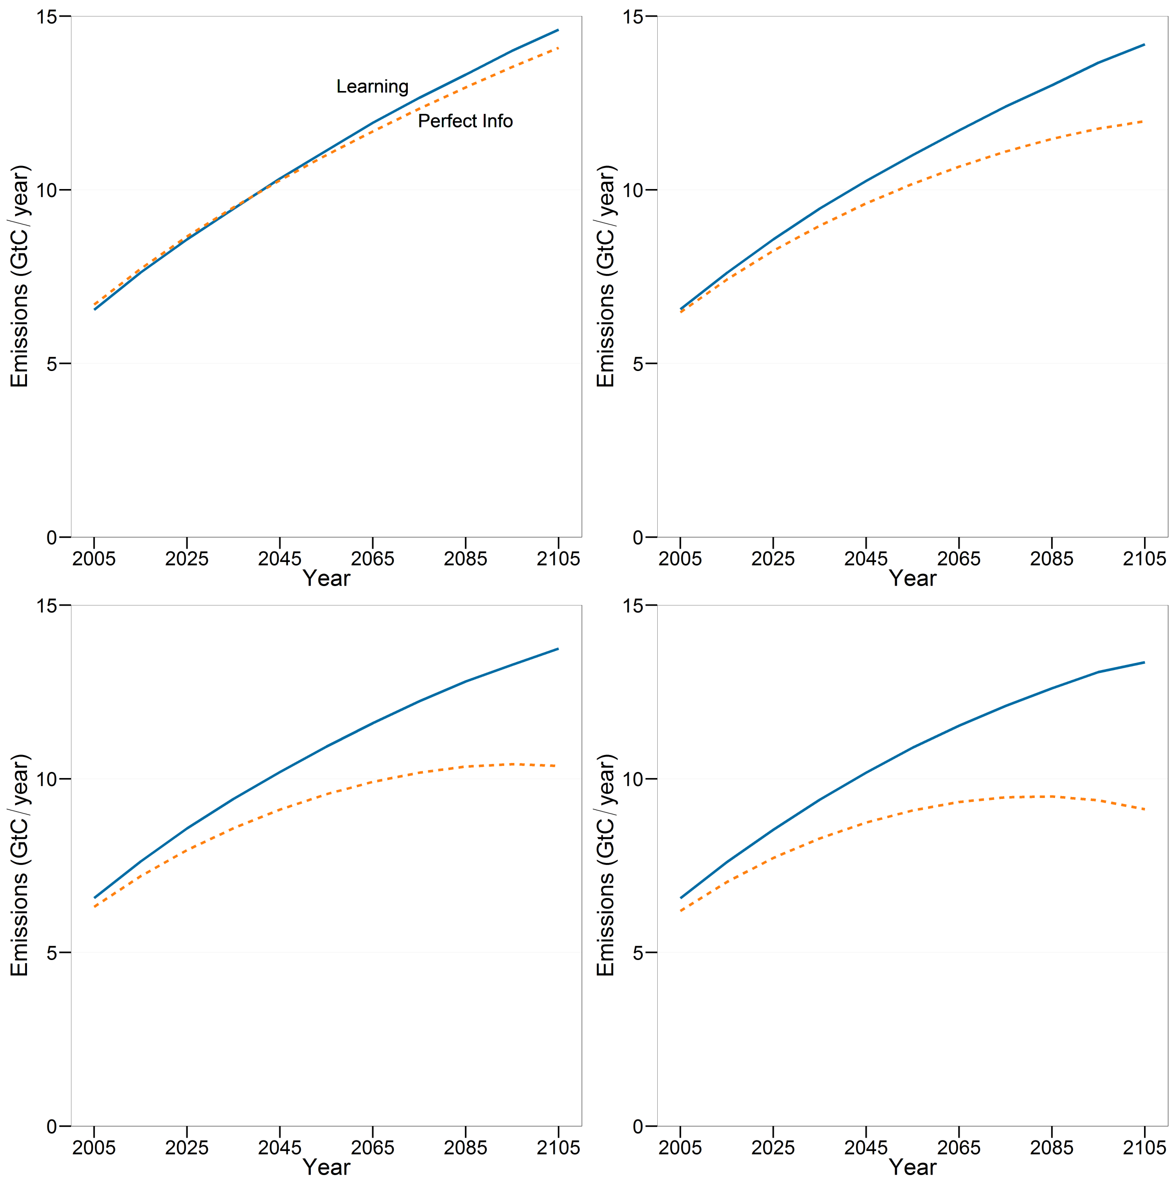

Figure A-3: Evolution of emissions (GtC per year) when the policymaker has perfect information (dashed orange) and when the policymaker is uncertain but learns about the climate sensitivity (solid blue). The panels correspond to a true climate sensitivity of 2°C (top left), 3°C (top right), 4°C (bottom left), and 5°C (bottom right). We sample 1000 Brownian bridges to compute the expected posterior and evaluate decadal temperature shocks at their mean.

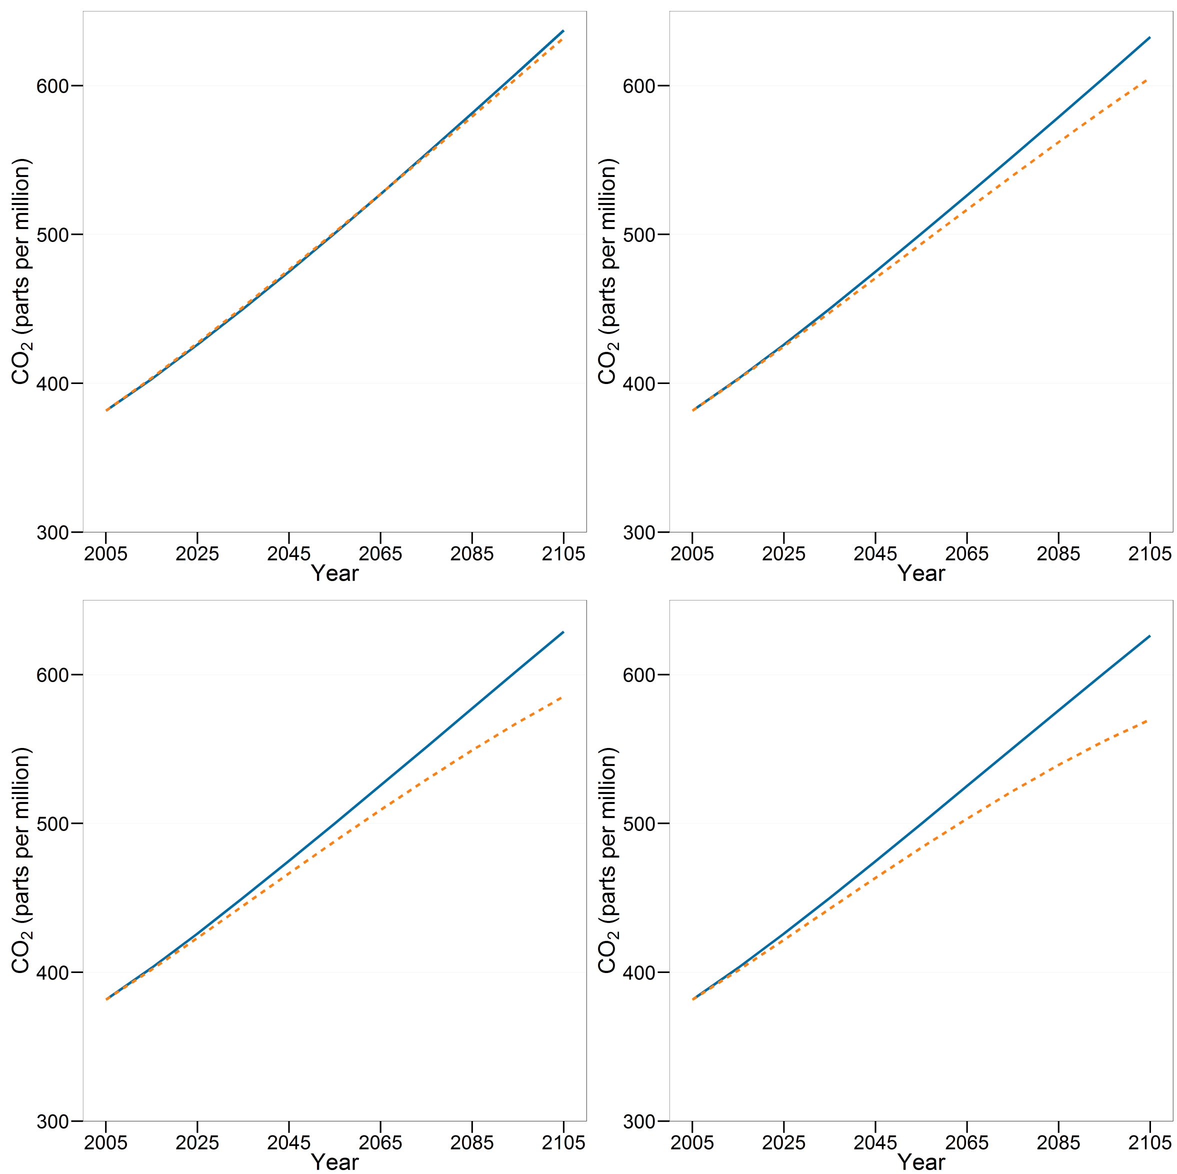

Figure A-4: Evolution of CO_2 concentration (parts per million) when the policymaker has perfect information (dashed orange) and when the policymaker is uncertain but learns about the climate sensitivity (solid blue). The panels correspond to a true climate sensitivity of 2°C (top left), 3°C (top right), 4°C (bottom left), and 5°C (bottom right). We sample 1000 Brownian bridges to compute the expected posterior and evaluate decadal temperature shocks at their mean.

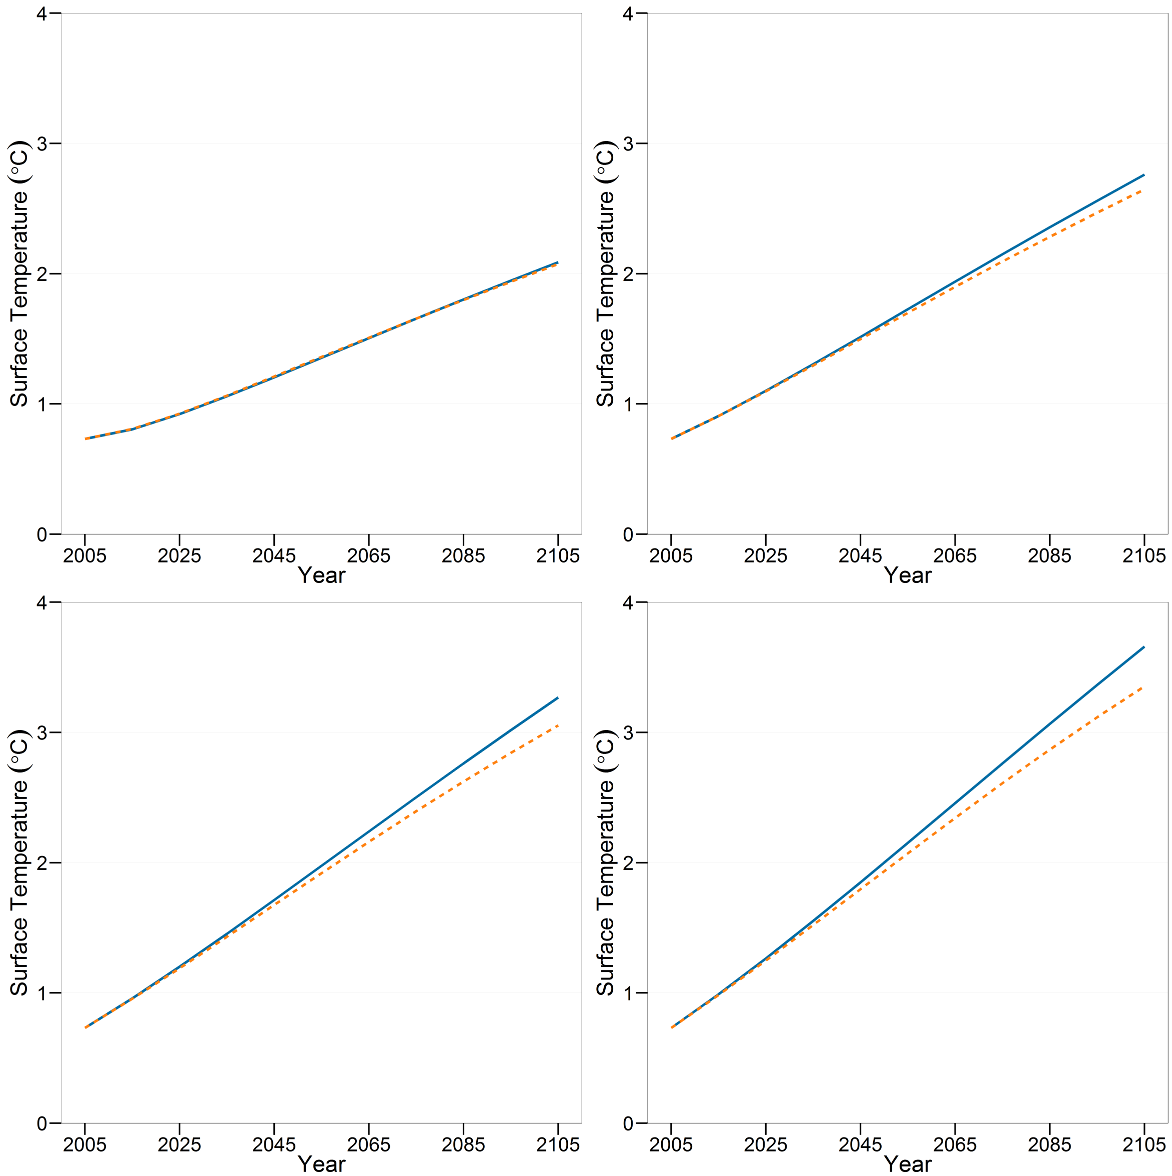

Figure A-5: Evolution of surface temperature (°C) when the policymaker has perfect information (dashed orange) and when the policymaker is uncertain but learns about the climate sensitivity (solid blue). The panels correspond to a true climate sensitivity of 2°C (top left), 3°C (top right), 4°C (bottom left), and 5°C (bottom right). We sample 1000 Brownian bridges to compute the expected posterior and evaluate decadal temperature shocks at their mean.

process constrained such that $W(T) = b$ at some end time T, here at the end of a decade, where $b \in \mathbb{R}$. The drift function in the Brownian bridge is the set of DICE climate equations but with the timestep downscaled to one year. This allows the interpolated CO_2 histories and previous year's interpolated temperatures to influence the interpolation of surface and ocean temperature. The volatility of the bridge is calibrated to the volatility in the historical temperature record, which we estimate using the historical observational record and the downscaled set of DICE climate equations.

The advantage of using the Brownian bridge is two fold. First, a stochastic simulation allows for the inclusion of noise in our interpolation. Fluctuations in annual temperature are consistent with the real world and are necessary to accurately capture the speed of learning. Learning on a deterministic interpolation would be artificially fast since a policymaker would observe a perfectly clear signal of the climate sensitivity. Second, the history observations at a given collocation node may be inconsistent with a deterministic interpolation using the annual version of the DICE temperature equations. The values of the nodes in the history state dimensions are determined by the domain of the state space and where Smolyak's method locates the nodes conditional on a chosen domain. One example of inconsistency between a grid and reality is if historical CO_2 was at extremely high levels, ocean temperature was higher than surface temperature every decade, but surface temperature was declining according to the history states. The CO_2 history and the warm ocean relative to the surface suggests that the surface should be warming, and a simple deterministic interpolation between decades would overshoot the end-of-decade surface temperature value. To reconcile this, the system requires volatility so that a sequence of negative surface temperature shocks can bring temperature back down to the end-of-decade value determined by the grid.[39] A Brownian bridge allows us to introduce volatility in the system that can explain why a deterministic DICE climate system does not fit the history data.

The Brownian bridge works as follows. We construct an annual drift function, $D(T_t, O_t, M_t, t) : \mathbb{R}^4 \to \mathbb{R}^2$, from downscaled equations of the DICE climate system. The DICE climate system is given by,

$$
\begin{aligned}
T_{t+10} &= T_t + C1\left[F(M_{t+10}, t) - \frac{f}{s}T_t + C3\,(O_t - T_t)\right] \\
O_{t+10} &= O_t + C4\,(T_t - O_t)
\end{aligned}
$$

[39] For example, the stochasticity may represent a La Niña event.

The downscaled annual transitions are then,[40]

$$D(T_t, O_t, M_t, t) \equiv \begin{bmatrix} T_{t+1} \\ O_{t+1} \end{bmatrix} =$$

$$\begin{bmatrix} \left(1 - C1\frac{f}{s}\right)^{1/10} T_t + \left(1 - \left(1 - C1\frac{f}{s}\right)^{1/10}\right) \left(\frac{s}{f}\right) F(t+1, M_{t+1}) - \sigma_{ocean}\left(T_t - O_t\right) \\ (1-\sigma_4)O_t + \sigma_4 T_t \end{bmatrix}$$

where $\sigma_4 = 1 - (1 - C4)^{1/10}$, $\sigma_{ocean} = 1 - (1 - C1 \cdot C3)^{1/10}$. Next, let $W(t)$ be a Wiener process from time 0 to time 10 and let $B(t)$ be a Brownian bridge from time 0 to time 10 starting at a value of $\begin{bmatrix} T_0 \\ O_0 \end{bmatrix}$ and ending at $\begin{bmatrix} T_T \\ O_T \end{bmatrix}$. $B(t)$ is governed by the following equation,

$$B(t) = D(T_t, O_t, M_t, t) + \sigma_B W(t) - \frac{t}{10}\left[W(T) - (b - D(T_t, O_t, M_t, t))\right],$$

where σ_B is the volatility of the Wiener process. As was described in the main text, the volatility of the Wiener process was estimated from the observational climate record.

A.4 Learning

The observational record for atmospheric CO_2 and surface temperature goes back to 1880, while the record for ocean temperature only goes back to 1955. We estimate the distribution of missing ocean temperature observations from 1880-1954 as a part of the learning procedure. However, this introduces many more parameters and has a significant computational cost. In the main text, we insert the mean of the missing observation estimates as actual observations to reduce the computational burden and find this has nearly no effect on the posterior distribution of climate sensitivity.

Last, we estimate the variance of the stochasticity of surface temperature and ocean temperature using the observational records. As with the missing ocean observations, we use the mean of the variance estimates as fixed values and find that this approximation also has little effect on posterior estimates even as it reduces the number of parameters we must estimate to only a single parameter (i.e., to the climate sensitivity parameter).

A.5 Hamiltonian Monte Carlo

Even if the prior $p(\theta)$ has a closed form, the proportionality constant in Bayes' rule

$$\int p(S_t, H_t|\theta)p(\theta)\mathrm{d}\theta$$

[40]Note that this is not an exact mapping due to the nonstationarity of the temperature transitions (Traeger, 2014).

may be difficult to compute so the posterior $p(\theta|S_t, H_t)$ will not be guaranteed to have a closed form. We use a Hamiltonian Monte Carlo (HMC) algorithm to sample the posterior and construct a numerical approximation to the posterior (Neal, 2011).[41] Here we describe the basic HMC algorithm and how it maps into Bayesian inference.

HMC is rooted in Hamiltonian dynamics, a formalization of the deterministic motion of molecules. Before we demonstrate how Hamiltonian dynamics maps into Bayesian inference, we must describe Hamiltonian dynamics. First, we define the Hamiltonian $H(\theta, p)$ of our system as

$$H(\theta, p) \equiv U(\theta) + K(p),$$

where θ is a position vector, p is a momentum vector, $U(\theta)$ is potential energy, and $K(p)$ is kinetic energy. θ is our variable of interest in HMC, and p is an auxiliary variable we use to facilitate proper sampling of our posterior density. When performing inference, $U(\theta)$ will be the negative log density of the random variable θ from which we wish to sample. We proceed with a description of a one-dimensional system since we only are sampling from a univariate density over the climate sensitivity.

In one dimension, the two energy functions are:

$$U(\theta) = \frac{1}{2}\theta^2, \qquad K(p) = \frac{1}{2m}p^2,$$

where m is a constant. The quadratic form of the kinetic energy function implies that the negative log probability density of p is normal with zero mean and variance m. The partial derivatives of the Hamiltonian tell us how the position and momentum of the molecule change over time,

$$\frac{\partial \theta}{dt} = \frac{\partial H(\theta, p)}{\partial p}, \qquad \frac{\partial p}{dt} = -\frac{\partial H(\theta, p)}{\partial \theta}.$$

We will simulate a discrete approximation to the system forward in time, which will yield us our new proposed position (climate sensitivity) for sampling. We discretize the simulation by using the *leapfrog* method.[42] In leapfrog (and in other Euler approaches), we define a stepsize ϵ for how far we move forward in time at each step of the simulation. Beginning at some time t, we determine the new position and momentum at time $t+\epsilon$ by first determining momentum at time $t + \epsilon/2$,

$$p(t + \epsilon/2) = p(t) + \epsilon/2 \frac{\partial U(q(t))}{\partial q(t)}.$$

Then, using this new momentum value, we determine the position a full step ahead:

$$\theta(t + \epsilon) = \theta(t) + \epsilon \frac{p(t + \epsilon/2)}{m}.$$

[41] Specifically, we employ the No U-Turn Sampler (NUTS), a more computationally efficient version of the standard HMC algorithm (Hoffman and Gelman, 2014).

[42] Leapfrog has better stability in approximation than standard Euler approaches.

Finally, we move momentum another half step to be at time $t + \epsilon$,

$$p(t + \epsilon) = p(t + \epsilon/2) + \epsilon/2 \frac{\partial U(q(t + \epsilon))}{\partial q(t + \epsilon)}.$$

We will perform this simulation a total of L steps.

A.5.1 The HMC Algorithm

We can relate the target climate sensitivity posterior to the potential energy function using the *canonical distribution*:

$$P(x) = \frac{1}{z} exp \left(-\frac{U(\theta)}{T} \right),$$

where z is a normalizing constant and T is the temperature of the system.[43] In our case, the climate sensitivity is the parameter of interest and is equivalent to the position variable in the Hamiltonian system. The posterior of the climate sensitivity can be put in the form of a canonical distribution if we define the potential energy function as

$$U(\theta) = -log\left(p(\theta)\, p(D|\theta)\right).$$

We set $T = 1$.

Using this form of the canonical distribution, we can proceed with the HMC algorithm which has two main steps. In the first step we draw a value for the momentum variable from its normal distribution, $p \sim N(0, m)$. Note that this value is independent of the current position. In the second step we perform a Metropolis update using Hamiltonian dynamics to propose a new state (θ, p). Starting from (θ, p), we simulate the Hamiltonian system forward in time L steps with a stepsize of ϵ using the leapfrog method. Once we have moved forward L timesteps, we negate the momentum variable to obtain our proposed state (θ^*, p^*).[44] We accept this as our next state (and sample θ^*) with probability $\min\{1, \exp(-H(\theta^*, p^*) + H(\theta, p))\}$. If the proposal is accepted, then (θ^*, p^*) is taken to be the new state. If the proposal is rejected, then the next state remains (θ, p). These two steps are repeated until a desired number of samples are obtained.

A.6 Approximation of Beliefs

We approximate our posterior distributions at each collocation node p in some decade d using a nonparametric kernel smoother and Gauss-Legendre quadrature. We use quadrature since directly using the HMC samples to take expectations would result in orders of magnitude more evaluations of the continuation value when maximizing the Bellman equation.

[43] T here does not relate to temperature in the actual climate but of the Hamiltonian system.

[44] The negation preserves symmetry in the proposal for the Metropolis update.

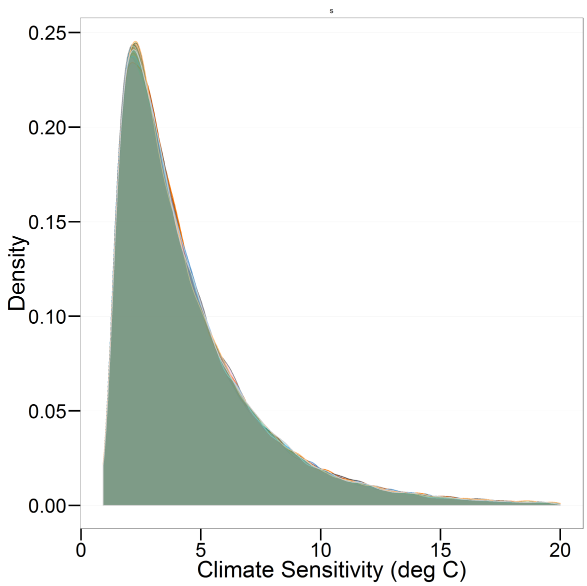

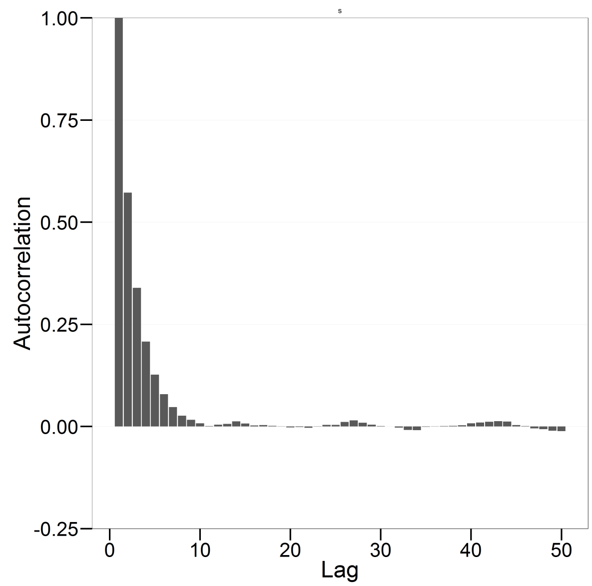

Figure A-6: Left panel: Ten different expected posterior densities using the same observational record data. Right panel: The autocorrelation between our post-burn in samples. The size of the post-burn in sample is 52,500 draws with a burn in of 2,500 draws.

To obtain the posterior approximation, we perform the interpolation and learning procedure N=1,000 times for each node p.[45] One Brownian bridge introduces randomness into the learning procedure, so we perform N interpolations and then use these stochastic interpolations obtain N sets of posterior samples. Each posterior sample contains 50,000 draws after dropping the first 2,500 for burn in.[46] We then combine all N samples into one large sample to obtain an *expected posterior* that marginalizes out the noise from the Brownian bridge.[47] After combining the draws from all N posterior samples, we use a Gaussian kernel smoother to generate a nonparametric density for the climate sensitivity. Next, we evaluate the smoothed density at a set of Gauss-Legendre quadrature nodes to get the density values for taking expectations.[48] We select the bandwidth for the kernel smoother as follows. We estimate the posterior distribution for the climate sensitivity using only the observational

[45] We selected N=1,000 by increasing N until additional Brownian bridges had negligible effects on beliefs and optimal policy.

[46] Extending the burn in period does not meaningfully reduce autocorrelation in the samples.

[47] Standard conjugate updating rules also yield an expected posterior, as the policymaker forms expectations about the future values of her belief parameters.

[48] The Gauss-Legendre quadrature weights are predetermined for integrating some function of the climate sensitivity s like $\int V(s)ds = \sum_{i=1}^{N} w(i)\,V(i)$. However we are aiming to take an expectations of a function, the continuation value, with respect to the density of the climate sensitivity $f(s)$: $\int V(s)\,dF(s)$. We can approximate this expectation using the following quadrature scheme $\sum_{i=1}^{N} w(i)\,f(i)\,V(i)$.

record 1,000 times where each estimate consists of 10,000 samples. Then we compute the mean climate sensitivity estimate for each of the 1000 sets of samples. Our selected bandwidth h^* for the Gaussian kernel minimizes the average absolute difference between the mean climate sensitivity estimates given by the quadrature scheme and the actual sample from the posterior across all 1,000 sets of samples.

The left panel of Figure A-6 displays nonparametric estimates of the climate sensitivity posterior density for 10 different HMC chains when using the observational record. The plot shows that the 10 nonparametric posterior densities are extremely close together, indicating that a chain length of 52,500 can closely approximate the posterior. The right panel of Figure A-6 displays the autocorrelation in the HMC between a given draw and a lagged draw. The autocorrelation drops to zero rapidly, which suggests that our sample is approximately drawn independently from the target distribution.[49]

Next we use value function iteration to approximate our value functions for each decade conditional on the approximated climate sensitivity distributions for each node in that decade. The value function algorithm is as follows,

1. At the terminal decade D, solve the Bellman equation at all collocation nodes. The continuation value function, the value function corresponding to time $D+1$, is a predetermined terminal value function.

2. Using the maximized values, solve a linear system of equations to recover the coefficients for polynomial approximating the value function corresponding to decade D.

3. For each decade $D-1$ to 0:

   (a) Solve the Bellman equation at all nodes.

   (b) Using the maximized values, solve a linear system of equations to recover the coefficients for the value function approximating the value function for that decade.

[49] Note that the autocorrelation plot starts with a lag of 0, so the first autocorrelation value is by definition 1.

A.7 Pseudocode for the Full Algorithm

In this section we provide pseudocode for the full algorithm. The algorithm is as follows,

```
BELIEF PRECOMPUTATION

For each decade from terminal time to initial time

    For each collocation node
      1. Linearly interpolate the node's CO2 history vector
      2. Stochastically interpolate the node's temperature history vector
         1,000 times with i.i.d. Brownian bridges
      3. Sample the 1,000 posteriors using a 52,500-length HMC chain, with
         the first 2,500 samples used for burn-in
      4. Combine the samples to generate an expected posterior across bridges
      5. Estimate a nonparametric density by applying a Gaussian kernel smoother
         to the combined sample
      6. Evaluate the nonparametric density at the Gauss-Legendre quadrature
         nodes to recover the weight vector for taking expectations
      7. Save the weight vector
    end

end

VALUE FUNCTION ITERATION

For each decade from terminal time to initial time

    For each collocation node
      1. Maximize the right hand side of the Bellman using the precomputed
         nonparametric density to take expectations over the one-decade-ahead
         value function approximant
    end

  1. Project the maximized Bellman values onto a value function basis

end
```

A.8 DICE Equations

Here we describe the DICE model. The version of DICE we implement follows that of Lemoine and Rudik (2017) except for here we have nonconjugate learning and we also have ocean temperature evolving stochastically. DICE is a Ramsey growth model with a climate system coupled to it. DICE has a representative policymaker who aims to maximize her stream of expected discounted utility from consumption. Each period t, the policymaker begins with a stock of capital K_t, an exogenous amount of labor L_t, and exogenous technology A_t. She combines these in a Cobb-Douglas production function to generate gross output, $Y_t^g = A_t L_t^{1-\kappa} K_t^{\kappa}$. In each period, the surface temperature of the Earth is given by T_t^s, and higher surface temperature damages the economic and reduces the amount of output from factor production in the following way

$$Y_t^n = \frac{Y_t^g}{1 + b_2 \, T_t^{s\,b_3}}.$$

The policymaker can use her net output after damages Y_t^n three different ways. The first is for consumption to increase utility where her utility function is the standard constant relative risk aversion form form

$$u(c_t; L_t) = \frac{(c_t/L_t)^{1-\eta}}{1-\eta}.$$

Second, she can invest it to increase her future capital stock. Last, she can use it to abate industrial CO_2 emissions. The net amount of emissions entering the atmospheric CO_2 stock M_t^{atm} each period is

$$e_t = 10 \left[\sigma_t (1-\alpha_t) Y_t^g + B_t\right],$$

where σ_t is the emissions intensity of production, α_t is the abatement rate, and B_t captures exogenous emissions. CO_2 in the atmosphere can be exchanged with the land/upper ocean M_t^{up} and the lower ocean M_t^{lo} according to a linear three box carbon model,

$$\begin{bmatrix} M_{t+10}^{atm} \\ M_{t+10}^{up} \\ M_{t+10}^{lo} \end{bmatrix} = \begin{bmatrix} \phi_{11} & \phi_{21} & 0 \\ \phi_{12} & \phi_{22} & \phi_{32} \\ 0 & \phi_{23} & \phi_{33} \end{bmatrix} \begin{bmatrix} M_t^{atm} \\ M_t^{up} \\ M_t^{lo} \end{bmatrix} + \begin{bmatrix} e_t \\ 0 \\ 0 \end{bmatrix}$$

Radiative forcing is the amount of heat trapped by by atmospheric CO_2

$$F_t(M_t^{atm}, t) = f_{2x} log_2 \left(M_t^{atm}/M_{pre}\right) + EF_t,$$

where M_{pre} is the preindustrial level of atmospheric CO_2, and EF_t is the amount of forcing from exogenous sources. Additional radiative forcing increases surface temperature which is given by

$$T_{t+10}^s = T_t^s + C_1 \left[F_{t+10}(M_{t+10}^{atm}, t+10) - \frac{f_{2x}}{s} T_t^s + C_3 (T_t^o - T_t^s)\right] + \epsilon_t^s,$$

where s is the uncertain climate sensitivity, T_t^o is the ocean temperature, and ϵ_t^s is the surface temperature volatility term. Ocean temperature is given by

$$T_{t+10}^o = C_4 T_t^o + (1 - C_4) T_t^s + \epsilon_t^o,$$

where ϵ_t^o is the ocean temperature volatility term.

Table 1: The model parameters.

Parameter	Value	Description
A_{2005}	0.027	Initial production technology
$g_{A,2005}$	0.009	Initial growth rate of production technology
δ_A	0.001	Change in growth rate of production technology
L_{2005}	6514	Initial population (millions)
L_∞	8600	Asymptotic population (millions)
δ_L	0.35	Rate of approach to asymptotic population level
σ_{2005}	0.13	Initial emission intensity of output (Gigatons of carbon per unit output)
$g_{\sigma,2005}$	-0.073	Initial growth rate of decarbonization
δ_σ	0.003	Change in growth rate of emissions intensity
a_{2005}	1.17	Initial cost of backstop technology (\$1000 per ton of carbon)
a_1	2	Ratio of initial backstop technology cost to final backstop technology cost
a_2	2.8	Abatement cost function exponent
g_Ψ	0.05	Growth rate of backstop technology cost
B_{2005}	1.1	Initial non-industrial CO_2 emissions (Gigatons of carbon)
g_B	0.9	Growth rate of non-industrial emissions
b_2	0.0028	Damage coefficient
b_3	2	Damage exponent
EF_{2005}	-0.06	Year 2005 exogenous forcing (W/m^2)
EF_{2105}	0.30	Year 2105 exogenous forcing (W/m^2)
κ	0.3	Capital elasticity in production
δ_κ	0.1	Annual capital depreciation rate
M_{pre}	596.4	Pre-industrial atmospheric CO_2 (Gigatons of carbon)
β	$1/1.015^{10}$	Discount factor
η	2	Relative risk aversion and relative inequality aversion
ϕ_{11}	0.811	Carbon transfer coefficient for atmosphere to atmosphere
ϕ_{12}	0.189	Carbon transfer coefficient for atmosphere to upper ocean
ϕ_{21}	0.097	Carbon transfer coefficient for upper ocean to atmosphere
ϕ_{22}	0.853	Carbon transfer coefficient for upper ocean to upper ocean
ϕ_{23}	0.050	Carbon transfer coefficient for upper ocean to lower ocean
ϕ_{32}	0.003	Carbon transfer coefficient for lower ocean to upper ocean
ϕ_{33}	0.997	Carbon transfer coefficient for lower ocean to lower ocean
C_1	0.22	Warming delay parameter

Continued on next page

Table 1 – continued from previous page

Parameter	Value	Description
C_3	0.3	Parameter governing transfer of heat from ocean to surface
C_4	0.05	Parameter governing transfer of heat from surface to ocean
f_{2x}	3.8	Forcing from doubling of CO_2 (W/m^2)
λ_0	1.2	Reference system's climate sensitivity (°C/[W/m^2])
σ_T^2	0.2^2	Surface temperature shock variance
σ_O^2	0.015^2	Ocean temperature shock variance
K_0	137	Year 2005 capital (trillions of USD)
M_{2005}^{atm}	808.9	Year 2005 atmospheric CO_2 (Gigatons of carbon)
M_{2005}^{up}	1255	Year 2005 biosphere and upper ocean CO_2 (Gigatons of carbon)
M_{2005}^{lo}	18365	Year 2005 lower ocean CO_2 (Gigatons of carbon)
T_{2005}^{atm}	0.7307	Year 2005 atmospheric temperature (°C)
T_{2005}^{ocean}	0.0068	Year 2005 ocean temperature (°C)
μ_{2005}	1.28	Location parameter for year 2005 lognormal climate sensitivity distribution
Σ_{2005}	0.602	Scale parameter for year 2005 lognormal climate sensitivity distribution
μ_{1880}	1.73	Location parameter for year 1880 lognormal climate sensitivity distribution
Σ_{1880}	0.604	Scale parameter for year 1880 lognormal climate sensitivity distribution
$\underline{s}$	1.2	Climate sensitivity distribution lower bound[50] (°C/[W/m^2])
$\bar{s}$	20	Climate sensitivity distribution upper bound (°C/[W/m^2])

A.9 Error Analysis

Here we explore error in our model. We do this in two ways. The first is to compare deterministic results obtained from our dynamic programming approach to results obtained from a highly accurate optimal control approach (Cai et al., 2015). We compare trajectories for our control variables when solving a deterministic version of DICE using the exact same collocation grid we use for our main results against trajectories obtained when solving the same model but with using an optimal control approach with a highly accurate non-linear solver as in Nordhaus (2008). If our dynamic programming solution is accurate, we should be able to closely replicate policy solutions from an alternative solution method that directly solves for policy instead of the value function.

Figure A-7 plots abatement and consumption trajectories for deterministic model runs with climate sensitivity values of 2, 3, 4, and 5. The black lines are the trajectories from our dynamic programming solution and the gray lines are trajectories from the optimal control solution using modified Nordhaus (2008) code.[51] The left panel displays abatement

[50]This corresponds to nonnegative total feedback.

[51]We must modify the code to make the problem Markov. This results in only minor differences from the original Nordhaus (2008) results.

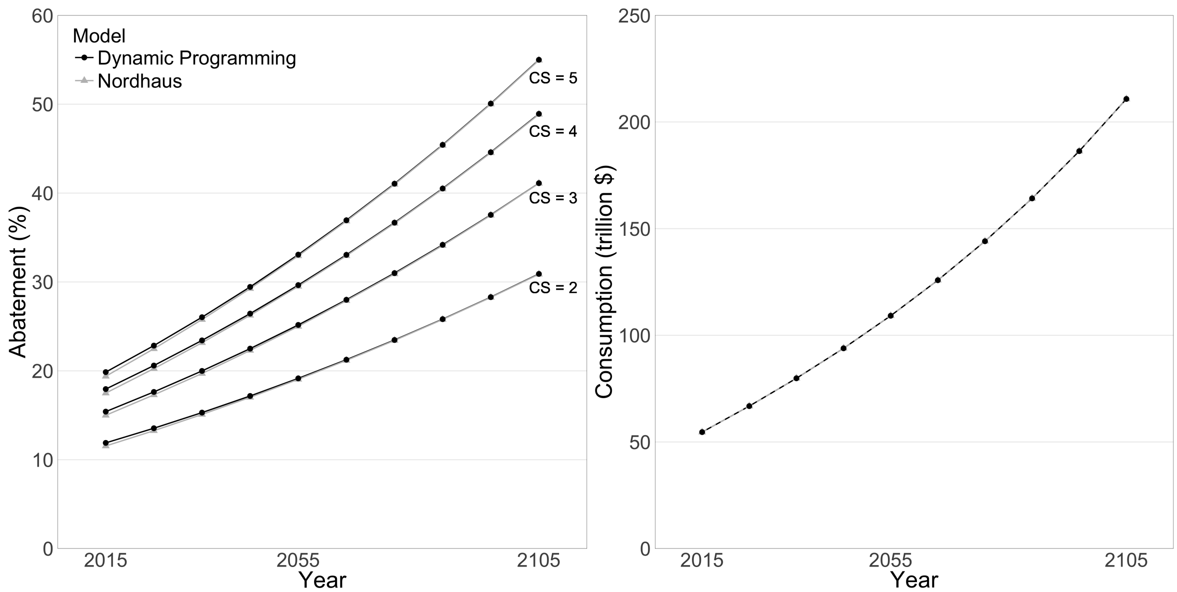

Figure A-7: Abatement and consumption trajectories for a deterministic DICE model with climate sensitivity values of 2, 3, 4, and 5. Black lines are solutions when using a dynamic programming approach with the same collocation grid as our main results. Gray lines are solutions when using the same optimal control approach as Nordhaus (2008). We only show one set of trajectories for consumption since consumption is insensitive to climate sensitivity and all trajectories are very similar.

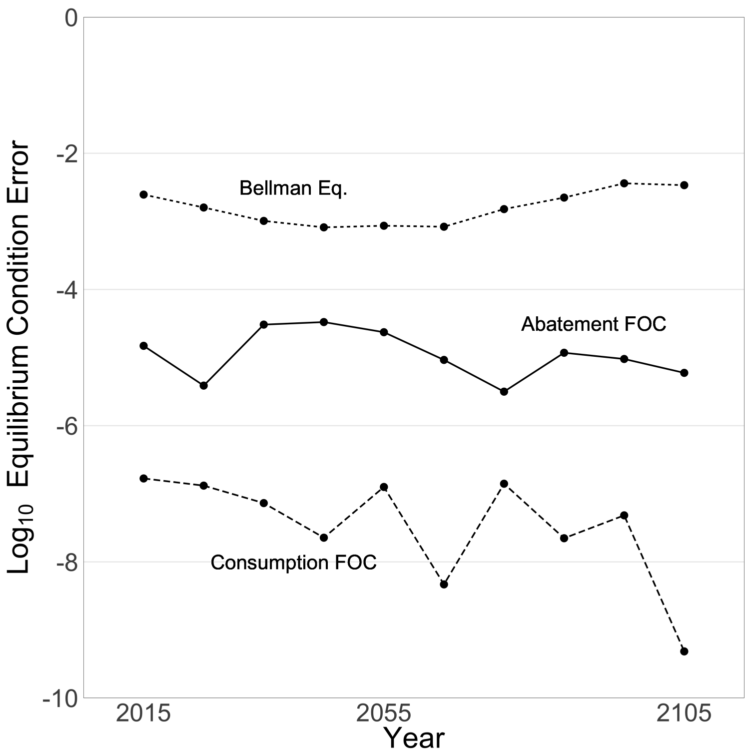

Figure A-8: Relative equilibrium condition errors along the optimal trajectory in $\log_{10}$ units when the true climate sensitivity value is 3. Larger negative numbers imply less error in the equilibrium condition.

trajectories. In early years, there is small amounts of error in the abatement policy, but this quickly approaches zero. The right panel displays one set of consumption trajectories for a climate sensitivity value of 3, we omit the other climate sensitivity values for clarity since they are all extremely similar. The solutions are nearly identical across the two solution methods, and any differences between the dynamic programming approach and optimal control approach are extremely small regardless of the climate sensitivity value.

Another test of the accuracy of our results is to explore error in the equilibrium conditions of the model. If we have accurate policies then the error in our first-order conditions should be small, conditional on our value function approximation. Similarly, if our approximation to the true value function is accurate then error in the Bellman equation should be small, conditional on our policies being chosen accurately.[52] Figure A-8 plots the relative error in the first-order conditions and the Bellman equation in $\log_{10}$ units along the optimal trajectory for a true climate sensitivity value of 3. Results are insensitive to which true climate sensitivity

[52]Collocation approaches force error in the Bellman to be zero on the collocation grid. This is not guaranteed off the grid.

value we select. The errors are relative in the sense that we rearrange all the equilibrium conditions so that there is a collection of terms on the left hand side, and 1 on the right hand side. We then subtract 1 from both sides, take the absolute value, and then take the base 10 logarithm. In the limit that error in the conditions goes to zero, the logarithm will approach $-\infty$. For any non-zero amount of error, it will return a finite number where numbers that are more negative imply greater accuracy. We can see that error is small for all equilibrium conditions and within the range reported in the existing numerical literature. **GET CITES, JMMV, OTHERS?**