Decarbonization Investment Strategies in an Uncertain Climate

Adam Michael Bauer*[1,2], Florent McIsaac[2], and Stéphane Hallegatte[2]

[1] *Department of Physics, University of Illinois Urbana-Champaign, 1110 W Green St Loomis Laboratory, Urbana, IL 61801*
[2] *World Bank Group, Washington DC*

Forthcoming in *Earth's Future*
4/8/2025

Abstract

The Paris Agreement established that global warming should be limited to "well below" 2 °C and encouraged efforts to limit warming to 1.5 °C. Achieving this goal presents a challenge, especially given (i) adjustment costs, which penalize a swift transition away from fossil fuels owing to, e.g., skilled labor scarcity, and (ii) climate uncertainty that complicates the link between emissions reductions and global warming. This paper presents a modeling framework that explores optimal decarbonization investment strategies with adjustment costs and climate uncertainty. The findings show that climate uncertainty impacts investment in three ways: (i) the cost of policy increases, especially when adjustment costs are present; (ii) abatement investment is front-loaded relative to a scenario without uncertainty; and (iii) the sectors with the largest changes in investment are those that are "hard-to-abate", such as heavy industry and agriculture, each of which have high investment costs and annual emission rates. The longer learning about climate uncertainty is delayed, the more these impacts are amplified. Each of these effects can be traced back to the carbon price distribution inheriting a "heavy tail" when climate uncertainty is present. The paper highlights how climate uncertainty and adjustment costs combined lead to heightened urgency for near-term investments in decarbonization.

Plain Language Summary: Achieving international climate goals will require transitioning the global economy from greenhouse gas-emitting technologies to clean technologies. How one allocates effort towards this transition over time is complicated: going "too fast" runs the risk of imposing extra costs – often referred to as "adjustment costs" – related to, say, a scarcity of electricians that are trained to install solar panels (here, the cost of training new electricians would be an adjustment cost). Going "too slow", however, risks missing climate targets altogether. Further complicating this process is uncertainty in how the climate responds to carbon emissions, which limits policymakers' ability to know how much time they have to complete the transition. To understand how adjustment costs interact with climate uncertainty in the green transition, this paper presents a model for financing the least-cost green transition under climate uncertainty. The paper shows how, to hedge against worst-case climate

*Corresponding author: adammb4@illinois.edu

outcomes, investment efforts need to be front-loaded, with more early investments in expensive-to-abate sectors like heavy industry and agriculture. The paper discusses how climate uncertainty increases the carbon price needed to meet climate goals, raising the profitability of investments in green technologies. Overall, the paper points towards climate uncertainty increasing the urgency of climate policies.

Key Points:

- The optimal policy under climate uncertainty leads to more early investments to hedge against worst-case outcomes.
- Adjustment costs slow the rate of abatement (via, e.g., labor scarcity), thus increasing near-term investment needs to meet climate goals.
- The interaction of climate uncertainty and adjustment costs leads to higher carbon prices and larger near-term investment needs.

1 Introduction

The signing of the Paris Agreement marked a breakthrough commitment and consensus in the international community about the serious threat of human-caused climate change (United Nations Framework Convention on Climate Change, 2015). Coincident with this watershed moment in global climate policy arose an equally daunting challenge: confining planetary warming to the targets enshrined in the Paris Agreement, namely, "well-below" 2 °C with effort towards limiting warming to 1.5 °C. These goals are ambitious, and present a major challenge to the international community. For example, a recent estimate suggests that while gross investments in climate financing amassed to about 1% of global GDP – about $1.3 trillion USD – in 2023, in order to reach the Paris Agreement targets, climate finance must increase by at least five-fold annually relative to current levels (Climate Policy Initiative, 2023).

Clearly, transitioning from an economy based on "dirty" capital, such as coal-fired power plants, to "clean" capital, such as renewable energy and electricity-powered vehicles, will require significant investment in abatement. A number of factors exist that complicate this process.

For example, adjustment costs (Lucas, 1967; Mussa, 1977) considerably raise the price tag of a swift transition to a greenhouse gas emissions-free economy. This is because an induced transition from dirty to clean capital creates an opportunity cost to direct scarce resources, such as skilled labor or productive capacity, from polluting production to greenhouse gas-abating economic activity and tends to slow the diffusion of new technologies (Morris et al., 2019). As an example, if one was to transition a region's electrical grid from relying on natural gas to primarily relying on renewables with battery storage, completing this transition in one year would be far more expensive than if it was done over the course of a decade. In this example, the "adjustment cost" might be the cost of training additional electricians to install the required solar panels, or the cost of buying more expensive batteries to accommodate the speed of the transition. Additional adjustment costs could arise from supply chain and trade constraints that place limits on the rate certain technologies can be bought, imported, and deployed. Models including adjustment costs tend to favor more near-term investments in clean technologies (Ha-Duong et al., 1997; Campiglio et al., 2022) conflicting with a number of models in the literature that do not consider these effects (e.g., Nordhaus (2017)) and recommend an initially low, rising investment pathway.

Another complicating factor is that there are a number of options one can choose to abate fossil fuel emissions; how does one choose between abating emissions in the energy sector, say, versus heavy industry? And how should these efforts be allocated over time? Conventional marginal abatement cost-based approaches would suggest that one should start with cheap mitigation options and progressively move towards more expensive technologies; however, this policy advice can be challenged if one includes the impact of adjustment costs, which has been shown to lead to more investment happening earlier in expensive-to-abate sectors (Vogt-Schilb and Hallegatte, 2014). Throughout, we will refer to sectors with both high marginal abatement costs and high emissions intensities as "hard-to-abate" (e.g., heavy industry and agriculture).

The challenge to financing decarbonization posed by each of these economic factors is compounded by the presence of uncertainty in the physical climate system. While the targets in the Paris Agreement are deterministic, in the sense that the targets themselves are not uncertain, the geophysical timing of their actualization is unclear owing to uncertain climate feedbacks (Sherwood et al., 2020). For example, the true value of the remaining carbon budget (or simply "the carbon budget"), a geophysical quantity that corresponds to the amount of emissions one has left to emit before a given long-term global temperature target is nearly certain to be reached, is ambiguous owing to uncertainty in the zero-emissions commitment, future aerosol emissions, and the transient climate response to emissions (Matthews et al., 2009, 2018, 2021; Jenkins et al., 2022). The presence of climate uncertainty has been generally shown to increase the stringency of climate policy (see Lemoine and Rudik (2017), Lemoine (2021), Cai (2021), and Bauer et al. (2024b) for a few examples) and imply more climate damages (Calel et al., 2020); this is especially the case if so-called "climate tipping points" are considered (Lenton et al., 2008; Lemoine and Traeger, 2016; Cai and Lontzek, 2019; Dietz et al., 2021).

In this paper, we amend an economic model of abatement investment that includes convex adjustment costs and heterogenous sectors (Vogt-Schilb et al., 2018) with a representation of climate uncertainty to jointly explore the influence of adjustment costs and climate uncertainty on optimal decarbonization investment strategies. In particular, we task the social planner with decarbonizing the economy prior to breaching some temperature target for the least cost when the true value of the carbon budget is hidden until some point in time; once the carbon budget is known, the planner's policy can be adjusted (either to be more stringent or lax) in accordance with the true value. The result is that the social planner has to formulate a policy which is robust under a number of potential future risk states that will be revealed later on (similar to the approach of, e.g., Ackerman et al. (2013), Crost and Traeger (2014), Morris et al. (2018), and Okullo (2020)). Our approach allows us to vary the time when the carbon budget is learned, and analyze how learning the true value of the carbon budget in 2030, for example, impacts the resulting policy in comparison to learning about the carbon budget in 2080 or later. We compare the results of this experiment with a "strawman" model, where the effects of adjustment costs are neglected, to isolate how climate uncertainty interacts with adjustment costs.

We find that the effect of climate uncertainty has three main impacts on optimal decarbonization investment strategies: (i) the presence of adjustment costs magnifies the impact of climate uncertainty, especially for stringent temperature targets, (ii) abatement investment is front-loaded relative to an equivalent policy without uncertainty, and (iii) sectoral allocation of abatement investment is affected to prioritize hard-to-abate sectors relative to the certainty policy. We further find that if learning about climate uncertainty is delayed, each of these effects is magnified because the planner has to formulate a policy that is robust to the worst-case climate outcome (i.e., a scenario with high climate sensitivity or, equivalently, a low carbon budget) for longer. Our findings suggest that uncertainty and adjustment costs interact to exacerbate the influence of uncertainty, and overall advocate for a more aggressive least-cost investment strategy for decarbonization. We find these results are robust to a number of alternative model specifications, such as allowing for the deployment of direct air capture technologies, different emissions baselines, and alternative cost of investment calibrations (Sections 1, 2, and 3 of the *Supplementary Information*, respectively).

2 Model Description

In this section, we outline each of our model frameworks. In Section 2.1, we define the base models: one that uses abatement investment and captures adjustment costs, and one that does not. In Section 2.2, we add climate uncertainty to the base models. Note the solutions to the models highlighted in Section 2.1 are the certainty solutions of the uncertain models in Section 2.2. In Section 2.3, we describe our general calibration routine.

2.1 Base models

2.1.1 Abatement investment model

The abatement investment model tasks a policymaker with decarbonizing a set of economic sectors, $\mathcal{I}$, for the least cost over some time horizon, $\mathcal{T} := [0, t_f]$, such that the carbon budget, $B > 0$, is not exceeded. Each sector has an abatement potential (or annual emissions rate), $\bar{a}_i > 0$, and a capital depreciation rate, $0 < \delta_i \leq 1$. No sector has the capacity for "negative emissions"; we explore the influence of "negative emissions" technologies on our results in Section 1 of the *Supplementary Information*, and find that none of the results are significantly altered.

We represent abatement of greenhouse gases via the accumulation of abatement capital stocks. Here, abatement capital stocks represent any abatement technologies that, once invested in, reduce greenhouse gas emissions over their lifetime. In this way, our representation of abatement capital stocks follows the "committed emissions" framework originally developed by Davis and Socolow (2014). As an example, if one builds a green steel plant with a capital lifetime of 30 years which replaces a dirty steel plant that emitted 1 $MtCO_2$ per year, then the green steel plant represents an abatement investment of 1 $MtCO_2$ per year, per year the green plant is in operation. The task of the social planner is therefore to build up sufficient abatement capital stocks over the investment horizon (while also replacing capital as it depreciates) such that the economy's greenhouse gas emissions rate reaches zero when the carbon budget is depleted.

Mathematically, the policymaker controls the amount of investment in abatement capital stocks, $x_{i,t}$, in each sector $i \in \mathcal{I}$ at each point in time $t \in \mathcal{T}$; abatement investment leads to the abatement of emissions, $a_{i,t}$, via the accumulation of abatement capital. Global cumulative emissions are given by ψ_t. The cost of abatement investment in each sector $i \in \mathcal{I}$, $c_i(x_{i,t})$, is assumed to be an increasing and convex function of investment (i.e., its first and second derivatives with respect to investment are positive).

The convexity requirement on our cost function endows the system with adjustment costs (Lucas, 1967; Mussa, 1977). Intuitively, if one wants to install a large amount of solar panels in a short amount of time, completing the installations will require training new technicians to make the installations, and retraining them after the installations are complete. These training and re-training costs would not be incurred if installations happened smoothly over time, where the demand for solar panel installations is equal to the supply of available technicians. Here, the training and re-training costs are the "adjustment costs", and demonstrate how the *marginal cost of abatement capital*, $c'(x)$, is a function of the *rate of capital installed*, x; since $c'(x)$ is an increasing function of x, $c(x)$ must be convex, and therefore a convex cost of investment captures adjustment costs.

The problem facing the social planner can therefore be formulated as,

$$\min_{\{x_{i,t}\}_{i\in\mathcal{I}}} \left[\sum_{t\in\mathcal{T}} \beta^t \sum_{i\in\mathcal{I}} c_i(x_{i,t})\right], \tag{2.1}$$

$$\begin{aligned}
\text{Subject to}: a_{i,t+1} &= a_{i,t} + \Delta t\,(x_{i,t} - \delta_i a_{i,t}),\\
\psi_{t+1} &= \psi_t + \Delta t\left(\sum_{i\in\mathcal{I}} (\bar{a}_i - a_{i,t})\right),\\
&0 \le a_{i,t} \le \bar{a}_i,\\
&0 \le \psi_t \le B,\\
&x_{i,t} \ge 0,\\
&a_{i,0}, \psi_0 \text{ given},
\end{aligned}$$

where $\Delta t > 0$ is the timestep (assumed to be annual throughout), $\beta := (1+r)^{-1}$ is the discount factor, and $0 < r \le 1$ is the social discount rate. While an exogenous rate of technology change could be incorporated into the discount factor (via the transformation $\beta = (1+r)^{-1} \to (1-\phi)/(1+r)$ where $\phi \ge 0$ is the exogenous rate of technology change), we ignore it to focus on model dynamics. This model was extensively studied in Vogt-Schilb et al. (2018), and we refer readers to this paper for a full analytical treatment of the model.

Two theoretical features of the model are important for our discussion. First, supposing that t_f is sufficiently large that the economy is decarbonized at some time $\tilde{t} < t_f$, the model will reach a steady state such that $\dot{\psi} = \dot{a}_i = 0$. This implies that for all $t > \tilde{t}$, $a_{i,t} = \bar{a}_i$, $\psi_t = B$, and $x_{i,t} = \delta_i \bar{a}_i$, meaning the steady-state investment effort is given by $c_i(\delta_i \bar{a}_i)$.

The second important result is that, along the optimal path of (2.1), investment is either bell-shaped or decreasing over time (see Proposition 1 in Vogt-Schilb et al. (2018)). The shape of the investment pathway in a given sector relies on the inequality,

$$\mu \gtreqless (r + \delta_i) c_i'(\delta_i \bar{a}_i), \tag{2.2}$$

that weighs the carbon price (given by μ) against the marginal abatement investment costs. Note the carbon price is determined endogenously in the model as the shadow value of emissions reductions (i.e., is the Lagrange dual of the cumulative emissions time derivative, $\dot{\psi}_t := \sum_{i \in \mathcal{I}} (\bar{a}_i - a_{i,t})$). When $\mu > (r + \delta_i) c_i'(\delta_i \bar{a}_i)$, paying the carbon price is more expensive than a marginal unit of abatement investment, and it follows that the optimal investment path is declining. If $\mu < (r + \delta_i) c_i'(\delta_i \bar{a}_i)$, then initially paying the carbon price is more cost-effective than a marginal unit of abatement investment, forcing the investment pathway to start low. The path then rises over time, commensurate with a rising carbon price (Hotelling, 1931), before declining again to the steady-state. This dynamic gives rise to the bell-shaped investment pathway shown in Vogt-Schilb et al. (2018). These considerations will be important when we discuss the sectoral allocation of abatement investment in Section 4.2.3.

2.1.2 "Strawman" model

We formulate our "strawman" model as an analogous problem to that of (2.1) without adjustment costs and abatement capital stock accumulation. In this formulation, the social planner controls the abatement *rate*, rather than abatement *investment*; this approach is analogous to conventional approaches in climate-economic integrated assessment models (Nordhaus, 2017; Barrage and Nordhaus, 2024) and has been used as a counterfactual case to isolate the effects of adjustment costs on decarbonization policy (see, e.g., Campiglio et al. (2022)). Rather than using investment costs, the social planner uses an abatement cost function in each sector $i \in \mathcal{I}$, $\gamma_i(a_{i,t})$, that maps abatement to monetary cost. The abatement cost functions γ_i are increasing and convex functions of abatement, but since capital accumulation dynamics are not present, the convexity of the cost function does not endow the model with adjustment costs. The remaining constraints are identical to those of the abatement investment model.

We can formulate this problem as a cost-minimizing optimization problem such that,

$$\begin{aligned} &\min_{\{a_{i,t}\}_{i \in \mathcal{I}}} \left[\sum_{t \in \mathcal{T}} \beta^t \sum_{i \in \mathcal{I}} \gamma_i(a_{i,t}) \right], \qquad (2.3) \\ \text{Subject to : } &\psi_{t+1} = \psi_t + \Delta t \left(\sum_{i \in \mathcal{I}} (\bar{a}_i - a_{i,t}) \right), \\ &0 \leq a_{i,t} \leq \bar{a}_i, \\ &0 \leq \psi_t \leq B, \\ &\psi_0 \text{ given.} \end{aligned}$$

This model is also treated in Vogt-Schilb et al. (2018) and we refer interested readers to their paper for a full analytical treatment, but note two theoretical findings analogous to those discussed for the abatement investment model.

This first relevant theoretical finding is that (2.3) admits a steady-state solution when $\psi_t = B$, but here the steady state is simply that $a_{i,t} = \bar{a}_i$ and the steady-state investment effort level is $\gamma_i(\bar{a}_i)$. The second finding is that the shape of the optimal abatement path follows an exponentially rising profile, as here the planner simply equalizes marginal abatement costs with the shadow price of emissions, the latter of which follows an exponential path analogous to a Hotelling rule (Hotelling, 1931).

2.2 Adding climate uncertainty

2.2.1 Abatement investment model with climate uncertainty

We now add uncertainty about the carbon budget to the model structure. Consider the abatement investment model outlined by (2.1), and allow the remaining carbon budget B to be uncertain. Since the carbon budget distribution is approximately Gaussian (Intergovernmental Panel on Climate Change, 2021), let $p_{B_{T^*}} \sim \mathcal{N}\left(\bar{B}_{T^*}, \sigma^2_{B_{T^*}}\right)$ be the Gaussian distribution of the carbon budget for some temperature target $T^* > 0$ where $\bar{B}_{T^*} > 0$ and $\sigma^2_{B_{T^*}} > 0$ are the average and variance of $p_{B_{T^*}}$, respectively. Further let $\tilde{B} := \{B_1, B_2, ..., B_N\}$ be a discrete set of quadrature nodes approximated from $p_{B_{T^*}}$, $N := |\tilde{B}|$ be the number of quadrature nodes, and $\mathcal{J} := \{1, 2, ..., N\}$ index each node. (Equivalently, $\tilde{B}$ can be thought of a set of samples from $p_{B_{T^*}}$, but since we use a quadrature approach in our numerical solution, it is more consistent to introduce $\tilde{B}$ as a set of quadrature nodes, see Section 4 of the *Supplementary Information*.) Assume the social planner learns the true value of the remaining carbon budget at a time t^* such that $0 < t^* < t_f$ (see Kelly and Kolstad (1999) and Fitzpatrick and Kelly (2017) for treatments where learning is continuous in time). Then we can write the social planner decarbonization problem as a 2-stage stochastic optimization problem, such that,

$$\min_{\{x_{i,t}\}_{i\in\mathcal{I}}} \left(\sum_{t=0}^{t^*-1} \beta^t \sum_{i\in\mathcal{I}} c_i(x_{i,t}) + \mathbf{E}_{\tilde{B}} \left[\mathcal{Q}\left(x_{i,t}; B^{(j)}\right) \right] \right), \tag{2.4}$$

$$\begin{aligned} \text{Subject to}: \psi_{t+1} &= \psi_t + \Delta t \sum_{i\in\mathcal{I}} \left(\bar{a}_i - a_{i,t}\right), \quad (2.5) \\ a_{i,t+1} &= a_{i,t} + \Delta t \left(x_{i,t} - \delta_i a_{i,t}\right), \\ 0 &\leq a_{i,t} \leq \bar{a}_i, \\ x_{i,t} &\geq 0, \\ a_{i,0}&, \psi_0, t^* \text{ given}, \end{aligned}$$

with

$$\mathcal{Q}\left(x_{i,t}; B^j\right) := \min_{\left\{x_{i,t}^{(j)}\right\}_{i\in\mathcal{I}, j\in\mathcal{J}}} \left(\sum_{t=t^*}^{t_f-1} \beta^t \sum_{i\in\mathcal{I}} c_i\left(x_{i,t}^{(j)}\right)\right), \tag{2.6}$$

$$\begin{aligned}
\text{Subject to : } \psi_{t+1}^{(j)} &= \psi_t^{(j)} + \Delta t \sum_{i\in\mathcal{I}} \left(\bar{a}_i - a_{i,t}^{(j)}\right), \\
a_{i,t+1}^{(j)} &= a_{i,t}^{(j)} + \Delta t\left(x_{i,t}^{(j)} - \delta_i a_{i,t}^{(j)}\right), \\
0 &\le a_{i,t}^{(j)} \le \bar{a}_i, \\
0 &\le \psi_t^{(j)} \le B^{(j)}, \\
x_{i,t}^{(j)} &\ge 0, \\
a_{i,t^*} &= a_{i,t^*}^{(j)}, \\
\psi_{t^*} &= \psi_{t^*}^{(j)}.
\end{aligned}$$

Note $a_{i,t}^{(j)}$ represents the abatement in sector $i \in \mathcal{I}$ at time $t \in \mathcal{T}$ in future state $j \in \mathcal{J}$ (the same notation is applied for cumulative emissions and investment, where the sectoral index is dropped for cumulative emissions). The final two constraints in (2.6) require continuity in abatement and cumulative emissions (the state variables) across the learning time. Note that if $t^* = 0$, there is no delayed learning, but the policymaker still minimizes the *expected* policy cost; this is distinct from the certainty solution of the model given by (2.1) where uncertainty is neglected entirely.

2.2.2 "Strawman" model with climate uncertainty

The "strawman" model can be formulated with delayed learning in an analogous way to the abatement investment model; using (2.3), we can write

$$\min_{\{a_{i,t}\}_{i\in\mathcal{I}}} \left(\sum_{t=0}^{t^*-1} \beta^t \sum_{i\in\mathcal{I}} \gamma_i(a_{i,t}) + \mathbf{E}_{\tilde{B}}\left[\mathcal{Q}\left(a_{i,t}; B^{(j)}\right)\right]\right), \tag{2.7}$$

$$\begin{aligned}
\text{Subject to : } \psi_{t+1} &= \psi_t + \Delta t \sum_{i\in\mathcal{I}} \left(\bar{a}_i - a_{i,t}\right), \\
0 &\le a_{i,t} \le \bar{a}_i, \\
&\psi_0, t^* \text{ given,}
\end{aligned}$$

with

$$\mathcal{Q}\left(a_{i,t}; B^{(j)}\right) := \min_{\left\{a_{i,t}^{(j)}\right\}_{i\in\mathcal{I}, j\in\mathcal{J}}} \left(\sum_{t=t^*}^{t_f-1} \beta^t \sum_{i\in\mathcal{I}} \gamma_i \left(a_{i,t}^{(j)}\right)\right), \tag{2.8}$$

$$\begin{aligned}
\text{Subject to : } \psi_{t+1}^{(j)} &= \psi_t^{(j)} + \Delta t \sum_{i\in\mathcal{I}} \left(\bar{a}_i - a_{i,t}^{(j)}\right), \\
0 &\leq a_{i,t}^{(j)} \leq \bar{a}_i, \\
0 &\leq \psi_t^{(j)} \leq B^{(j)}, \\
\psi_{t^*} &= \psi_{t^*}^{(j)}.
\end{aligned}$$

2.3 Calibrating marginal investment costs

Throughout the main text, we assume that our cost functions $c_i(x_{i,t})$ and $\gamma_i(a_{i,t})$ are quadratic, such that

$$c_i(x_{i,t}) = \frac{1}{2}\bar{c}_i x_{i,t}^2, \tag{2.9}$$

$$\gamma_i(a_{i,t}) = \frac{1}{2}\bar{\gamma}_i a_{i,t}^2. \tag{2.10}$$

In Section 3 of the *Supplementary Information*, we test this assumption with a higher-order cost function, and find the main results of the paper are not changed. Existing data for marginal abatement costs broken down by economic sectors (as first done by McKinsey & Company (2013)) allow us to directly fit values for $\bar{\gamma}_i$; later, we will utilize the recent IPCC estimates for marginal abatement costs (see Intergovernmental Panel on Climate Change (2022), Summary for Policymakers Figure 7, SPM.7). However, we are not currently aware of data for calibrating the marginal investment costs, $\bar{c}_i$. We therefore use the following prescription for translating marginal abatement costs into marginal investment costs. First, we require that the relative marginal abatement cost and marginal investment cost between each pair of sectors (i, j) are equal, such that

$$\frac{\bar{\gamma}_i}{\bar{\gamma}_j} = \frac{\bar{c}_i}{\bar{c}_j}, \tag{2.11}$$

as done in past work (Grubb, 1997; Vogt-Schilb et al., 2018). This defines marginal costs up to one remaining parameter. The final cost coefficient is calibrated by demanding that the total discounted costs between both approaches are equal for a given target under the assumption of certainty about the carbon budget (Grubb, 1997). Therefore, any differences in total policy costs once we add in the effects of uncertainty must come from the uncertainty itself, as the models are calibrated to be equally costly in the certainty case.

3 Simple Example of Model Behavior

To demonstrate the generic behavior of the abatement investment model and "strawman" model responses to uncertainty, we develop a simple numerical experiment. The goal of this section is not

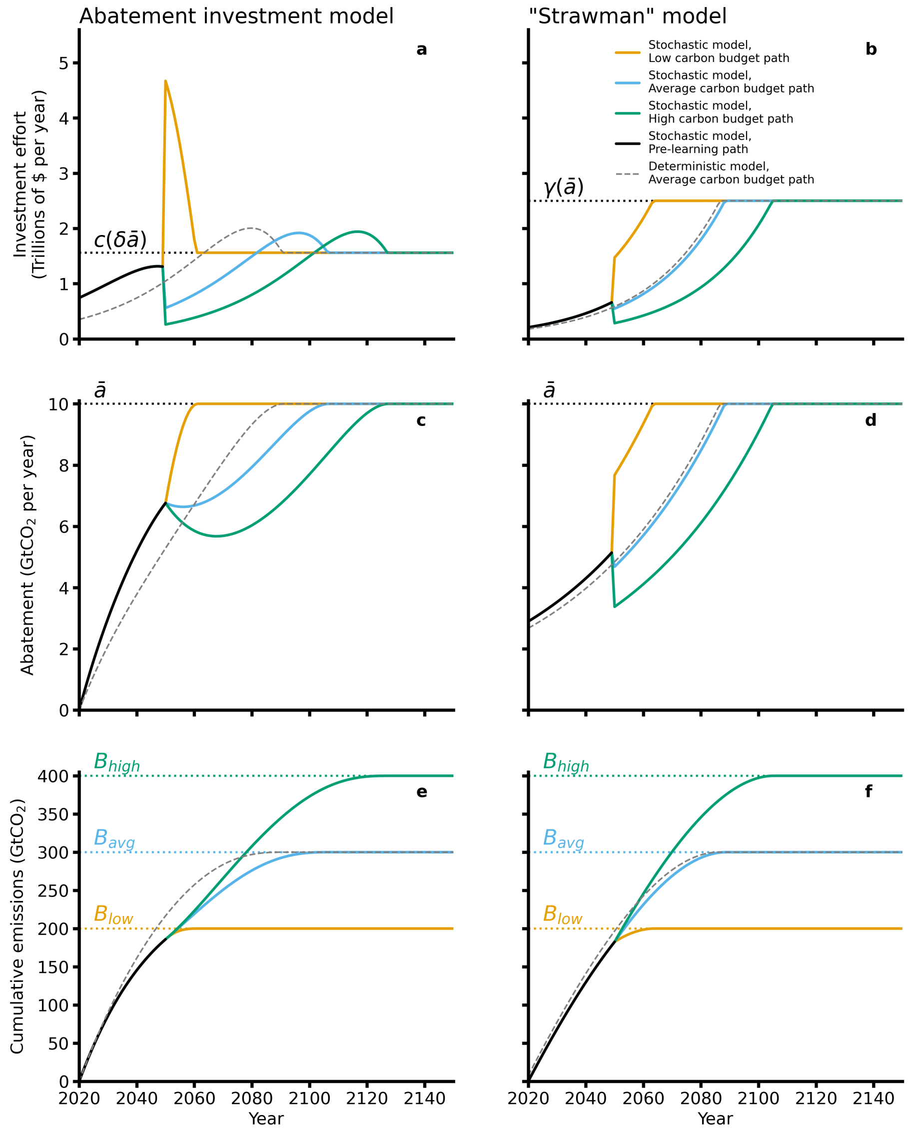

Figure 1: Example abatement investment model and "strawman" model output. The left-hand column shows example output from the abatement investment model, while the right-hand column shows example output from the "strawman" model. Panel **a** shows the optimal investment pathway in the deterministic abatement investment model with an average carbon budget (grey, dashed line), as well as the paths for the stochastic model with three carbon budgets: high (green solid line), average (blue solid line), and low (yellow solid line). Note the solid black line shows the stochastic model path prior to learning where all three possible outcomes for the carbon budget follow the same policy. The dotted black line shows the steady-state investment effort level. Panel **c** shows optimal abatement pathways, and here the black dotted line shows the economy-wide annual emissions rate. Panel **e** shows the optimal cumulative emissions paths, where the three dotted lines are the three carbon budgets, labeled by their colors. Panels **b**, **d**, and **f** mimic that of **a**, **c**, and **e**, respectively, for the "strawman" model.

to emulate real-world economies, but rather to provide intuition for how both the deterministic and stochastic models behave before moving to a more complicated setup with multiple, heterogeneous economic sectors.

For the following discussion, we assume there is one economic sector with a marginal abatement cost, $\bar{\gamma}$, of 50 (\$ / tCO_2) / ($GtCO_2$ / yr), annual emissions rate, $\bar{a}$, of 10 $GtCO_2$ per year, and a capital deprecation rate, δ, of 5% annually. Following the procedure outlined in Section 2.3 implies the marginal investment cost, $\bar{c}$, of 12500 (\$ / tCO_2) / ($GtCO_2$ / yr^2). The social planner uses a social discount rate, r, of 2%. The expected carbon budget, B, is 300 $GtCO_2$. We assume that thirty years after the start time, uncertainty in the carbon budget is resolved, and is either 200 $GtCO_2$ (the worst case), 300 $GtCO_2$ (is in line with our expectations before learning), or 400 $GtCO_2$ (the best case). Each of these outcomes are equally likely. We use this parameter set to solve (2.4) and (2.7), and in our deterministic model runs, we solve (2.1) and (2.3) with a carbon budget of 300 $GtCO_2$ for comparison with the stochastic case.

We visualize the results of this experiment in Figure 1, where the optimal investment, abatement, and cumulative emissions pathways for the abatement investment model ("strawman", resp.) are shown in the left-hand (right-hand, resp.) column. In each case, the results from solving the deterministic models are shown in grey, dashed lines. Prior to learning, the optimal stochastic path is shown in the solid black line, while the three possible stochastic model paths after learning are shown in the solid, colored lines. Each color corresponds to a different carbon budget. Notice that before the learning date (here, 2050), there is one path for the stochastic model, that then branches into three after learning (marked by the change in line color from black to colored), showing how the policymaker can adjust the policy course after learning the true value of the carbon budget.

Starting with the deterministic models (grey, dashed lines in Figure 1), Fig. 1**a–b** show how the abatement investment model and "strawman" model result in markedly different investment pathways. The abatement investment model starts low and rises over time, before declining to the steady-state investment value of $c(\delta\bar{a})$, therefore following the bell-shaped path mentioned in Section 2.1. On the other hand, the investment pathway for the "strawman" model follows an exponential path, since the planner equalizes marginal abatement costs with an exponentially rising carbon price (akin to a Hotelling rule (Hotelling, 1931)).

The optimal abatement pathways between the abatement investment model and the "strawman" model also differ substantially. In the abatement investment model, since the level of abatement capital stocks are initially zero and must be built over time, the optimal abatement path starts at zero and grows steadily. On the other hand, the "strawman" model has a discontinuous jump at the initial time point because such capital stock dynamics are neglected; the planner can essentially "snap their fingers" and eliminate 30% of greenhouse gas emissions "overnight". This leads to an initially high and convex path of abatement in the "strawman" model, as opposed to the initially low and concave abatement path in the abatement investment model. Finally, the cumulative emissions path between both models are qualitatively similar, albeit with the abatement investment model path being slightly more concave than the "strawman" model path.

These qualitative differences between model behavior are accentuated by the addition of climate uncertainty (see the colored lines in Figure 1). In the abatement investment model, we find a massive spike in investment in the worst-case carbon budget scenario (yellow line in Fig. 1**a**); this large increase in investment is required to overcome adjustment costs while still limiting emissions below the lowest carbon budget (B_{low} in Fig. 1**e**). Such a large injection of investment, however, is not required in the "strawman" framework because of the lack of adjustment costs (yellow line in Fig. 1**b**), thus limiting the additional costs imposed by uncertainty in the "strawman" model.

Comparing the average and high carbon budget investment paths between models further highlights how neglecting adjustment costs and capital stock dynamics changes the model response to uncertainty. In the abatement investment model, after learning about the carbon budget the investment paths for the average and high carbon budgets decrease before tracing out bell-shaped paths. This causes abatement to decrease slowly and continuously, as some abatement capital stocks depreciate (i.e., $\delta a > x$ because $\dot{a} < 0$) during this period. This is not the case in the "strawman" model, where we find that after learning, both the investment effort and the abatement decrease instantaneously by a small amount for the average carbon budget path and decrease substantially in the high carbon budget case. This shows how, when adjustment costs and transition dynamics are neglected, the planner can "snap their fingers" to decrease abatement in response to good news about the climate. This contributes to far less precautionary spending in the "strawman" model relative to the abatement investment model.

These differences between the investment pathways manifest in different optimal abatement and cumulative emissions pathways, as well as total policy costs. We find that precautionary spending in the abatement investment model amounts to about 1 $GtCO_2$ of precautionary abatement (i.e., the difference between the green solid line and grey dashed line in 2050 in Fig. 1**c**) compared to about 0.2 $GtCO_2$ per year in the "strawman" case. We further find that the total policy cost in the stochastic abatement investment model exceeds that of the stochastic "strawman" model by about 7% (despite the models being calibrated to have the same total policy cost in the deterministic case), highlighting how the interaction between adjustment costs and climate uncertainty leads to higher policy costs. We now move to design numerical experiments in-line with recent estimates of marginal abatement costs, carbon budgets and their uncertainties, and social discount rates.

4 Calibrated Numerical Experiments

4.1 Calibration

4.1.1 Global parameters

The remaining carbon budget. We use a recent study (Dvorak et al., 2022) that estimates the carbon budget probability distribution using a state-of-the-art, simplified climate emulator, FaIR (Smith et al., 2018; Leach et al., 2021), which was also used by the IPCC for the same purpose (Intergovernmental Panel on Climate Change, 2021). We utilize two temperature targets in our analysis: 1.7 °C and 2 °C. We explore a 2 °C temperature target because of its relevance in the Paris Agreement targets, while 1.7 °C represents a more ambitious, yet achievable target under the worst-case scenario. Indeed, in the worst case, we cannot achieve the 1.5 °C temperature target without considerable overshoot of the target temperature because the worst-case carbon budget has already been depleted. As such, we

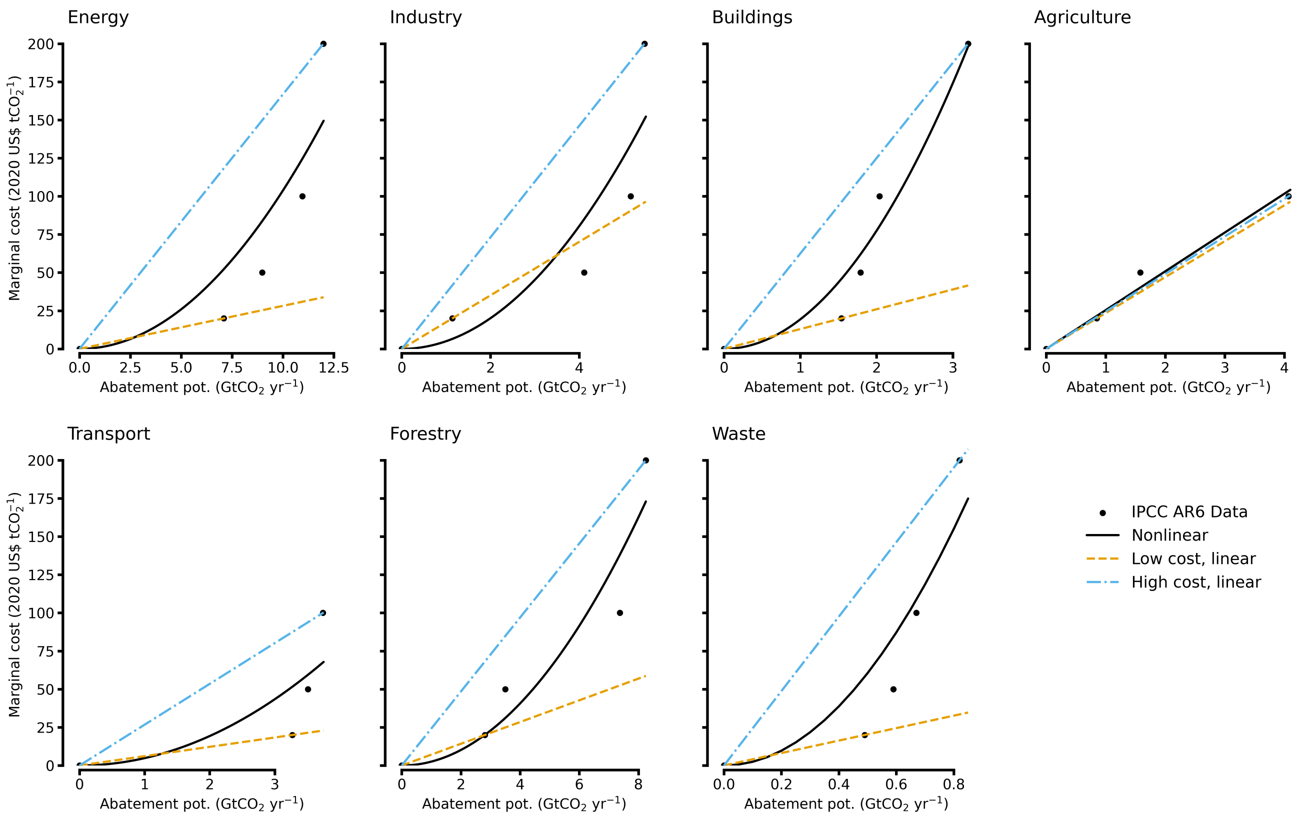

Figure 2: Calibrating marginal abatement costs. Marginal abatement costs broken down by sector using IPCC AR6 data (Intergovernmental Panel on Climate Change, 2022). Black lines represent a nonlinear marginal abatement cost (in particular, a quadratic fit for all sectors other than agriculture); gold lines represent a linear fit to lowest cost bracket marginal cost values; blue lines represent a linear fit to highest cost bracket marginal abatement cost data.

explore the 1.5 °C temperature target in Section 3 of the *Supplementary Information* using a heavily truncated distribution, but exclude the numerical results from the main analysis.

The carbon budget distribution is approximately Gaussian (Intergovernmental Panel on Climate Change, 2021; Matthews et al., 2021; Dvorak et al., 2022), and we sample the $1^{\text{st}} - 99^{\text{th}}$ percentile range of their distribution using a truncated normal distribution, see Table 1. We truncate our carbon budget distribution for two reasons. The first is that, in our framework, negative values of the carbon budget make the model infeasible to solve, and hence are excluded. The second is that the carbon budget distribution estimated in Dvorak et al. (2022) is found using a parameter ensemble approach that limits the ability of the model to capture the tails of the carbon budget distribution; indeed, the paper only quotes results up to the 99^{th} percentile. We apply a Gauss-Hermite quadrature rule (Cai and Judd, 2010) to approximate the expectation operator, following the approach outlined in Cai and Judd (2015) for applying quadrature rules to truncated normal distributions (see Section 4 of the *Supplementary Information* for the technical details of this approach).

Social discount rate. We set the social discount rate to 2% yr^{-1}, in line with a recent international expert elicitation (Drupp et al., 2018) and the US EPA's prevailing rate for their social cost of carbon estimates (National Center for Energy Economics, 2022).

Emissions scenario. It remains to define a "business-as-usual" emissions scenario, a notoriously difficult task. Remarkably, the total abatement potential described by the IPCC is $\sim$40 $\text{GtCO}_2\ \text{yr}^{-1}$, which is approximately equal to the peak emissions of SSP2–4.5 (the "middle of the road" emissions scenario used by the IPCC (Intergovernmental Panel on Climate Change, 2021)) and the Resources for the Future-socioeconomic projects used in the United States' Environmental Protection Agency's estimates of the social cost of carbon (Riahi et al., 2017; National Center for Energy Economics, 2022). In each of these baselines, emissions are expected to rise to just above 40 $\text{GtCO}_2\ \text{yr}^{-1}$ before declining. We therefore assume that, without abatement, emissions will maintain their peak levels in each of these emissions scenarios, i.e., at $\sim$40 $\text{GtCO}_2\ \text{yr}^{-1}$. This describes our "business-as-usual" scenario; we probe the impact of assuming an increasing emissions baseline in Section 2 of the *Supplementary Information*. One caveat with our approach is that some green technologies are already cost-competitive with their dirty equivalents, and therefore do not necessarily require a carbon price to be assimilated into the economy; indeed, these "costless" abatement technologies are a key source of uncertainty in projecting the costs of the green transition (Kotchen et al., 2023). Therefore, our assumption that emissions will remain at peak levels without a carbon price can be viewed as conservative, as all abatement in our approach is spurred by the presence of a carbon price. We solve the model for 300 years to approximate an infinite horizon solution, as studied in Vogt-Schilb et al. (2018).

4.1.2 Sectoral parameters

Abatement potentials. We lift the abatement potentials for each sector from the third working group's contribution to the IPCC's AR6 (Intergovernmental Panel on Climate Change, 2022). These abatement potentials are estimated via bottom-up studies of the potential for a variety of abatement technologies to be deployed at scale, including technological interventions (e.g., solar and wind energy to lower energy sector emissions) and behavioral changes (e.g., altering dietary habits to lower agricultural

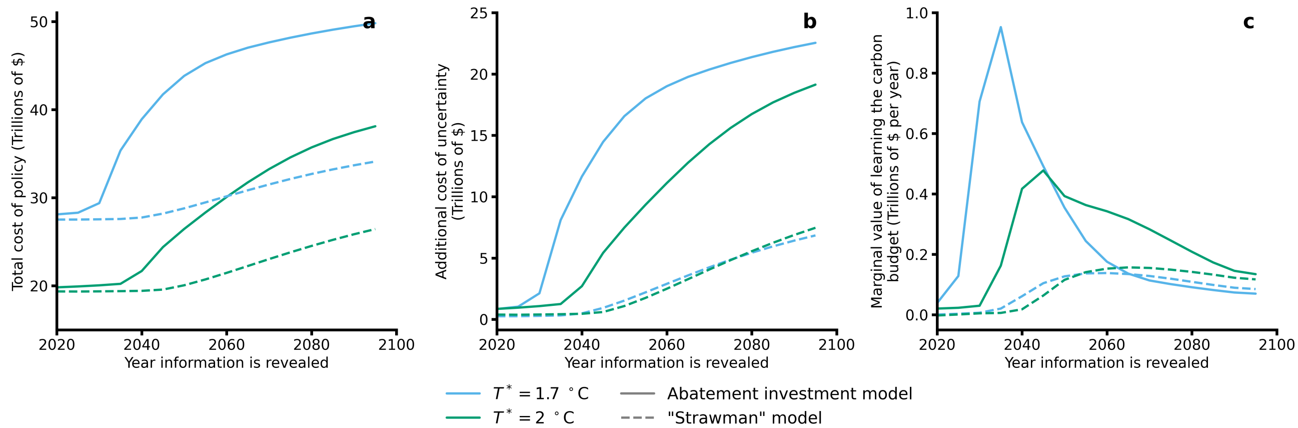

Figure 3: Effect of uncertainty and delayed learning on aggregate policy cost. In Panel **a**, we show the total discounted policy cost for each temperature target (see the legend) for the abatement investment model that includes the impact of adjustment costs (solid lines) and "strawman" model, with does not include adjustment costs (dashed lines). Panel **b** is as **a**, net of the cost of the corresponding policy without uncertainty. Panel **c** shows the marginal cost of delaying learning by one additional year.

emissions).

Marginal abatement and investment costs. We use the procedure outlined in Section 2.3 to compute the marginal investment cost coefficients using the IPCC marginal abatement cost data (see Intergovernmental Panel on Climate Change (2022), Summary for Policymakers Figure 7, SPM.7) shown in Figure 2; throughout the main text, we use the low cost, linear calibration from Figure 2. Our results are summarized in Table 1. Note that while we use the same discount rate, sectoral marginal abatement costs, emissions rates, and capital depreciation rates for each temperature target, the sectoral marginal investment costs are different owing to the total policy cost being different for each target; since equalizing the total cost of policy between the abatement investment model and the "strawman" model is our final constraint on the sectoral marginal investment cost parameters (in addition to the relative marginal investment cost being equal to the relative marginal abatement cost for every pair of sectors, as shown in (2.11)), the final column in Table 1 changes depending on the temperature target under consideration.

4.2 Results

4.2.1 Aggregate policy costs

We first explore the impact of uncertainty in the carbon budget on the aggregate cost of optimal policy for each temperature target (1.7 °C and 2 °C) in Figure 3. In Fig. 3**a**, we show the total discounted cost of the optimal policy using the abatement investment model (solid lines) and "strawman" model (dashed lines); we show the policy cost net of the equivalent certain policy cost in Fig. 3**b**. We find that across models and temperature targets, the later the policymaker learns the true value of the carbon budget, the higher the total policy cost relative to certainty.

The strongly increasing relationship between total policy cost and the delay in learning about the

Table 1: Calibration parameters for the abatement investment and "strawman" models for each temperature target using the low cost, linear calibration.

		$T^* = 1.7$ °C		
$t_f = 300$ yr	$r = 2$ % yr^{-1}	$\bar{B} = 770$ GtCO_2	$\sigma_B = 220$ GtCO_2	
Sector	$\bar{\gamma}$ $\left[\frac{\$\ \text{tCO}_2^{-1}}{\text{GtCO}_2\ \text{yr}^{-1}}\right]$	$\bar{a}$ [GtCO_2 yr^{-1}]	δ [% yr^{-1}]	$\bar{c}$ $\left[\frac{\$\ \text{tCO}_2^{-1}}{\text{GtCO}_2\ \text{yr}^{-3}}\right]$
Waste	40.82	0.82	3.3	17470
Industry	17.54	5.47	4	7510
Forestry	7.12	8.25	0.8	3050
Agriculture	23.53	4.07	5	10070
Transport	6.12	3.74	6.7	2620
Energy	2.82	11.99	2.5	1210
Buildings	12.99	3.2	1.7	5560
		$T^* = 2$ °C		
$t_f = 300$ yr	$r = 2$ % yr^{-1}	$\bar{B} = 1340$ GtCO_2	$\sigma_B = 380$ GtCO_2	
Sector	$\bar{\gamma}$ $\left[\frac{\$\ \text{tCO}_2^{-1}}{\text{GtCO}_2\ \text{yr}^{-1}}\right]$	$\bar{a}$ [GtCO_2 yr^{-1}]	δ [% yr^{-1}]	$\bar{c}$ $\left[\frac{\$\ \text{tCO}_2^{-1}}{\text{GtCO}_2\ \text{yr}^{-3}}\right]$
Waste	40.82	0.82	3.3	20900
Industry	17.54	5.47	4	8980
Forestry	7.12	8.25	0.8	3644
Agriculture	23.53	4.07	5	12040
Transport	6.12	3.74	6.7	3130
Energy	2.82	11.99	2.5	1440
Buildings	12.99	3.2	1.7	6650

carbon budget can be explained by how long the social planner has to formulate a policy that is robust to the worst-case scenarios for the carbon budget. If the policymaker learns about the carbon budget 10 years after the start date, then their policy must be robust to the worst-case carbon budgets for a relatively short amount of time. Therefore, there is some increased spending to preemptively plan for the worst-case, but the planner finds out relatively quickly if the worst-case ever materializes, and can relax the policy in the likely event it does not. On the other hand, if the policymaker learns about the carbon budget 50 years after the start date, then the policymaker must plan for the worst-case for much longer; by then, a significantly higher amount of precautionary spending has occurred, which drives up the total policy cost, especially for ambitious temperature targets.

A similar logic explains the "S"-shaped dynamics of the additional cost of uncertainty. Using the 1.7 °C target as an example in Fig. 3**b**, the difference between learning in 2020 and 2030 is small owing to only a small amount of precautionary spending required for a robust policy. Yet this difference gets large after 2030 when the worst-case carbon budgets must be planned for over an increasingly long period of time. Finally, the difference between 2080 and 2090 is negligible, as there is only so much additional precautionary spending that can make the policy marginally more robust to learning of a catastrophe later on. This dynamic is present across temperature targets and model types, though the curves are less steep for more lenient temperature targets (as catastrophes are not so extreme).

The "S"-shaped behavior of total policy cost can be clearly understood by computing the marginal value of learning about the carbon budget a year earlier (see Fig. 3**c**); to this end, we compute the marginal value of learning the carbon budget a year earlier by taking the derivative of Fig. 3**b** with respect to the year the information is revealed. We find that the value of learning the carbon budget a year earlier can be as high as $1t per year. This implies a potentially substantial cost-savings from learning the true value of the carbon budget earlier rather than later, and furthermore, indicates that the value of waiting to invest in abatement in order to learn more about the carbon budget must be smaller than the increase in the expected net-present value of the social planner's costs if they delay (Dixit and Pindyck, 1994). This implies a positive "value of learning" about the carbon budget, making reductions in carbon budget uncertainty valuable for policymakers. In simple terms: bad news is difficult, bad news late is worse.

A final insight from Figure 3 is that adjustment costs accentuate the impact of delayed learning about the carbon budget. We find that for both temperature targets, the additional cost of uncertainty in the abatement investment model far exceeds that of the "strawman" model, showcasing the additional price of delayed learning when adjustment costs are present. This is easily explained: adjustment costs significantly penalize rapid decarbonization, which is exactly what is necessary if the worst-case arises. Put differently, adjustment costs make catastrophes far worse than when they are absent. Integrated assessment models that do not capture the capital accumulation and adjustment cost dynamics described in our abatement investment model are therefore missing important investment dynamics and may be underestimating near-term decarbonization investment needs, as well as the impact of uncertainty on decarbonization pathways.

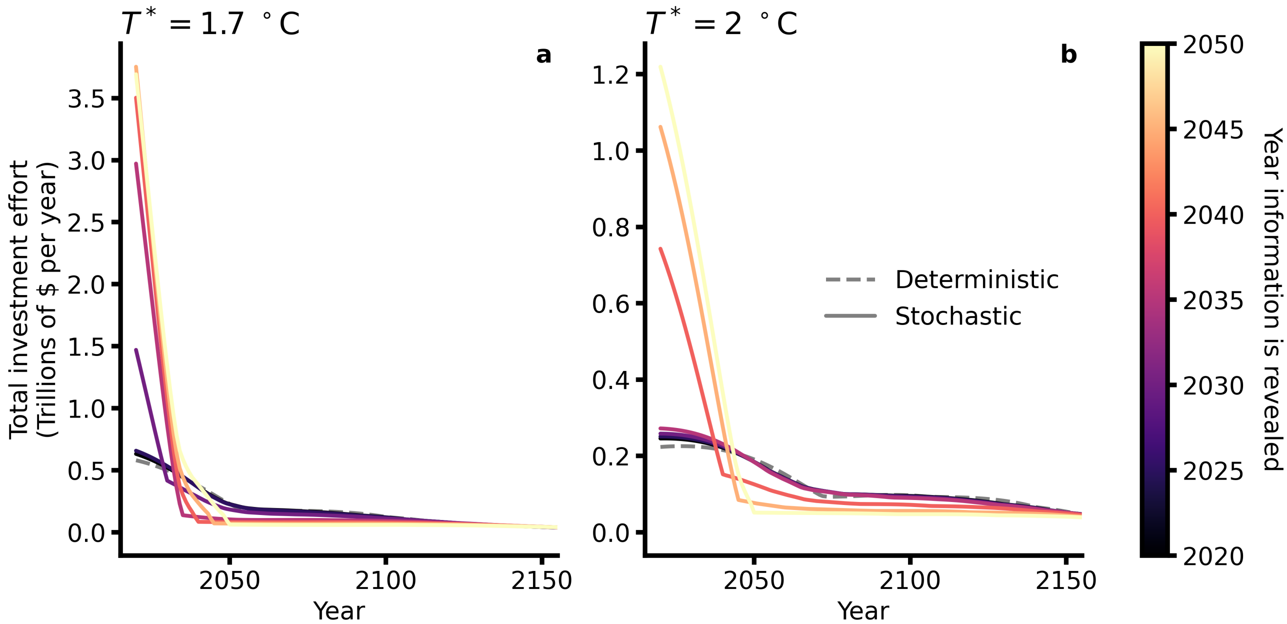

Figure 4: Effect of uncertainty and delayed learning on the temporal distribution of spending. Panel **a** shows the average, total discounted investment effort, $\beta^t \mathbb{E}\left[\sum_{i \in \mathcal{I}} c_i(x_{i,t})\right]$, for a 1.7 °C temperature target (solid lines), where the colors represent different years where information about the carbon budget is revealed. The deterministic policy is shown in the grey, dashed line. Panel **b** is as **a**, but for the 2 °C temperature target.

4.2.2 Temporal distribution of investments

If uncertainty about the carbon budget increases the total cost of policy, how does the social planner distribute the additional spending in time? We explore the temporal shift in spending in Figure 4, where we show the average, discounted, total investment effort (i.e., the sum of investment costs for each sector discounted to the present-day, $\beta^t \mathbb{E}\left[\sum_{i \in \mathcal{I}} c_i(x_{i,t})\right]$) over time. We find that the inclusion of uncertainty leads to a front-loading of abatement investment effort relative to the certainty policy across temperature targets; while the deterministic and stochastic investment paths both exhibit a declining profile over time, we find that the stochastic investment paths feature much higher near-term investments and decline faster than the deterministic paths.

In addition, we find that the rate at which spending is shifted depends strongly on the desired temperature target. This can be explained by the difference in the worst-case carbon budgets between targets. In the 1.7 °C target case, the worst-case carbon budget could be as low as $\sim$ 300 $\mathrm{GtCO_2}$, which given our baseline levels of emissions ($\sim$ 40 $\mathrm{GtCO_2}$ yr^{-1}) would imply the temperature target could be reached in $\sim$ 7.5 years in the absence of abatement. Hence, a small delay in learning the carbon budget leads to a large amount of front-loaded precautionary spending, given the overall urgency of the temperature target. On the other hand, for the 2 °C target, in the worst-case the policymaker would have $\sim$ 13 years to decarbonize in the absence of abatement, which delays the strong front-loading of precautionary investment by $\sim$10 years relative to the 1.7 °C temperature target.

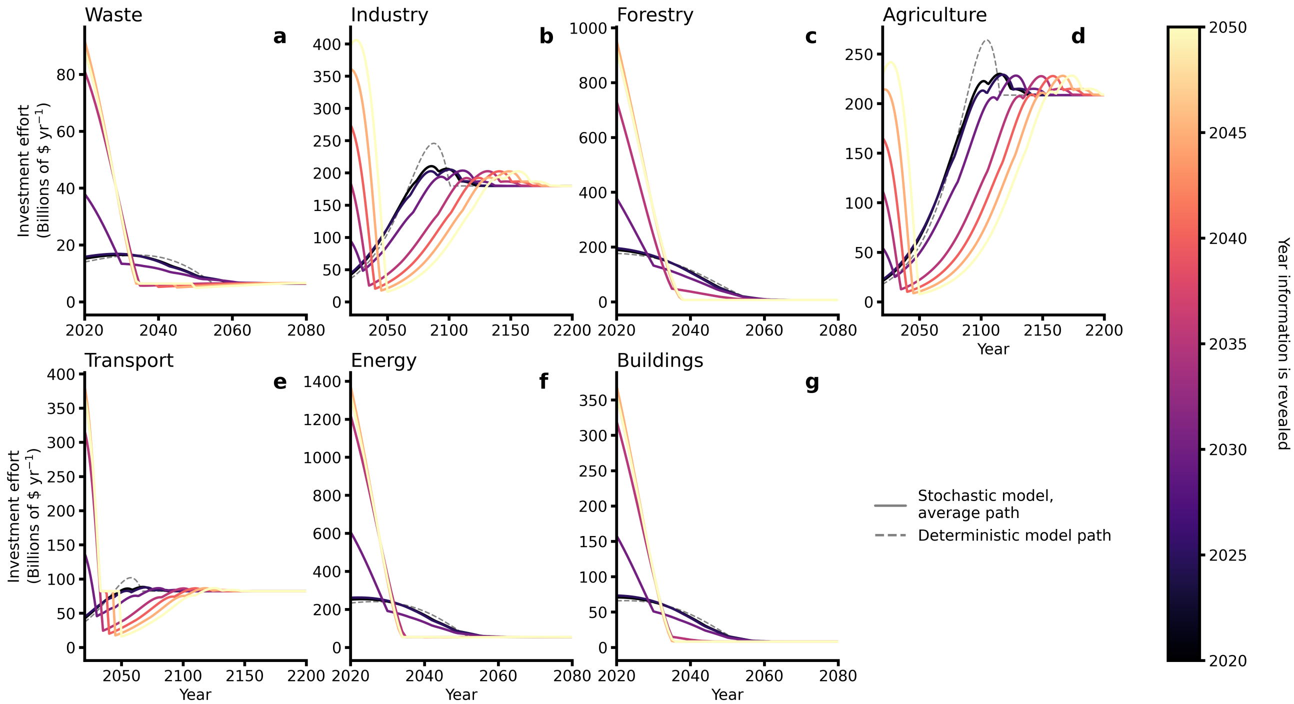

Figure 5: Effect of uncertainty and delayed learning on sectoral allocation of abatement investment. In each panel we show the average, total investment effort in each sector for the stochastic abatement investment model (solid lines), where the color of each curve corresponds to the year the information about the carbon budget is revealed. The grey, dashed line is the deterministic investment path in each sector. Notice how uncertainty in the carbon budget creates more urgency for near-term investments, causing bell-shaped investment paths to shift to declining paths in agriculture and industry, while already declining investment path sectors such as energy get steeper relative to their no uncertainty case. Note these investment pathways are for the 1.7 °C temperature target, and that panels **b**, **d**, and **e** show the investment effort until to 2200, while the remaining panels show the investment effort until 2080.

The front-loading of abatement investment effort shown in Figure 4 provides further insight into the finding that delaying learning about the carbon budget causes total policy costs to increase, as shown in Figure 3. As an increasing amount of abatement investment is front-loaded to account for uncertainty, discounting becomes less relevant for the total policy cost, as spending occurs much sooner than in the certainty case.

4.2.3 Sectoral allocation of investments

If climate uncertainty leads to additional up-front costs, which sectors are getting the lion's share of the additional investment? We sectorally disaggregate investment effort in Figure 5 to explore this question. We find that across sectors, as learning about the carbon budget happens later, more investment effort occurs earlier; this is a form of precautionary spending, as the social planner prepares for the worst-case scenario, as discussed above. For example, easy-to-abate sectors such as energy and forestry see a ∼4-5× increase in investment effort relative to certainty in 2020 when the carbon budget is revealed in mid-century. In hard-to-abate sectors, such as industry, agriculture, and transport (Fig. 5**b,d,e**), we observe that the optimal investment pathway shifts from a bell-shaped path to a declining path prior to

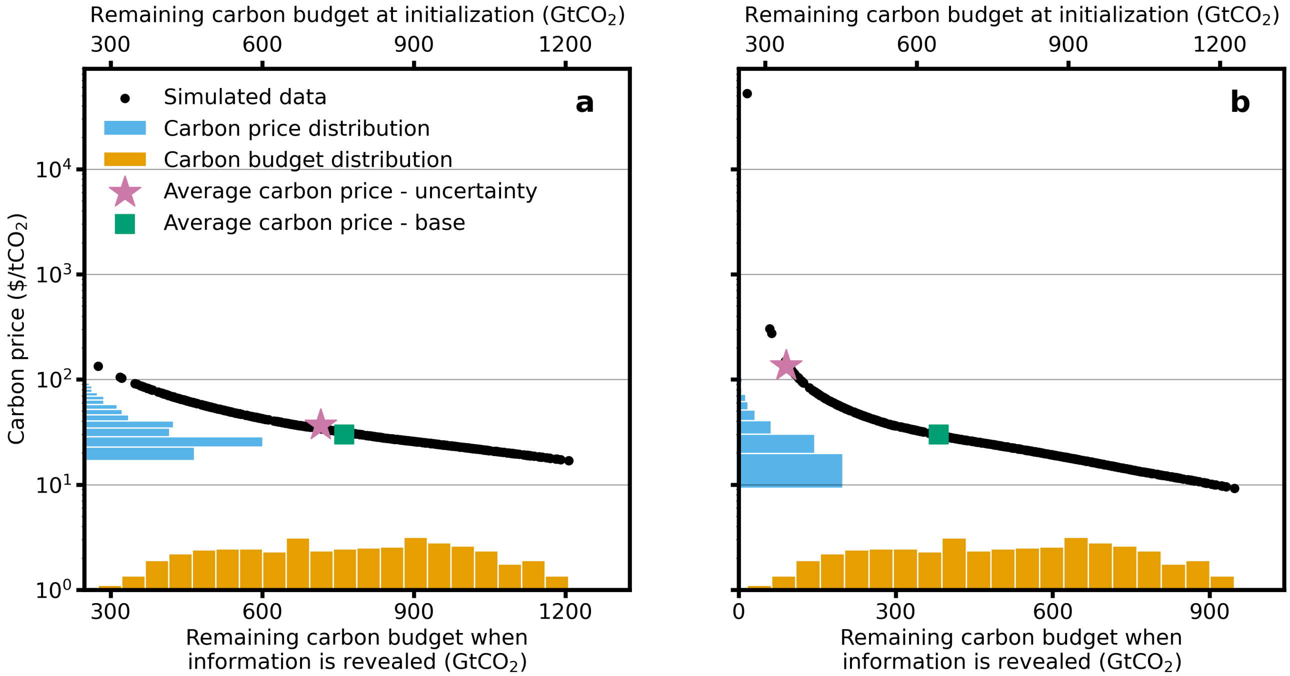

Figure 6: Effect of uncertainty and delayed learning on the carbon price. In Panel **a**, we plot the remaining carbon budget distribution when the information is revealed (yellow bars) against the carbon price distribution (blue bars) for the case where information is revealed in 2020. The black dots are the simulated values of the carbon price for each carbon budget; the pink star is the average carbon price with uncertainty; the green square is the carbon price in the certainty case for comparison with the uncertainty case. Note reading the gold bars with the top x-axis dictates the initial distribution of the carbon budget, while reading the gold bars with the bottom x-axis dictates how much of the carbon budget remains when the true value is revealed. Panel **b** is as **a**, but for when the carbon budget is revealed in 2030. Note the data used in this figure is for a 1.7 °C temperature target.

learning (further buttressing the front-loading of investment effort observed in Figure 4) before relaxing to a bell-shaped path after learning, overall forming a "U"-shaped path. This behavior is similar to that of the average carbon budget and high carbon budget investment paths shown in Figure 1**a**.

The change in sectoral allocation of abatement investment can be explained by the influence of climate uncertainty on the carbon price. In Figure 6, we plot the distribution of carbon prices after information is revealed on the vertical axis as a function of the remaining carbon budget when information is revealed on the horizontal. In Fig. 6**a**, information is revealed in 2020, while in Fig. 6**b** information is revealed in 2030. On the top horizontal axis, we also show the carbon budget at the beginning of the simulation for comparison with the carbon budget when information is revealed. Notice that when information is revealed in 2020, the worst-case carbon budget is about 300 $GtCO_2$, but when information is revealed in 2030, the worst-case carbon budget is near zero. While the carbon prices shown in Figure 6 are relatively low relative to recent estimates (e.g., the US EPA estimate of $190 per tCO_2 (Rennert et al., 2022)), the overall dynamics of our results are robust to a number of alternative specifications, including higher cost specifications that drive higher carbon prices (see Section 3 and 5 of the *Supplementary Information*).

We find that, if the policymaker learns about the carbon budget in 2020, there is an approximately linear relationship between the carbon budget and the log of the carbon price (Figure 6**a**); it follows that since the underlying carbon budget distribution is Gaussian that the carbon price is approximately log-normally distributed. Note here learning about the carbon budget in 2020 is equivalent to minimizing the expected, total investment cost across a distribution of possible carbon budgets. This procedure results in a slightly higher carbon price than the deterministic case because the planner still has to consider worst-case outcomes, but since uncertainty is resolved immediately, there is less hedging behavior when the carbon budget is learned in 2020 as is found when learning is delayed.

If the policymaker learns in 2030, we find the log-normal distribution of carbon prices breaks down; the carbon price distribution inherits a "heavy tail" and becomes positively skewed, as shown in Figure 6**b**. This is because the amount of the carbon budget left when information is revealed can be as low as ~ 30 $GtCO_2$ in the worst case, since some amount of emissions has occurred between 2020 and 2030. If this worst-case scenario materializes – akin to a so-called "green swan" scenario (Bolton et al., 2020) – the social planner has no choice but to impose an incredibly stringent carbon price to meet the 1.7 °C temperature target. (Recall the social planner is constrained to always meet the target in our framework; in actuality, the planner may simply choose to target a higher temperature threshold if the required carbon price is prohibitively high.) This heavy tail makes the average carbon price higher than in the case of learning in 2020, and explains why the average carbon price (pink star in Figure 6**b**) is outside the main body of the carbon price distribution in 2030.

Furthermore, Figure 6 demonstrates why the "value of learning" – that is, the value in learning the carbon budget early in the investment path (Dixit and Pindyck, 1994; Lemoine and Rudik, 2017; Barnett et al., 2021) – in our framework is positive: the carbon price distribution is only heavy tailed on the high end, leading to an asymmetry between the worst-case and best-case responses. In the worst-case, the social planner must impose an incredibly stringent carbon price to accommodate the lower-than-expected carbon budget, whereas in the best-case, the social planner only relaxes the carbon price slightly. This asymmetry gives rise to the precautionary investment strategy undertaken by the social planner. Furthermore, learning about the carbon budget sooner limits the degree to which the carbon price distribution becomes heavy-tailed, supporting a positive value of learning.

The carbon price influences the shape of the optimal investment pathway via the inequality shown in (2.2). When the carbon price increases, the relative cost of either paying the carbon price or a marginal unit of abatement investment is shifted in favor of investment. This leads to the social planner investing more aggressively in easy-to-abate sectors, which causes their already declining investment pathways to have an even steeper profile. In hard-to-abate sectors, the bell-shaped profile shifts to a declining path as information is learned later and the average carbon price increases. This explains how investments across economic sectors become more urgent in the near-term as a result of delayed learning about climate uncertainty, as shown in Figure 5.

This behavior sheds further light on our previous findings that later learning about the carbon budget leads to the front-loading of abatement investment, shown in Figure 4. The front-loading of investment relative to certainty is caused by the increased urgency in investments to accommodate the worst-case carbon budget that causes: (i) the optimal investment pathway for hard-to-abate sectors to change from a bell-shaped path to a "U"-shaped path that declines prior to learning and is bell-shaped thereafter, and (ii) the already-declining investment pathways of easy-to-abate sectors to decline more steeply for higher carbon prices. The high carbon prices, in turn, are caused by the remaining carbon budget when the information is revealed to be smaller than the initial distribution; this leads to a risk of extremely high carbon prices when the carbon budget is revealed to be low, which drives up the average carbon price. In other words: the later bad news arrives, the more extreme the policy response must be to accommodate the bad news, here in the form of an extremely high carbon price that drives up the cost of policy.

5 Discussion

In this work, we explored the impact of climate uncertainty and delayed learning about this uncertainty on optimal abatement investment pathways when adjustment costs are present. We posit three main takeaway points: (i) that climate uncertainty increases the overall costliness of policy, especially when adjustment costs are included, (ii) climate uncertainty causes abatement investment to be front-loaded relative to an equivalent policy without uncertainty, and (iii) the sectoral allocation of abatement is significantly impacted by climate uncertainty, especially in hard-to-abate sectors, where we find a shift from a bell-shaped investment pathway to a "U"-shaped path that declines prior to learning and is bell-shaped after learning. Each of these effects is increased when uncertainty about the carbon budget is resolved later in the century, as the policymaker has to formulate policies that are robust to the worst-case climate outcome (i.e., a high climate sensitivity) for a longer period of time.

This work provides many insights that are relevant for climate policymakers. The first is, as shown in Figure 3, adjustment costs exacerbate the effects of climate uncertainty on optimal decarbonization investment pathways and the overall cost of decarbonization. In essence, this is because adjustment costs limit the ability of policymakers to hasten the green transition without incurring potentially prohibitively high costs in the event of a worst-case carbon budget. Including the effects of capital accumulation dynamics and adjustment costs in integrated assessment models is therefore essential, so that crucial transition dynamics can be fully accounted for in climate policy recommendations. Models that do not capture these dynamics likely underestimate near-term decarbonization investment needs and policy costs, especially in the presence of climate risk.

A second takeaway is that, if one takes meeting a given temperature target seriously, then considering the worst-case scenario for carbon budgets can have a sizable influence on optimal policy. This is especially the case when considering stringent temperature targets, when learning the precise value of the carbon budget (or its uncertainty is lessened to the point of near-certainty) is significantly delayed, or when both of these effects are present. Note this argument would still hold even in a cost-benefit framework, where rather than learning about the carbon budget we learn about how climate damages depend on cumulative emissions (which are linked to temperature via the transient climate response to emissions (Matthews et al., 2009)). The intuition here is analogous to using greenhouse gas targets as insurance against worst-case climate damages, as argued by Weitzman (2012). Here we add to this narrative by quantifying how the presence of adjustment costs makes the worst-case scenario even more important to consider, as if the worst-case scenario materializes, adjustment costs will make an aggressive decarbonization strategy especially expensive. These transition dynamics would therefore increase the "insurance value" of stringent climate targets, to borrow the language of the Weitzman

(2012) framework.

An additional takeaway from our work is that narrowing the uncertainty in the carbon budget could lead to non-trivial monetary savings for policymakers. Indeed, the logic of the marginal cost of uncertainty shown in Figure 3**c** cuts both ways; learning a year earlier about the carbon budget could also lead to high levels of savings, especially if we are able to learn before 2030. This could be a fruitful direction for research for climate scientists that would complement many policy discussions.

Two final takeaways are demonstrated in Sections 1 and 2 of the *Supplementary Information*, respectively. The first is that the inclusion of direct air capture only leads to marginal cost-savings as a function of learning time for both the 1.7 °C target and the 2 °C target. Importantly, the role of direct air capture in this case is as a cost-offset for more expensive sectors such as agriculture and industry. This gives context to the role of direct air capture in the decarbonization discussion: direct air capture can offset the costs of decarbonizing hard-to-abate sectors, but should not be considered a substitute for decarbonizing other relatively easy-to-abate sectors such as energy. The second point is that solving the model with an increasing baseline level of emissions reveals that the investment pathway for achieving the 2 °C target exceeds that of the 1.7 °C investment path without the increasing baseline. This highlights the significant cost of building more "dirty" capital stocks that have to be replaced later down the line if climate goals are still to be realized; see Rozenberg et al. (2020) for a full treatment of how "dirty" capital is phased out (or stranded) as a result of decarbonization, and how these phase-out dynamics interact with policy instrument choice.

We note that our framework is idealized, and lends itself to a number of future studies. For example, we consider here the effects of moving from a deterministic framework with static and dynamic abatement costs (our abatement investment and "strawman" models, respectively) to one with uncertainty about the carbon budget, conditional on the deterministic costs being identical. One could carry out a future study where the expected costs between the two frameworks are identical, and analyze the dynamic differences when expected costs are set to be the same (i.e., when the dashed and solid lines in Figure 3 are required to be equal). Moreover, a number of additional considerations beyond those mentioned in this paper are necessary in formulating a robust climate policy suite. As an example, we compute a globally optimal investment strategy, and future work would consider the cost implications of sub-optimal policies, or heterogenous policy mixes (as explored in Clausing and Wolfram (2023)).

Open Research Section

All code and data used in this study are publicly available by two methods:

- The World Bank Reproducible Research Repository (Bauer et al., 2024a)
  - https://reproducibility.worldbank.org/index.php/catalog/129 (last accessed: 2/17/2025)
- The lead author's Github
  - https://github.com/adam-bauer-34/BMH-delayed-learning-reprod (last accessed: 2/17/2025)

Acknowledgments

The authors thank two anonymous reviewers for helpful comments and reviews that greatly improved the manuscript. The authors further thank Yongyang Cai, Kevin Carey, Simon Dietz, Roy Dong, Jen Morris, Sam Okullo, Cristian Proistosescu, Adrien Vogt-Schilb, and Gernot Wagner and participants of the American Geophysical Union's 2024 Fall Meeting for helpful discussions related to this work. Finally, the authors thank María Reyes Retana Torre for verifying the reproducibility of our results. The authors acknowledge financial support from the Climate Support Facility of the World Bank. AMB acknowledges financial support from a National Science Foundation Graduate Research Fellowship (Grant No. DGE 21-46756). The views expressed in this paper are the sole responsibility of the authors. They do not necessarily reflect the views of the World Bank, its executive directors, or the countries they represent. All remaining errors are the sole responsibility of the authors.

References

F. Ackerman, E. A. Stanton, and R. Bueno. Epstein–Zin Utility in DICE: Is Risk Aversion Irrelevant to Climate Policy? *Environmental and Resource Economics*, 56(1):73–84, Sept. 2013. ISSN 0924-6460, 1573-1502. doi: 10.1007/s10640-013-9645-z. URL http://link.springer.com/10.1007/s10640-013-9645-z.

M. Barnett, W. Brock, and L. Hansen. Climate Change Uncertainty Spillover in the Macroeconomy. Technical Report w29064, National Bureau of Economic Research, Cambridge, MA, July 2021. URL http://www.nber.org/papers/w29064.pdf.

L. Barrage and W. Nordhaus. Policies, projections, and the social cost of carbon: Results from the DICE-2023 model. *Proceedings of the National Academy of Sciences*, 121(13):e2312030121, Mar. 2024. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.2312030121. URL https://pnas.org/doi/10.1073/pnas.2312030121.

A. M. Bauer, F. McIsaac, and S. Hallegatte. Reproducibility package for "How Delayed Learning about Climate Uncertainty Impacts Decarbonization Investment Strategies", Apr. 2024a. URL https://reproducibility.worldbank.org/index.php/catalog/129.

A. M. Bauer, C. Proistosescu, and G. Wagner. Carbon Dioxide as a Risky Asset. *Climatic Change*, 177(5):72, May 2024b. ISSN 0165-0009, 1573-1480. doi: 10.1007/s10584-024-03724-3. URL https://link.springer.com/10.1007/s10584-024-03724-3.

P. Bolton, M. Després, L. A. Pereria da Silva, F. Samama, and R. Svartzman. *The green swan: Central banking and financial stability in the age of climate change*, volume 31. BIS Books, Bank for International Settlements, 2020. URL https://www.bis.org/publ/othp31.pdf.

Y. Cai. The Role of Uncertainty in Controlling Climate Change. In *Oxford Research Encyclopedia of Economics and Finance*. Oxford University Press, Feb. 2021. ISBN 9780190625979. doi: 10.1093/acrefore/9780190625979.013.573. URL https://oxfordre.com/economics/view/10.1093/acrefore/9780190625979.001.0001/acrefore-9780190625979-e-573.

Y. Cai and K. L. Judd. Stable and Efficient Computational Methods for Dynamic Programming. *Journal of the European Economic Association*, 8(2-3):626–634, Apr. 2010. ISSN 15424766. doi: 10.1111/j.1542-4774.2010.tb00532.x. URL https://academic.oup.com/jeea/article-lookup/doi/10.1111/j.1542-4774.2010.tb00532.x.

Y. Cai and K. L. Judd. Dynamic programming with Hermite approximation. *Mathematical Methods of Operations Research*, 81(3):245–267, June 2015. ISSN 1432-2994, 1432-5217. doi: 10.1007/s00186-015-0495-z. URL http://link.springer.com/10.1007/s00186-015-0495-z.

Y. Cai and T. S. Lontzek. The Social Cost of Carbon with Economic and Climate Risks. *Journal of Political Economy*, 127(6):2684–2734, Dec. 2019. ISSN 0022-3808, 1537-534X. doi: 10.1086/701890. URL https://www.journals.uchicago.edu/doi/10.1086/701890.

R. Calel, S. C. Chapman, D. A. Stainforth, and N. W. Watkins. Temperature variability implies greater economic damages from climate change. *Nature Communications*, 11(1):5028, Oct. 2020. ISSN 2041-1723. doi: 10.1038/s41467-020-18797-8. URL https://www.nature.com/articles/s41467-020-18797-8.

E. Campiglio, S. Dietz, and F. Venmans. Optimal Climate Policy as If the Transition Matters. Working Paper 10139, CESifo, Munich, 2022. URL https://www.cesifo.org/en/publications/2022/working-paper/optimal-climate-policy-if-transition-matters.

K. A. Clausing and C. Wolfram. Carbon Border Adjustments, Climate Clubs, and Subsidy Races When Climate Policies Vary. *Journal of Economic Perspectives*, 37(3):137–162, Sept. 2023. ISSN 0895-3309. doi: 10.1257/jep.37.3.137. URL https://www.aeaweb.org/articles?id=10.1257/jep.37.3.137.

Climate Policy Initiative. Global Landscape of Climate Finance 2023. Technical report, Climate Policy Initiative, Washington DC, Nov. 2023.

B. Crost and C. P. Traeger. Optimal CO_2 mitigation under damage risk valuation. *Nature Climate Change*, 4(7):631–636, July 2014. ISSN 1758-678X, 1758-6798. doi: 10.1038/nclimate2249. URL https://www.nature.com/articles/nclimate2249.

S. J. Davis and R. H. Socolow. Commitment accounting of CO_2 emissions. *Environmental Research Letters*, 9(8):084018, Aug. 2014. ISSN 1748-9326. doi: 10.1088/1748-9326/9/8/084018. URL https://iopscience.iop.org/article/10.1088/1748-9326/9/8/084018.

S. Dietz, J. Rising, T. Stoerk, and G. Wagner. Economic impacts of tipping points in the climate system. *PNAS*, 118(34):e2103081118, Aug. 2021. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.2103081118. URL https://pnas.org/doi/full/10.1073/pnas.2103081118.

A. K. Dixit and R. S. Pindyck. *Investment under uncertainty*. Princeton University Press, Princeton, N.J, 1994. ISBN 9780691034102.

M. A. Drupp, M. C. Freeman, B. Groom, and F. Nesje. Discounting Disentangled. *American Economic Journal: Economic Policy*, 10(4):109–134, Nov. 2018. ISSN 1945-7731, 1945-774X. doi: 10.1257/pol.20160240. URL https://pubs.aeaweb.org/doi/10.1257/pol.20160240.

M. T. Dvorak, K. C. Armour, D. M. W. Frierson, C. Proistosescu, M. B. Baker, and C. J. Smith. Estimating the timing of geophysical commitment to 1.5 and 2.0 °C of global warming. *Nature Climate Change*, 12(6):547–552, June 2022. ISSN 1758-678X, 1758-6798. doi: 10.1038/s41558-022-01372-y. URL https://www.nature.com/articles/s41558-022-01372-y.

L. G. Fitzpatrick and D. L. Kelly. Probabilistic Stabilization Targets. *Journal of the Association of Environmental and Resource Economists*, 4(2):611–657, June 2017. ISSN 2333-5955, 2333-5963. doi: 10.1086/691687. URL https://www.journals.uchicago.edu/doi/10.1086/691687.

M. Grubb. Technologies, energy systems and the timing of CO_2 emissions abatement. *Energy Policy*, 25(2):159–172, Feb. 1997. ISSN 03014215. doi: 10.1016/S0301-4215(96)00106-1. URL https://linkinghub.elsevier.com/retrieve/pii/S0301421596001061.

M. Ha-Duong, M. J. Grubb, and J.-C. Hourcade. Influence of socioeconomic inertia and uncertainty on optimal CO_2-emission abatement. *Nature*, 390(6657):270–273, Nov. 1997. ISSN 0028-0836, 1476-4687. doi: 10.1038/36825. URL http://www.nature.com/articles/36825.

H. Hotelling. The Economics of Exhaustible Resources. *Journal of Political Economy*, 39(2):137–175, 1931. ISSN 0022-3808. URL https://www.jstor.org/stable/1822328.

Intergovernmental Panel on Climate Change. *Climate Change 2021: The Physical Science Basis*. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge University Press, 2021.

Intergovernmental Panel on Climate Change. *Climate Change 2022: Mitigation of Climate Change*. Contribution of Working Group III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge University Press, 2022.

S. Jenkins, B. Sanderson, G. Peters, T. L. Frölicher, P. Friedlingstein, and M. Allen. The Multi-Decadal Response to Net Zero CO_2 Emissions and Implications for Emissions Policy. *Geophysical Research Letters*, 49(23):e2022GL101047, Dec. 2022. ISSN 0094-8276, 1944-8007. doi: 10.1029/2022GL101047. URL https://agupubs.onlinelibrary.wiley.com/doi/10.1029/2022GL101047.

D. L. Kelly and C. D. Kolstad. Bayesian learning, growth, and pollution. *Journal of Economic Dynamics and Control*, 23(4):491–518, Feb. 1999. ISSN 01651889. doi: 10.1016/S0165-1889(98)00034-7. URL https://linkinghub.elsevier.com/retrieve/pii/S0165188998000347.

M. J. Kotchen, J. A. Rising, and G. Wagner. The costs of "costless" climate mitigation. *Science*, 382(6674):1001–1003, Dec. 2023. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.adj2453. URL https://www.science.org/doi/10.1126/science.adj2453.

N. J. Leach, S. Jenkins, Z. Nicholls, C. J. Smith, J. Lynch, M. Cain, T. Walsh, B. Wu, J. Tsutsui, and M. R. Allen. FaIR v2.0.0: a generalized impulse response model for climate uncertainty and future scenario exploration. *Geoscientific Model Development*, 14(5):3007–3036, May 2021. ISSN 1991-9603. doi: 10.5194/gmd-14-3007-2021. URL https://gmd.copernicus.org/articles/14/3007/2021/.

D. Lemoine. The Climate Risk Premium: How Uncertainty Affects the Social Cost of Carbon. *Journal of the Association of Environmental and Resource Economists*, 8(1):27–57, Jan. 2021. ISSN 2333-5955, 2333-5963. doi: 10.1086/710667. URL https://www.journals.uchicago.edu/doi/10.1086/710667.

D. Lemoine and I. Rudik. Managing Climate Change Under Uncertainty: Recursive Integrated Assessment at an Inflection Point. *Annual Review of Resource Economics*, 9(1):117–142, Oct. 2017. ISSN 1941-1340, 1941-1359. doi: 10.1146/annurev-resource-100516-053516. URL http://www.annualreviews.org/doi/10.1146/annurev-resource-100516-053516.

D. Lemoine and C. P. Traeger. Economics of tipping the climate dominoes. *Nature Climate Change*, 6(5):514–519, May 2016. ISSN 1758-678X, 1758-6798. doi: 10.1038/nclimate2902. URL https://www.nature.com/articles/nclimate2902.

T. M. Lenton, H. Held, E. Kriegler, J. W. Hall, W. Lucht, S. Rahmstorf, and H. J. Schellnhuber. Tipping elements in the Earth's climate system. *Proceedings of the National Academy of Sciences*, 105(6):1786–1793, Feb. 2008. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.0705414105. URL https://pnas.org/doi/full/10.1073/pnas.0705414105.

R. E. Lucas. Adjustment Costs and the Theory of Supply. *Journal of Political Economy*, 75(4):321–334, 1967. ISSN 0022-3808. URL https://www.jstor.org/stable/1828594.

H. D. Matthews, N. P. Gillett, P. A. Stott, and K. Zickfeld. The proportionality of global warming to cumulative carbon emissions. *Nature*, 459(7248):829–832, June 2009. ISSN 0028-0836, 1476-4687. doi: 10.1038/nature08047. URL http://www.nature.com/articles/nature08047.

H. D. Matthews, K. Zickfeld, R. Knutti, and M. R. Allen. Focus on cumulative emissions, global carbon budgets and the implications for climate mitigation targets. *Environmental Research Letters*, 13(1): 010201, Jan. 2018. ISSN 1748-9326. doi: 10.1088/1748-9326/aa98c9. URL https://iopscience.iop.org/article/10.1088/1748-9326/aa98c9.

H. D. Matthews, K. B. Tokarska, J. Rogelj, C. J. Smith, A. H. MacDougall, K. Haustein, N. Mengis, S. Sippel, P. M. Forster, and R. Knutti. An integrated approach to quantifying uncertainties in the remaining carbon budget. *Communications Earth & Environment*, 2(1):7, Jan. 2021. ISSN 2662-4435. doi: 10.1038/s43247-020-00064-9. URL https://www.nature.com/articles/s43247-020-00064-9.

McKinsey & Company. *Pathways to a low-carbon economy: Version 2 of the global greenhouse gas abatement cost curve*. Stockholm, Sept. 2013. URL https://www.mckinsey.com/capabilities/sustainability/our-insights/pathways-to-a-low-carbon-economy.

J. F. Morris, V. Srikrishnan, M. D. Webster, and J. M. Reilly. Hedging Strategies: Electricity Investment Decisions under Policy Uncertainty. *The Energy Journal*, 39(1):101–122, Jan. 2018. ISSN 0195-6574, 1944-9089. doi: 10.5547/01956574.39.1.jmor. URL https://journals.sagepub.com/doi/10.5547/01956574.39.1.jmor.

J. F. Morris, J. M. Reilly, and Y.-H. H. Chen. Advanced technologies in energy-economy models for climate change assessment. *Energy Economics*, 80:476–490, May 2019. ISSN 01409883. doi: 10.1016/j.eneco.2019.01.034. URL https://linkinghub.elsevier.com/retrieve/pii/S0140988319300490.

M. Mussa. External and Internal Adjustment Costs and the Theory of Aggregate and Firm Investment. *Economica*, 44(174):163, May 1977. ISSN 00130427. doi: 10.2307/2553718. URL https://www.jstor.org/stable/10.2307/2553718?origin=crossref.

National Center for Energy Economics. *Supplementary Material for the Regulatory Impact Analysis for the Supplemental Proposed Rulemaking, "Standards of Performance for New, Reconstructed, and Modified Sources and Emissions Guidelines for Existing Sources: Oil and Natural Gas Sector Climate Review"*. U.S. Environmental Protection Agency, Washington D.C., Sept. 2022.

W. D. Nordhaus. Revisiting the social cost of carbon. *Proceedings of the National Academy of Sciences*, 114(7):1518–1523, Feb. 2017. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1609244114. URL https://pnas.org/doi/full/10.1073/pnas.1609244114.

S. J. Okullo. Determining the Social Cost of Carbon: Under Damage and Climate Sensitivity Uncertainty. *Environmental and Resource Economics*, 75(1):79–103, Jan. 2020. ISSN 0924-6460, 1573-1502. doi: 10.1007/s10640-019-00389-w. URL http://link.springer.com/10.1007/s10640-019-00389-w.

K. Rennert, F. Errickson, B. C. Prest, L. Rennels, R. G. Newell, W. Pizer, C. Kingdon, J. Wingenroth, R. Cooke, B. Parthum, D. Smith, K. Cromar, D. Diaz, F. C. Moore, U. K. Müller, R. J. Plevin, A. E. Raftery, H. Ševčíková, H. Sheets, J. H. Stock, T. Tan, M. Watson, T. E. Wong, and D. Anthoff. Comprehensive Evidence Implies a Higher Social Cost of CO2. *Nature*, Sept. 2022. ISSN 0028-0836, 1476-4687. doi: 10.1038/s41586-022-05224-9. URL https://www.nature.com/articles/s41586-022-05224-9.

K. Riahi, D. P. van Vuuren, E. Kriegler, J. Edmonds, B. C. O'Neill, S. Fujimori, N. Bauer, K. Calvin, R. Dellink, O. Fricko, W. Lutz, A. Popp, J. C. Cuaresma, S. Kc, M. Leimbach, L. Jiang, T. Kram, S. Rao, J. Emmerling, K. Ebi, T. Hasegawa, P. Havlik, F. Humpenöder, L. A. Da Silva, S. Smith, E. Stehfest, V. Bosetti, J. Eom, D. Gernaat, T. Masui, J. Rogelj, J. Strefler, L. Drouet, V. Krey, G. Luderer, M. Harmsen, K. Takahashi, L. Baumstark, J. C. Doelman, M. Kainuma, Z. Klimont, G. Marangoni, H. Lotze-Campen, M. Obersteiner, A. Tabeau, and M. Tavoni. The Shared Socioeconomic Pathways and their energy, land use, and greenhouse gas emissions implications: An overview. *Global Environmental Change*, 42:153–168, Jan. 2017. ISSN 09593780. doi: 10.1016/j.gloenvcha.2016.05.009. URL https://linkinghub.elsevier.com/retrieve/pii/S0959378016300681.

J. Rozenberg, A. Vogt-Schilb, and S. Hallegatte. Instrument choice and stranded assets in the transition to clean capital. *Journal of Environmental Economics and Management*, 100:102183, Mar. 2020. ISSN 00950696. doi: 10.1016/j.jeem.2018.10.005. URL https://linkinghub.elsevier.com/retrieve/pii/S009506961730623X.

S. C. Sherwood, M. J. Webb, J. D. Annan, K. C. Armour, P. M. Forster, J. C. Hargreaves, G. Hegerl, S. A. Klein, K. D. Marvel, E. J. Rohling, M. Watanabe, T. Andrews, P. Braconnot, C. S. Bretherton,

G. L. Foster, Z. Hausfather, A. S. Heydt, R. Knutti, T. Mauritsen, J. R. Norris, C. Proistosescu, M. Rugenstein, G. A. Schmidt, K. B. Tokarska, and M. D. Zelinka. An Assessment of Earth's Climate Sensitivity Using Multiple Lines of Evidence. *Reviews of Geophysics*, 58(4), Dec. 2020. ISSN 8755-1209, 1944-9208. doi: 10.1029/2019RG000678. URL https://onlinelibrary.wiley.com/doi/10.1029/2019RG000678.

C. J. Smith, P. M. Forster, M. Allen, N. Leach, R. J. Millar, G. A. Passerello, and L. A. Regayre. FAIR v1.3: a simple emissions-based impulse response and carbon cycle model. *Geoscientific Model Development*, 11(6):2273–2297, June 2018. ISSN 1991-9603. doi: 10.5194/gmd-11-2273-2018. URL https://gmd.copernicus.org/articles/11/2273/2018/.

United Nations Framework Convention on Climate Change. *Adoption of the Paris Agreement. I: Proposal by the President.* United Nations Office, Geneva, 2015.

A. Vogt-Schilb and S. Hallegatte. Marginal abatement cost curves and the optimal timing of mitigation measures. *Energy Policy*, 66:645–653, Mar. 2014. ISSN 03014215. doi: 10.1016/j.enpol.2013.11.045. URL https://linkinghub.elsevier.com/retrieve/pii/S030142151301152X.

A. Vogt-Schilb, G. Meunier, and S. Hallegatte. When starting with the most expensive option makes sense: Optimal timing, cost and sectoral allocation of abatement investment. *Journal of Environmental Economics and Management*, 88:210–233, Mar. 2018. ISSN 00950696. doi: 10.1016/j.jeem.2017.12.001. URL https://linkinghub.elsevier.com/retrieve/pii/S0095069617308392.

M. L. Weitzman. GHG Targets as Insurance Against Catastrophic Climate Damages. *Journal of Public Economic Theory*, 14(2):221–244, Mar. 2012. ISSN 1097-3923, 1467-9779. doi: 10.1111/j.1467-9779.2011.01539.x. URL https://onlinelibrary.wiley.com/doi/10.1111/j.1467-9779.2011.01539.x.

References from the Supporting Information

Agrawal, A., Diamond, S. & Boyd, S. (2019, July). Disciplined geometric programming. *Optimization Letters*, *13*(5), 961–976. doi: 10.1007/s11590-019-01422-z

Bellman, R. E. (1957). *Dynamic Programming.* Princeton University Press. doi: 10.1515/9781400835386

Cai, Y. (2019, October). Computational Methods in Environmental and Resource Economics. *Annual Review of Resource Economics*, *11*(1), 59–82. doi: 10.1146/annurev-resource-100518-093841

Cai, Y., & Judd, K. L. (2010, April). Stable and Efficient Computational Methods for Dynamic Programming. *Journal of the European Economic Association*, *8*(2-3), 626–634. doi: 10.1111/j.1542-4774.2010.tb00532.x

Cai, Y., & Judd, K. L. (2015, June). Dynamic programming with Hermite approximation. *Mathematical Methods of Operations Research*, *81*(3), 245–267. doi: 10.1007/s00186-015-0495-z

Diamond, S., & Boyd, S. (2016, January). CVXPY: A Python-embedded modeling language for convex optimization. *The Journal of Machine Learning Research*, *17*(1), 2909–2913.

Dvorak, M. T., Armour, K. C., Frierson, D. M. W., Proistosescu, C., Baker, M. B., & Smith, C. J. (2022, June). Estimating the timing of geophysical commitment to 1.5 and 2.0 °C of global warming. *Nature Climate Change*, *12*(6), 547–552. doi: 10.1038/s41558-022-01372-y

Gurobi Optimization, LLC. (2023). *Gurobi Optimizer Reference Manual.* Retrieved from `https://www.gurobi.com`

Intergovernmental Panel on Climate Change. (2021). *Climate change 2021: The physical science basis.* Cambridge University Press.

International Energy Agency. (2022, April). *Direct Air Capture 2022 – Analysis* (Tech. Rep.). International Energy Agency. Retrieved 2023-10-16, from `https://www.iea.org/reports/direct-air-capture-2022`

Judd, K. L. (1998). *Numerical methods in economics.* Cambridge, Mass: MIT Press.

Matthews, H. D., Tokarska, K. B., Rogelj, J., Smith, C. J., MacDougall, A. H., Haustein, K., ..., Knutti, R. (2021, January). An integrated approach to quantifying uncertainties in the remaining carbon budget. *Communications Earth & Environment*, *2*(1), 7. doi: 10.1038/s43247-020-00064-9

Riahi, K., van Vuuren, D. P., Kriegler, E., Edmonds, J., O'Neill, B. C., Fujimori, S., ... Tavoni, M. (2017, January). The Shared Socioeconomic Pathways and their energy, land use, and greenhouse gas emissions implications: An overview. *Global Environmental Change*, *42*, 153–168. doi: 10.1016/j.gloenvcha.2016.05.009

Terlouw, T., Treyer, K., Bauer, C., & Mazzotti, M. (2021, August). Life Cycle Assessment of Direct Air Carbon Capture and Storage with Low-Carbon Energy Sources. *Environmental Science & Technology*, *55*(16), 11397-11411. doi: 10.1021/acs.est.1c03263

Supplementary Information for "Decarbonization Investment Strategies in an Uncertain Climate"

Adam Michael Bauer*[1,2], Florent McIsaac[2], and Stéphane Hallegatte[2]

[1] *Department of Physics, University of Illinois Urbana-Champaign, 1110 W Green St Loomis Laboratory, Urbana, IL 61801*
[2] *World Bank Group, Washington DC*

Forthcoming in *Earth's Future*
4/8/2025

Contents of this file

1. Text S1 to S5
2. Figures S1 to S11
3. Tables S1 to S4

Introduction This Supplement includes sensitivity tests of the main results and details on our computational framework.

*Corresponding author: adammb4@illinois.edu

Text S1: Including direct air capture technologies. In this Section, we include direct air capture technologies (DAC) in our modeling framework. We treat DAC like an additional economic sector with zero annual emissions (i.e., $\bar{a}_{DAC} = 0$) and some abatement potential, that we call $Z > 0$ (in $GtCO_2$ yr^{-1}). DAC is unique in that its emissions rate is not the same as its abatement potential, as was assumed in every other sector.

Augmented model framework

A social planner tasked with decarbonization with DAC technologies at their disposal solves the following cost-minimization problem,

$$\min_{\{x_{i,t}\}_{t\in\mathcal{T},i\in\mathcal{I}}} \left[\sum_{t\in\mathcal{T}} \beta^t \sum_{i\in\mathcal{I}} c_i(x_{i,t})\right], \tag{0.1}$$

$$\begin{aligned}
\text{Subject to}: \; a_{i,t+1} &= a_{i,t} + \Delta t \left(x_{i,t} - \delta_i a_{i,t}\right), \\
\psi_{t+1} &= \psi_t + \Delta t \left(\sum_{i\in\mathcal{I}} \left(\bar{a}_i - a_{i,t}\right)\right), \\
0 &\leq a_{i,t} \leq \bar{a}_i, \\
0 &\leq \psi_t \leq B, \\
0 &\leq a_{DAC,t} \leq Z, \\
a_{i,0}, &\psi_0 \text{ given.}
\end{aligned}$$

Extending this formulation to a model with uncertainty about the RCB is straightforward. We take the marginal abatement cost of DAC from International Energy Agency (2022) as \$100–\$300 tCO_2^{-1}. In each simulation we show below, we assume a marginal abatement cost of \$200 tCO_2^{-1}. Any lower marginal cost we would judge to be unrealistically cheap, and any higher cost causes model behavior to be completely unchanged in all of our simulations. We translate this marginal abatement cost to a marginal investment cost using the approach in the main text. We assume the capital depreciation rate of DAC is 5% yr^{-1} and that, at maximum, the abatement potential from DAC technologies is 5 $GtCO_2$ yr^{-1} (Terlouw et al., 2021).

Results

We find that the impact of including DAC in our model for abatement investment is consistent across temperature targets, see Figure 1. For both temperature targets we explored, we find that there is practically no change in the value of learning; in other words, the investment strategies with and without DAC are essentially identical. This owes to the carbon price being too low to induce the social planner to invest heavily in DAC; it is cheaper to simply abate emissions at their source, given the relatively cheap cost of abatement investment in other sectors compared to DAC.

The marginal investments in DAC allows the social planner to invest slightly less resources in hard-to-abate sectors, see Figure 2. We find that most sectors – waste, forestry, transport, energy, and buildings – have minimal changes to their investment pathway when DAC technologies are introduced. However, hard-to-abate sectors, such as industry and agriculture, have non-trivial reductions in investment early on; these resources have been redirected to DAC. The additional resources dedicated to DAC serves to decrease the steady state level of investment in agriculture, hence the blue line in Figure 2**d** being

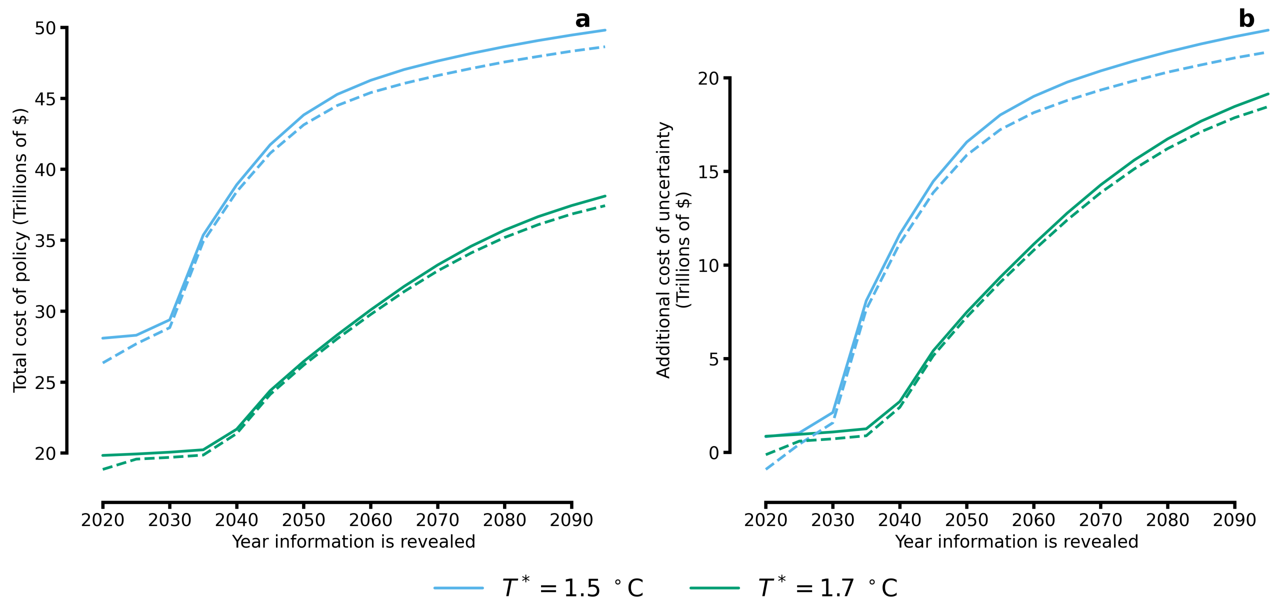

Figure 1: Effect of uncertainty and delayed learning on aggregate policy cost including direct air capture technologies. In panel **a**, we show the total discounted policy cost for each temperature target (see the legend) for the abatement investment model without DAC (solid lines) and abatement investment model with direct air capture technologies (dashed lines). Panel **b** is as **a**, net of the cost of the corresponding policy without uncertainty.

below the gold line. We conclude that the social planner uses DAC as a way to offset the costs in hard-to-abate sectors only, while relatively easy-to-abate sectoral investment is unchanged with and without DAC technologies. Note one caveat to these conclusions would be that, if the cost of other technologies are higher than we consider here while the cost of DAC is held constant, we would expect that DAC would have a more prominent role because it would offer more cost-savings than presented here.

Text S2: Increasing emissions baselines

In this Section, we allow for the baseline emissions to increase. We assume that each sector grows at some annual rate, $0 < \zeta \leq 1$, until 2050, after which emissions are kept at level. We can therefore write the emissions baseline in each sector as,

$$\bar{a}_{i,t} = \begin{cases} \bar{a}_{i,2020}(1+\zeta)^t, & t < 30, \\ \bar{a}_{i,2020}(1+\zeta)^{30}, & t \geq 30, \end{cases} \tag{0.2}$$

where $\bar{a}_{i,2020}$ is the 2020 emissions levels quoted in Table 1 in the main text. We choose $\zeta = 1.5\%$ yr^{-1} in the results that follow; this puts 2050 emissions on par with SSP3–7.0 (Riahi et al., 2017). The model equations with a time dependent emissions baseline are straightforward to formulate based off our discussion in §2 in the main text by simply replacing $\bar{a}_i$ with $\bar{a}_{i,t}$.

Results

After incorporating an increasing emissions baseline, we find that the cost of achieving a given temperature target unilaterally increases; this owes to the assumption of more 'dirty' capital stocks being

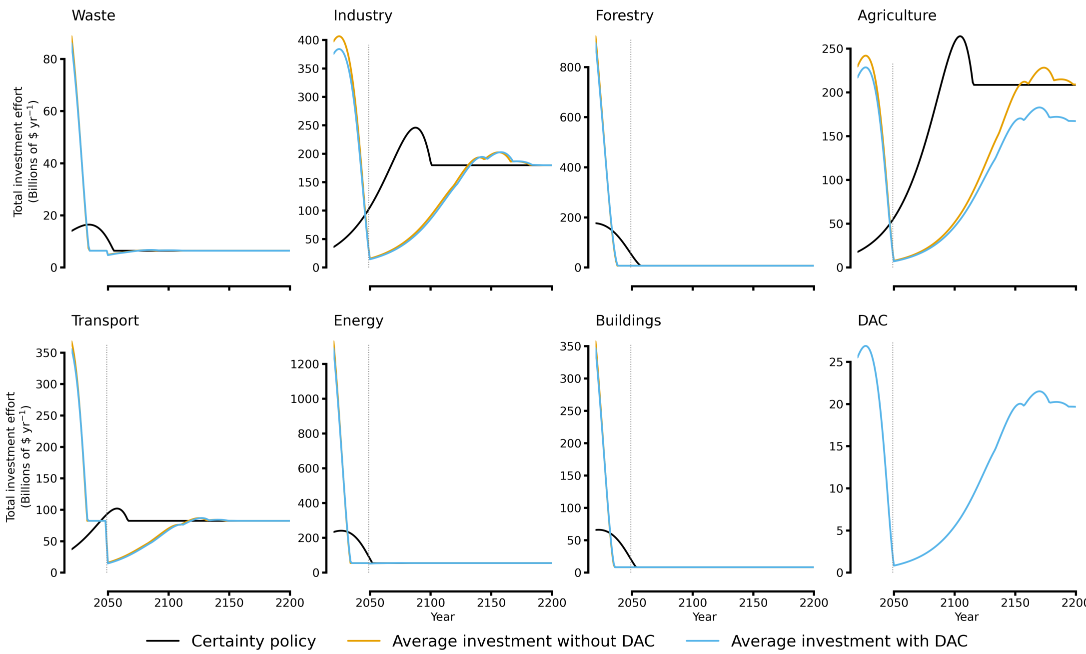

Figure 2: Impact of delayed learning on sectoral allocation of abatement investment when direct air capture technologies are present. In each panel we show the average investment effort in each sector (see the titles). The black line is the certainty policy, the gold line is the average investment when the carbon budget is learned in 2050 without direct air capture technologies (DAC), and the blue line is the average investment under the same experimental setup as the gold line but with DAC included. Note these investment pathways are for the 1.7 °C temperature target.

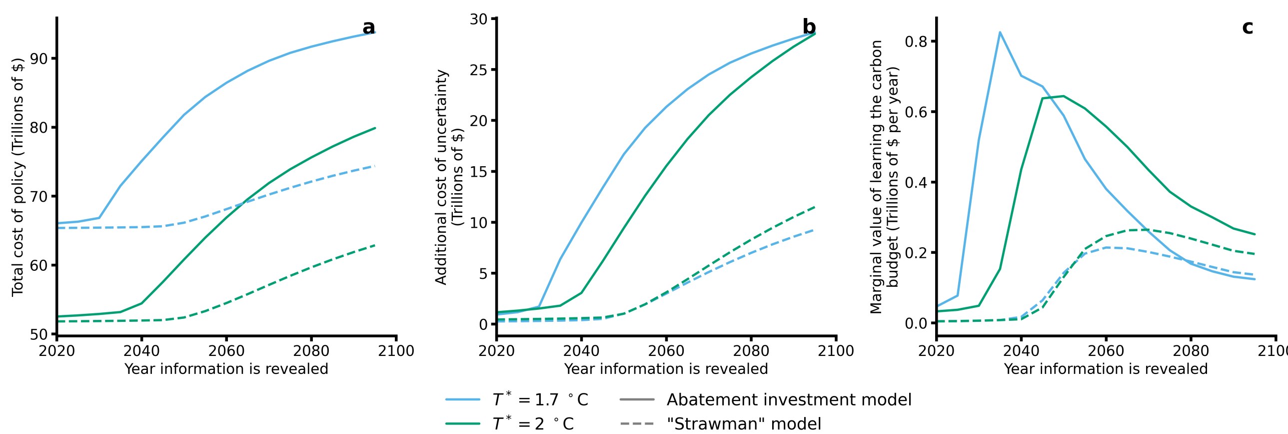

Figure 3: Effect of uncertainty and delayed learning on aggregate policy cost, growing emissions baseline. In Panel **a**, we show the total discounted policy cost for each temperature target (see the legend) for the abatement investment model that includes the impact of adjustment costs (solid lines) and "strawman" model, with does not include these factors (dashed lines). Panel **b** is as **a**, net of the cost of the corresponding policy without uncertainty. Panel **c** shows the marginal cost of delaying learning by one additional year.

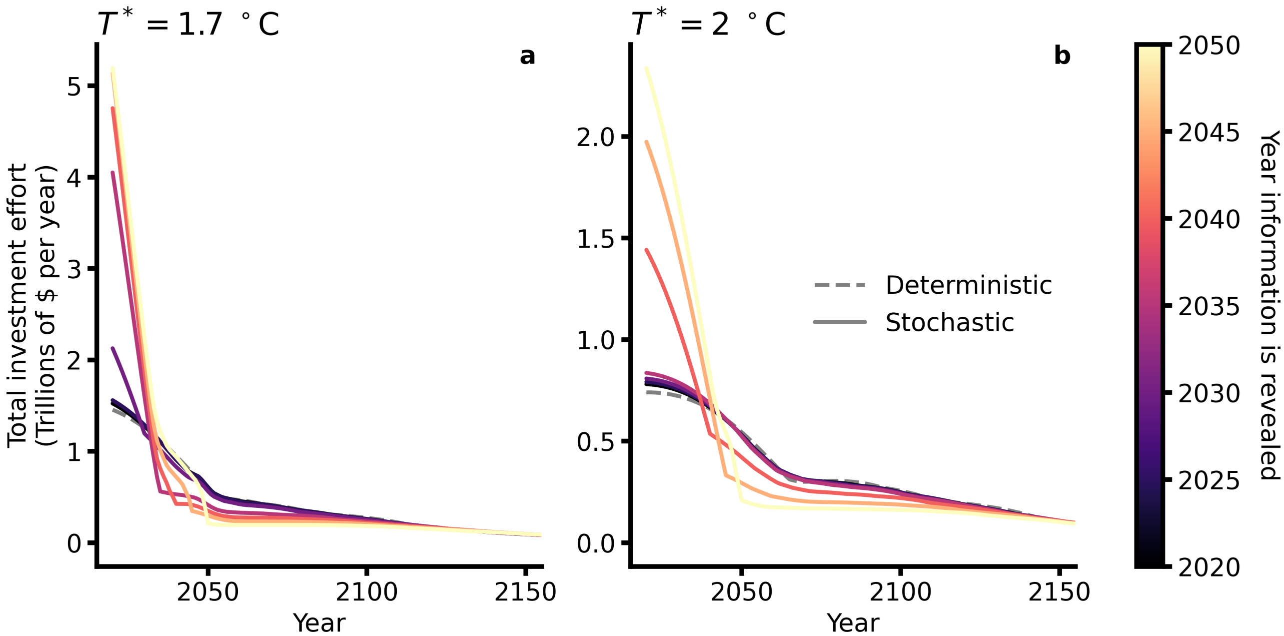

Figure 4: Effect of uncertainty and delayed learning on the temporal distribution of spending, growing emissions baseline. Panel **a** shows the average, total investment effort, $\sum_{i \in \mathcal{I}} c_i(x_{i,t})$, for a 1.7 °C temperature target (solid lines), where the colors represent different years where information about the carbon budget is revealed. The deterministic policy is shown in the grey, dashed line. Panel **b** is as **a**, but for the 2 °C temperature target.

turned over to 'clean' capital stocks, which of course come at a price. The key conclusions of our paper, however, still hold: learning later about the carbon budget causes the overall cost of policy to increase and spending is relatively front-loaded when compared to the certainty policy (see Figures 3 and 4).

One important note is that, when an increasing emissions baseline in considered (i.e., we consider a world trying to decarbonize while some sectors are building new 'dirty' capital simultaneously), the investment cost of decarbonizing and meeting the 2 °C policy exceeds the cost of decarbonizing a world with a constant emissions baseline that meets the 1.7 °C policy. This emphasizes the significant cost-savings of decarbonization that simply arise by not building any new 'dirty' capital stocks, such as coal-fired power plants or inefficient buildings.

Text S3: Alternative calibrations

High cost, linear calibration

In this subsection, we use the high cost, linear calibration of the marginal abatement cost curves (blue lines in Figure 1 in main text) in our optimization scheme. We find that while the quantitative implications of our analysis are altered, the qualitative results are robust to this alternative calibration.

Calibration

We use the same scheme as outlined in §2 in the main text to calibrate the abatement investment and "strawman" model. We use the approach outlined in the main text to ensure the relative costs between each sector are equal in both models, and that the total cost of each policy without uncertainty is equal. The numerical values for the marginal abatement cost and marginal investment cost coefficients

Table 1: Calibration parameters for the "strawman" and abatement investment models for each temperature target using the high cost, linear calibration.

	$T^* = 1.7$ °C			
$t_f = 300$ yr	$r = 2$ % yr^{-1}	$\bar{B} = 770$ $GtCO_2$	$\sigma_B = 220$ $GtCO_2$	
Sector	$\bar{\gamma}$ $\left[\frac{\$\ tCO_2^{-1}}{GtCO_2\ yr^{-1}}\right]$	$\bar{a}$ [$GtCO_2$ yr^{-1}]	δ [% yr^{-1}]	$\bar{c}$ $\left[\frac{\$\ tCO_2^{-1}}{GtCO_2\ yr^{-3}}\right]$
Waste	243.9	0.82	3.3	116240
Industry	36.6	5.47	4	17420
Forestry	24.24	8.25	0.8	11550
Agriculture	24.57	4.07	5	11710
Transport	26.7	3.74	6.7	12740
Energy	16.68	11.99	2.5	7950
Buildings	62.5	3.2	1.7	29790
	$T^* = 2$ °C			
$t_f = 300$ yr	$r = 2$ % yr^{-1}	$\bar{B} = 1340$ $GtCO_2$	$\sigma_B = 380$ $GtCO_2$	
Sector	$\bar{\gamma}$ $\left[\frac{\$\ tCO_2^{-1}}{GtCO_2\ yr^{-1}}\right]$	$\bar{a}$ [$GtCO_2$ yr^{-1}]	δ [% yr^{-1}]	$\bar{c}$ $\left[\frac{\$\ tCO_2^{-1}}{GtCO_2\ yr^{-3}}\right]$
Waste	243.9	0.82	3.3	133400
Industry	36.6	5.47	4	20000
Forestry	24.24	8.25	0.8	13260
Agriculture	24.57	4.07	5	13440
Transport	26.7	3.74	6.7	14620
Energy	16.68	11.99	2.5	9120
Buildings	62.5	3.2	1.7	34180

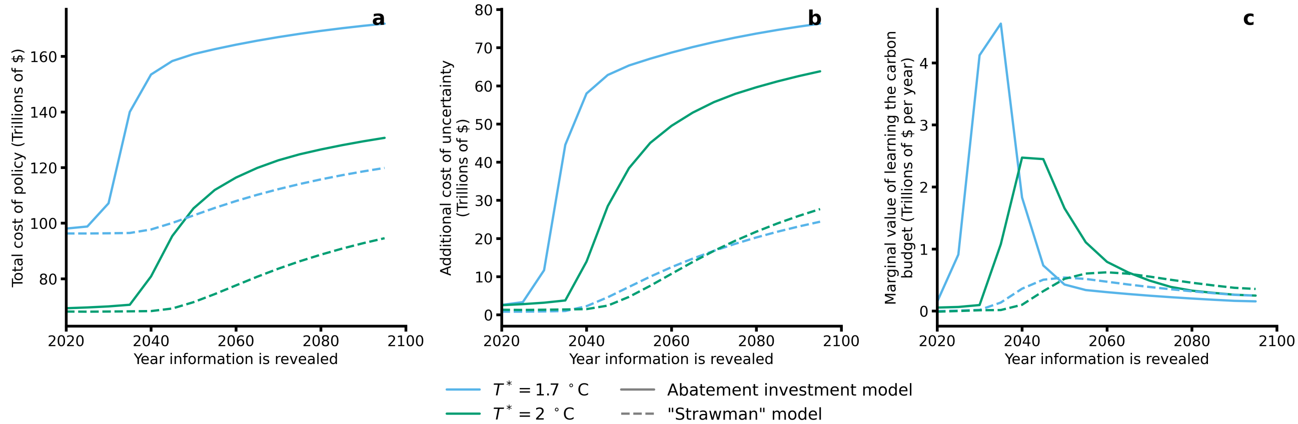

Figure 5: Effect of uncertainty and delayed learning on aggregate policy cost, high cost, linear calibration. In Panel **a**, we show the total discounted policy cost for each temperature target (see the legend) for the abatement investment model that includes the impact of adjustment costs (solid lines) and "strawman" model, with does not include these factors (dashed lines). Panel **b** is as **a**, net of the cost of the corresponding policy without uncertainty. Panel **c** shows the marginal cost of delaying learning by one additional year.

are shown in Table 1.

Results

In the high cost, linear calibration, we find that the key implications of our paper hold: learning later about the carbon budget causes the overall cost of policy to increase and spending is relatively front-loaded when compared to the equivalent certainty policy (see Figures 5 and 6). The only differences with the results of the low cost, linear calibration are quantitative, which is obviously the case, given that costs have increased in each sector significantly.

Nonlinear calibration

In this subsection, we use the nonlinear calibration of the marginal abatement cost curves (black lines in Figure 1 in main text) in our optimization scheme. Similar to the high cost, linear calibration shown above, using the nonlinear calibration causes our quantitative results to change, while leaving our qualitative conclusions about the effects of delayed uncertainty on policy unchanged.

Calibration

In our nonlinear calibration we fit a quadratic curve to the marginal abatement costs in Figure 1 to all sectors except agriculture, which is obviously linear. Therefore, we use

$$\gamma_i'(a_{i,t}) = \bar{\gamma}_i a_{i,t}^{\xi_i}, \tag{0.3}$$

in conjunction with relative costs being equal to calibrate our "strawman" and abatement investment models. See Table 2 for the numerical values for this calibration.

Results

In the nonlinear calibration, we find that the cost of achieving a given temperature target increases across each target relative to the low cost, linear calibration; this owes to the nonlinear case being

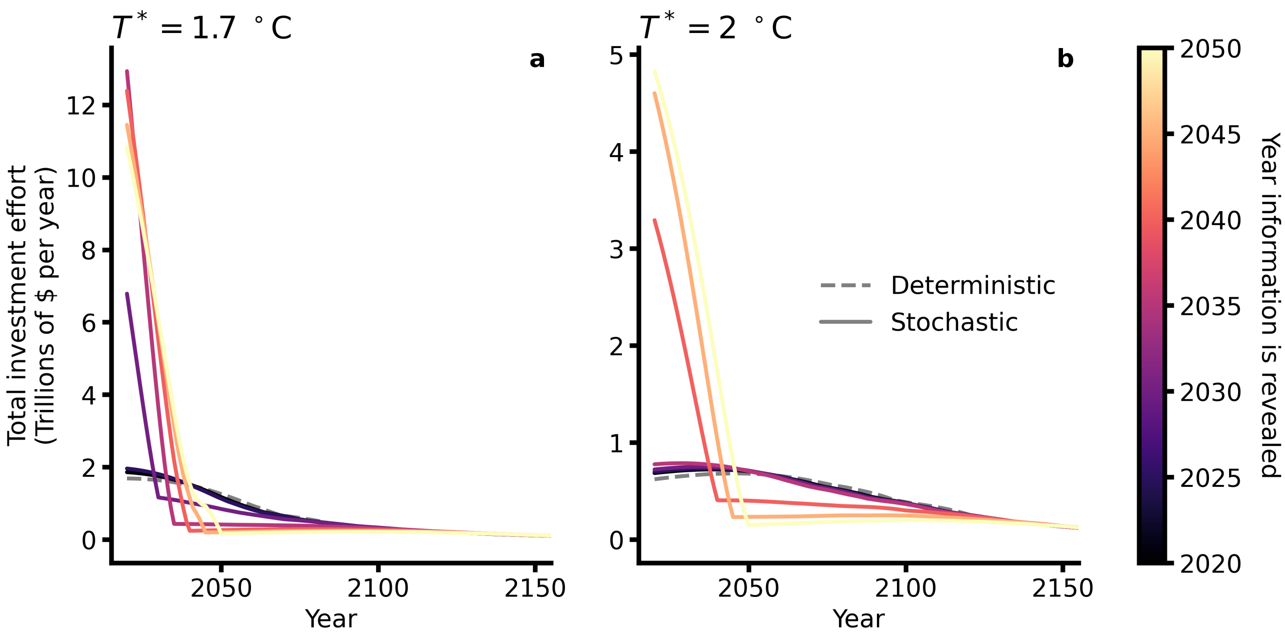

Figure 6: Effect of uncertainty and delayed learning on the temporal distribution of spending, high cost, linear calibration. Panel **a** shows the average, total investment effort, $\sum_{i \in \mathcal{I}} c_i(x_{i,t})$, for a 1.7 °C temperature target (solid lines), where the colors represent different years where information about the carbon budget is revealed. The deterministic policy is shown in the grey, dashed line. Panel **b** is as **a**, but for the 2 °C temperature target.

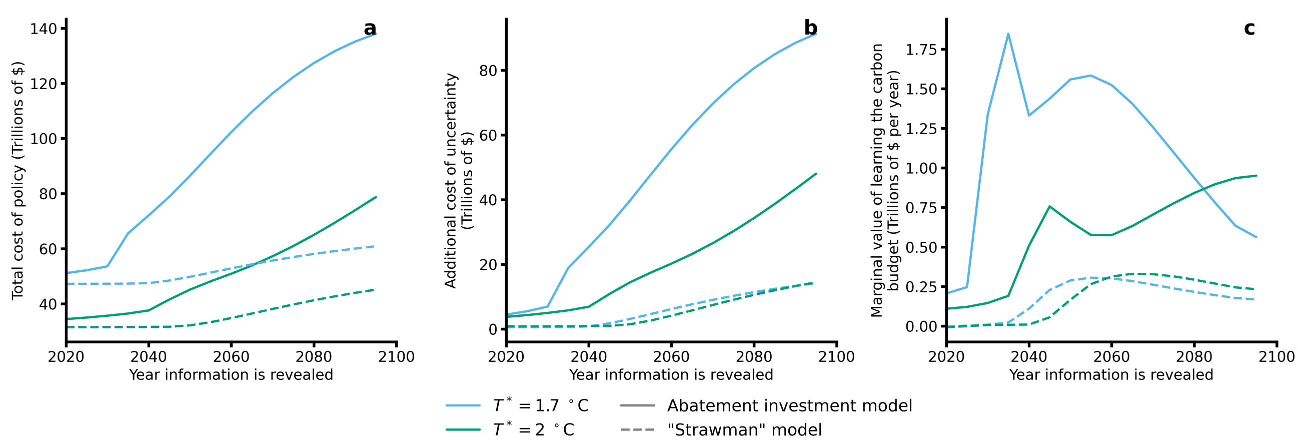

Figure 7: Effect of uncertainty and delayed learning on aggregate policy cost, nonlinear calibration. In Panel **a**, we show the total discounted policy cost for each temperature target (see the legend) for the abatement investment model that includes the impact of adjustment costs (solid lines) and "strawman" model, with does not include these factors (dashed lines). Panel **b** is as **a**, net of the cost of the corresponding policy without uncertainty. Panel **c** shows the marginal cost of delaying learning by one additional year.

Table 2: Calibration parameters for the "strawman" and abatement investment models for each temperature target using the nonlinear calibration.

$T^* = 1.7$ °C					
$t_f = 300$ yr	$r = 2$ % yr^{-1}	$\bar{B} = 770$ $GtCO_2$	$\sigma_B = 220$ $GtCO_2$		
Sector	$\bar{\gamma}$ $\left[\frac{\$\ tCO_2^{-1}}{GtCO_2^{\xi}\ yr^{-\xi}}\right]$	$\bar{a}$ [$GtCO_2$ yr^{-1}]	δ [% yr^{-1}]	$\bar{c}$ $\left[\frac{\$\ tCO_2^{-1}}{GtCO_2^{\xi}\ yr^{-3\xi}}\right]$	ξ [–]
Waste	242.1	0.82	3.3	976320	2
Industry	5.03	5.47	4	20290	2
Forestry	2.54	8.25	0.8	10250	2
Agriculture	25.4	4.07	5	102510	1
Transport	4.83	3.74	6.7	19460	2
Energy	1.04	11.99	2.5	4180	2
Buildings	19.4	3.2	1.7	78220	2

$T^* = 2$ °C					
$t_f = 300$ yr	$r = 2$ % yr^{-1}	$\bar{B} = 1340$ $GtCO_2$	$\sigma_B = 380$ $GtCO_2$		
Sector	$\bar{\gamma}$ $\left[\frac{\$\ tCO_2^{-1}}{GtCO_2^{\xi}\ yr^{-\xi}}\right]$	$\bar{a}$ [$GtCO_2$ yr^{-1}]	δ [% yr^{-1}]	$\bar{c}$ $\left[\frac{\$\ tCO_2^{-1}}{GtCO_2^{\xi}\ yr^{-3\xi}}\right]$	ξ [–]
Waste	242.1	0.82	3.3	1614400	2
Industry	5.03	5.47	4	33550	2
Forestry	2.54	8.25	0.8	16950	2
Agriculture	25.4	4.07	5	169510	1
Transport	4.83	3.74	6.7	32180	2
Energy	1.04	11.99	2.5	6920	2
Buildings	19.4	3.2	1.7	129020	2

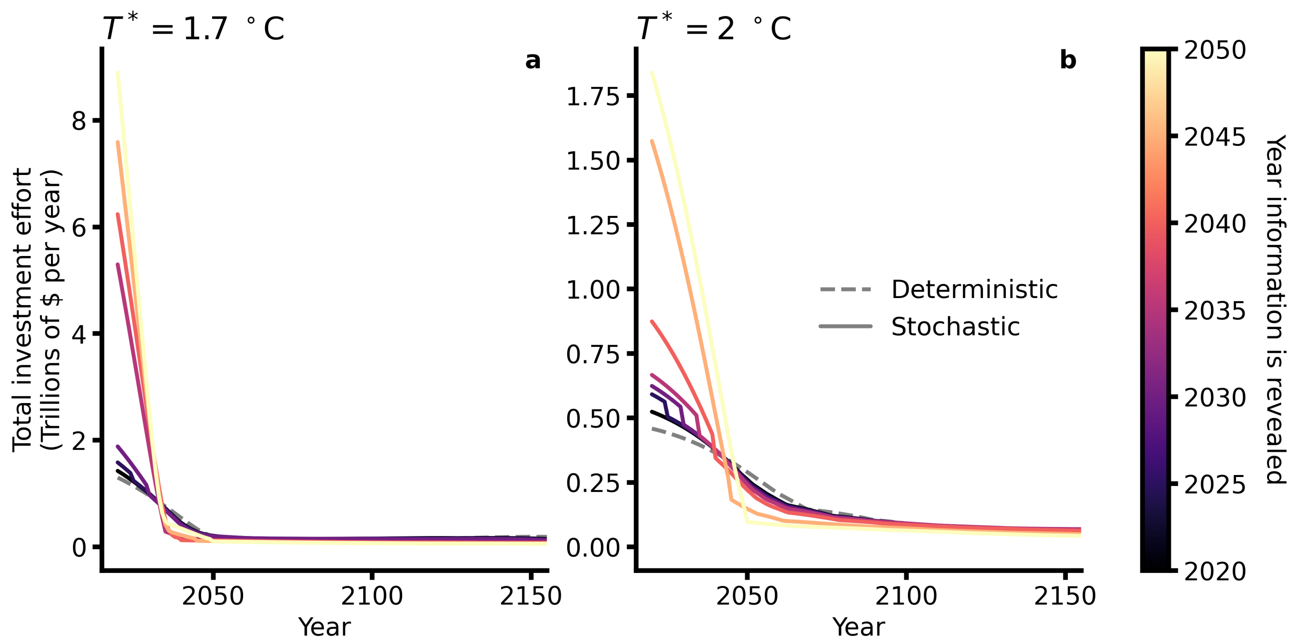

Figure 8: Effect of uncertainty and delayed learning on the temporal distribution of spending, nonlinear calibration. Panel **a** shows the average, total investment effort, $\sum_{i\in\mathcal{I}} c_i(x_{i,t})$, for a 1.7 °C temperature target (solid lines), where the colors represent different years where information about the carbon budget is revealed. The deterministic policy is shown in the grey, dashed line. Panel **b** is as **a**, but for the 2 °C temperature target.

significantly more expensive than the low cost, linear calibration. The key conclusions of our paper, however, still hold: learning later about the carbon budget causes the overall cost of policy to increase and spending is relatively front-loaded when compared to the equivalent certainty policy (see Figures 7 and 8).

1.5 °C temperature target

In this subsection, we explore our model implications when we use a remaining carbon budget that implies warming stays at or below 1.5 °C. We emphasize that our approach differs from the notion of "keeping 1.5 °C possible", that is, formulating a global investment strategy that maintains a nonzero probability of 1.5 °C warming by 2100. Our model goes much further, and enforces that the planner must formulate a strategy that is guaranteed to keep warming below 1.5 °C. Hence, in the results below we find that the cost of achieving the 1.5 °C target is incredibly high even for our lowest cost calibration, because the policy is robust to even very tiny carbon budgets, which causes the cost of policy to skyrocket. For context, the worst case carbon budget we sample is just higher than 100 $GtCO_2$, which would give the social planner just north of two years to decarbonize in the absence of any abatement. Furthermore, we cannot rule out zero or negative RCBs for the 1.5 °C temperature target (Matthews et al., 2021; Dvorak et al., 2022), implying that even our sampling of the RCB is optimistic from a geophysical perspective. This framing provides context to our results: they should not be interpreted as "ruling out" limiting warming to 1.5 °C, but rather should demonstrate the urgency of this target in terms of global investment needs.

Calibration

Table 3: Calibration parameters for the "strawman" and abatement investment models for the 1.5 °C temperature target using the low cost, linear calibration.

	$T^* = 1.5$ °C			
$t_f = 300$ yr	$r = 2$ % yr^{-1}	$\bar{B} = 440$ $GtCO_2$	$\sigma_B = 180$ $GtCO_2$	
Sector	$\bar{\gamma}$ $\left[\frac{\$ \text{ tCO}_2^{-1}}{\text{GtCO}_2 \text{ yr}^{-1}}\right]$	$\bar{a}$ [$GtCO_2$ yr^{-1}]	δ [% yr^{-1}]	$\bar{c}$ $\left[\frac{\$ \text{ tCO}_2^{-1}}{\text{GtCO}_2 \text{ yr}^{-3}}\right]$
Waste	40.82	0.82	3.3	14440
Industry	17.54	5.47	4	6200
Forestry	7.12	8.25	0.8	2520
Agriculture	23.53	4.07	5	8320
Transport	6.12	3.74	6.7	2160
Energy	2.82	11.99	2.5	998
Buildings	12.99	3.2	1.7	4600

We calibrate our 1.5 °C target model using the low cost, linear calibration for the marginal abatement costs, i.e., the gold lines in Figure 1 in the main text. We follow the same prescription as laid out in the main text, see Table 3.

Results

We find that for the 1.5 °C temperature target, the main qualitative findings of our paper hold, as was the case in each other alternative calibration; see Figures 9 and 10).

However, we do find some differences with other temperature targets. For example, the overall cost of policy as a function of learning time does not follow an "S"-curve as much as an inverted "L"-curve, owing to even a marginal delay in learning about the RCB causing a large increase in cost. This is because the worst-case carbon budget is as low as 100 $GtCO_2$, causing an immediate surge in precautionary investment. However, overall costs level out as the budget is learned later in the century, consistent with other temperature targets.

Furthermore, we find that the marginal cost of delay for the 1.5 °C target is more than twice that of the 1.7 °C or 2 °C targets, again owing to the urgency of the worst-case scenario.

Text S4: Computational details

Truncated Gauss-Hermite approximation In our formulation of the RCB, we use a truncated normal distribution, as negative values of the RCB make our optimization problem infeasible. Geophysical modeling suggests that the RCB distribution is approximately normal, justifying our approach (Intergovernmental Panel on Climate Change, 2021; Dvorak et al., 2022).

To approximate this distribution, we use Gauss-Hermite quadrature nodes (Cai and Judd, 2010) transformed to a truncated normal distribution following the approach of Cai and Judd (2015). We do so using the following prescription. For a random variable $x \sim \mathcal{N}(\bar{x}, \sigma_x)$, we can generate a random draw x^* using

$$x^* = \bar{x} + \vartheta \sigma_x, \tag{0.4}$$

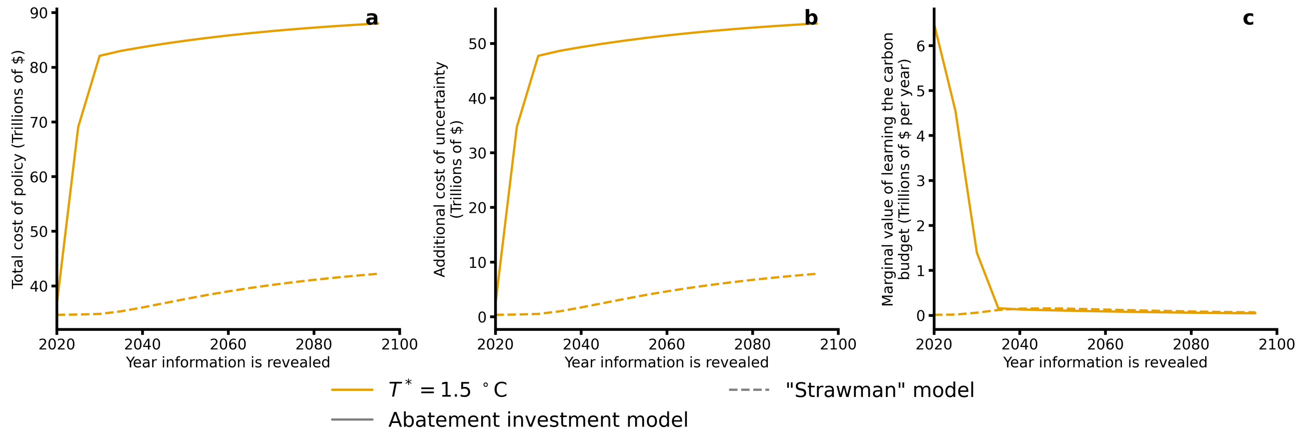

Figure 9: Effect of uncertainty and delayed learning on aggregate policy cost, $T^* = 1.5$C. In Panel **a**, we show the total discounted policy cost for each temperature target (see the legend) for the abatement investment model that includes the impact of adjustment costs (solid lines) and "strawman" model, with does not include these factors (dashed lines). Panel **b** is as **a**, net of the cost of the corresponding policy without uncertainty. Panel **c** shows the marginal cost of delaying learning by one additional year.

where $\vartheta \sim \mathcal{N}(0, 1)$. Therefore, formulating the Gauss-Hermite quadrature approximation for a truncated normal distribution, x, equates to formulating the Gauss-Hermite quadrature approximation for ϑ.

Therefore, consider the random variable ϑ. We can transform ϑ to a bounded random variable using the sigmoid function

$$\Psi = \frac{1 - e^{-\kappa\vartheta}}{1 + e^{-\kappa\vartheta}} \Upsilon, \tag{0.5}$$

where $\kappa, \Upsilon > 0$. One can immediately see that Ψ has zero mean, is symmetric around the mean, and is bounded between $(-\Upsilon, \Upsilon)$. We can choose a number such that our distribution is bounded between $\pm\Upsilon$ and select a κ such that Ψ has unit variance.[1] An intuitive way to select Υ is for Υ to be the number of standard deviations away from $\bar{x}$ one wishes to sample; for example, if x_1 is the first percentile of the distribution of x, then

$$\Upsilon = \frac{\bar{x} - x_1}{\sigma_x}. \tag{0.6}$$

Assuming one has calculated κ for their choice of Υ, a draw x^*_{trunc} can be generated using

$$x^*_{trunc} = \bar{x} + \Upsilon \left(\frac{1 - e^{-\kappa\vartheta}}{1 + e^{-\kappa\vartheta}} \right) \sigma_x. \tag{0.7}$$

It remains to apply the Gauss-Hermite approximation. We can approximate the distribution of ϑ using N Gauss-Hermite nodes with locations $\{\vartheta_1, ..., \vartheta_N\}$, where ϑ_i is the i^{th} root of the physicist's Hermite polynomials (Cai and Judd, 2010). Therefore, the locations of the Gauss-Hermite quadrature nodes

[1] This is done using a root-finding algorithm that can be found in the reproducibility package for this paper.

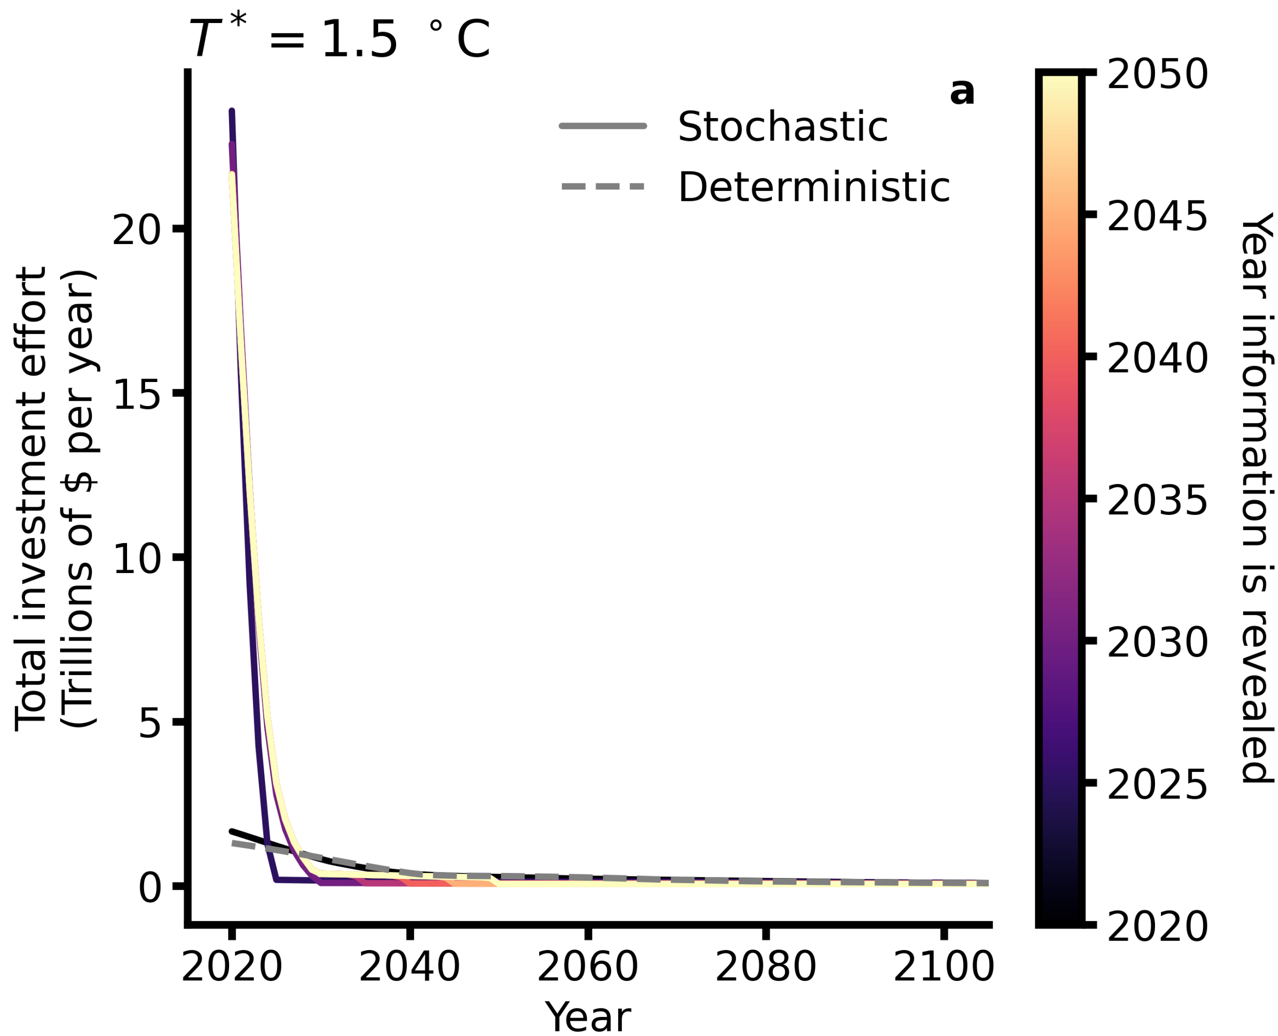

Figure 10: Effect of uncertainty and delayed learning on the temporal distribution of spending, $T^* = 1.5$ **°C.** Panel **a** shows the average, total investment effort, $\sum_{i \in \mathcal{I}} c_i(x_{i,t})$, for a 1.5 °C temperature target (solid lines), where the colors represent different years where information about the carbon budget is revealed. The deterministic policy is shown in the grey, dashed line.

Table 4: Truncated Gauss-Hermite parameters for each temperature target.

T^*	Υ	κ
1.5	1.8809	1.04536
1.7	2.3254	1.04534
2	2.3262	1.04505

for the truncated normal distribution are given by

$$x_{i,trunc} = \bar{x} + \Upsilon \left(\frac{1 - e^{-\kappa \vartheta_i}}{1 + e^{-\kappa \vartheta_i}} \right) \sigma_x, \tag{0.8}$$

as originally desired.

We choose Υ to be the 1^{st} percentile for the RCB distribution for the 1.7 °C and 2 °C targets, and the 3^{rd} percentile for the 1.5 °C target. We give the numerical values for Υ and κ in each case in Table 4.

Solving the model We solve our model using CVXPY, a domain-specific language written in Python and Julia for convex optimization problems (Diamond and Boyd, 2016). CVXPY uses disciplined geometric programming to interpret model equations and constraints and cast them into generic forms that can be used by open-source solvers (Agrawal et al., 2019). We use the GUROBI solver (Gurobi Optimization, LLC, 2023), which can be accessed for free by academics and researchers, that solves the base model in under 10^{-2} seconds and models with uncertainty in the range of 10^{-1} seconds, depending on how many quadrature nodes we use or samples we draw from the RCB distribution. Using this software allows us to circumvent solving the Bellman equation (Bellman, 1957) directly using techniques such as value function iteration (Judd, 1998), which can be computationally challenging (Cai, 2019), especially with high-dimensional state spaces and constraints on each state variable as is present in our model setup. Note that we use Gauss-Hermite quadrature methods to evaluate the expectation operator in our objective functions, and use 8 approximation nodes throughout (Cai, 2019); only in generating Figure 5 in the main text do we take a large number of samples (500) from the RCB distribution to solve the model for illustrative purposes. The results of this study were checked to be the same using both methods.

Text S5: Carbon prices under different cost calibrations

We show in Figure 11 how the carbon price changes with the learning date for all of our cost calibrations. Note in particular how calibrations with higher marginal investment costs (i.e., the high-linear and nonlinear calibrations) have higher carbon prices than simulations using the low-linear calibration as shown in the main text. The overall impact of uncertainty on policy costs and temporal allocation of investment is robust to these different carbon prices, as shown in Text S3 above.

References

A. Agrawal, S. Diamond, and S. Boyd. Disciplined geometric programming. *Optimization Letters*, 13(5):961–976, July 2019. ISSN 1862-4472, 1862-4480. doi: 10.1007/s11590-019-01422-z. URL

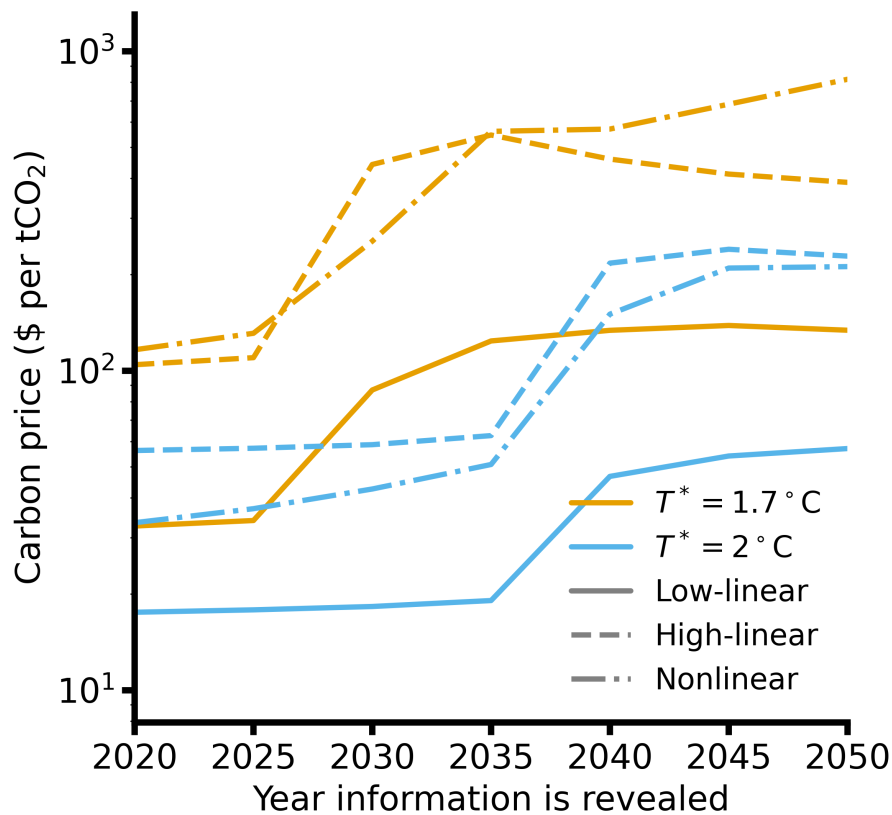

Figure 11: Carbon prices as a function of learning date. We show the logarithm of the average carbon price as a function of the learning date for each cost calibration: low-linear (solid lines), high-linear (dashed lines), and nonlinear (dashdot lines). For each cost calibration, we show the carbon prices for both temperature targets we consider in the main text: 1.7 °C (yellow lines) and 2 °C (blue lines).

http://link.springer.com/10.1007/s11590-019-01422-z.

R. E. Bellman. *Dynamic Programming*. Princeton University Press, 1957. ISBN 9781400835386. doi: 10.1515/9781400835386. URL https://www.degruyter.com/document/doi/10.1515/9781400835386/html.

Y. Cai. Computational Methods in Environmental and Resource Economics. *Annual Review of Resource Economics*, 11(1):59–82, Oct. 2019. ISSN 1941-1340, 1941-1359. doi: 10.1146/annurev-resource-100518-093841. URL https://www.annualreviews.org/doi/10.1146/annurev-resource-100518-093841.

Y. Cai and K. L. Judd. Stable and Efficient Computational Methods for Dynamic Programming. *Journal of the European Economic Association*, 8(2-3):626–634, Apr. 2010. ISSN 15424766. doi: 10.1111/j.1542-4774.2010.tb00532.x. URL https://academic.oup.com/jeea/article-lookup/doi/10.1111/j.1542-4774.2010.tb00532.x.

Y. Cai and K. L. Judd. Dynamic programming with Hermite approximation. *Mathematical Methods of Operations Research*, 81(3):245–267, June 2015. ISSN 1432-2994, 1432-5217. doi: 10.1007/s00186-015-0495-z. URL http://link.springer.com/10.1007/s00186-015-0495-z.

S. Diamond and S. Boyd. CVXPY: A Python-embedded modeling language for convex optimization. *The Journal of Machine Learning Research*, 17(1):2909–2913, Jan. 2016. ISSN 1532-4435.

M. T. Dvorak, K. C. Armour, D. M. W. Frierson, C. Proistosescu, M. B. Baker, and C. J. Smith. Estimating the timing of geophysical commitment to 1.5 and 2.0 °C of global warming. *Nature Climate Change*, 12(6):547–552, June 2022. ISSN 1758-678X, 1758-6798. doi: 10.1038/s41558-022-01372-y. URL https://www.nature.com/articles/s41558-022-01372-y.

Gurobi Optimization, LLC. Gurobi Optimizer Reference Manual, 2023. URL https://www.gurobi.com.

Intergovernmental Panel on Climate Change. *Climate Change 2021: The Physical Science Basis*. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge University Press, 2021.

International Energy Agency. Direct Air Capture 2022 – Analysis. Technical report, International Energy Agency, Apr. 2022. URL https://www.iea.org/reports/direct-air-capture-2022.

K. L. Judd. *Numerical methods in economics*. MIT Press, Cambridge, Mass, 1998. ISBN 9780262100717.

H. D. Matthews, K. B. Tokarska, J. Rogelj, C. J. Smith, A. H. MacDougall, K. Haustein, N. Mengis, S. Sippel, P. M. Forster, and R. Knutti. An integrated approach to quantifying uncertainties in the remaining carbon budget. *Communications Earth & Environment*, 2(1):7, Jan. 2021. ISSN 2662-4435. doi: 10.1038/s43247-020-00064-9. URL https://www.nature.com/articles/s43247-020-00064-9.

K. Riahi, D. P. van Vuuren, E. Kriegler, J. Edmonds, B. C. O'Neill, S. Fujimori, N. Bauer, K. Calvin, R. Dellink, O. Fricko, W. Lutz, A. Popp, J. C. Cuaresma, S. Kc, M. Leimbach, L. Jiang, T. Kram, S. Rao, J. Emmerling, K. Ebi, T. Hasegawa, P. Havlik, F. Humpenöder, L. A. Da Silva, S. Smith, E. Stehfest, V. Bosetti, J. Eom, D. Gernaat, T. Masui, J. Rogelj, J. Strefler, L. Drouet, V. Krey, G. Luderer, M. Harmsen, K. Takahashi, L. Baumstark, J. C. Doelman, M. Kainuma, Z. Klimont, G. Marangoni, H. Lotze-Campen, M. Obersteiner, A. Tabeau, and M. Tavoni. The Shared Socioeconomic Pathways and their energy, land use, and greenhouse gas emissions implications: An overview. *Global Environmental Change*, 42:153–168, Jan. 2017. ISSN 09593780. doi: 10.1016/j.gloenvcha.2016.05.009. URL https://linkinghub.elsevier.com/retrieve/pii/S0959378016300681.

T. Terlouw, K. Treyer, C. Bauer, and M. Mazzotti. Life Cycle Assessment of Direct Air Carbon Capture and Storage with Low-Carbon Energy Sources. *Environmental Science & Technology*, 55 (16):11397–11411, Aug. 2021. ISSN 0013-936X, 1520-5851. doi: 10.1021/acs.est.1c03263. URL https://pubs.acs.org/doi/10.1021/acs.est.1c03263.