ECONOMIC POLICY

Bundesministerium der Finanzen

72nd Economic Policy Panel Meeting

22-23 October 2020

Supported by the Federal Ministry of Finance, Germany

Pareto-Improving Carbon-Risk Taxation

Laurence Kotlikoff (Boston University and The Gaidar Institute)
Felix Kubler (University of Zurich and The Russian Presidential Academy of National Economy and Public Administration)
Andrey Polbin (The Gaidar Institute and The Russian Presidential Academy of National Economy and Public Administration)
Simon Scheidegger (HEC Lausanne and The Russian Presidential Academy of National Economy and Public Administration)

The organisers would like to thank the Federal Ministry of Finance, Germany, for their support. The views expressed in this paper are those of the author(s) and not those of the supporting organization.

CEPR CESifo SciencesPo.

Pareto-Improving Carbon-Risk Taxation*

Laurence Kotlikoff[a,b], Felix Kubler[c,d], Andrey Polbin[b,d], Simon Scheidegger[e,d] †

August 26, 2020

Abstract

Anthropogenic climate change produces two conceptually distinct negative economic externalities. The first is an expected path of climate damage. The second, the focus of this paper, is an expected path of economic risk. To isolate the climate-risk problem, we consider mean-zero, symmetric shocks in our 12-period, overlapping generations model. These shocks impact dirty energy usage (carbon emissions), the relationship between carbon concentration and temperature, and the connection between temperature and damages. Our model exhibits a de minimis climate problem absent its shocks. However, due to non-linearities, symmetric shocks deliver negatively skewed impacts, including the potential for climate disasters. As we show, Pareto-improving carbon taxation can dramatically lower climate risk, in general, and disaster risk, in particular. The associated climate-risk tax, which is focused exclusively on limiting climate risk, can be as large as, or larger than, the carbon average-damage tax, which is focused exclusively on limiting average damage.

JEL classification: F0, F20, H0, H2, H3, J20

Keywords: climate change, uncertainty, carbon taxes, environmental policy, clean energy, externalities, generational equity, economic efficiency, Pareto optimality

---

*We thank Doris Follini for helpful discussions as well as three anonymous referees for many helpful comments. We thank the Gaidar Institute, Boston University, the University of Lausanne, the University of Zurich and the Swiss National Science Foundation (SNF), under project ID "Can Economic Policy Mitigate Climate-Change?", and the Swiss Platform for Advanced Scientific Computing (PASC), under project ID "Computing Equilibria in Heterogeneous Agent Macro Models on Contemporary HPC Platforms", for research support, and the Swiss National Supercomputing Center (CSCS) under project ID 995. Simon Scheidegger gratefully acknowledges support from the MIT Sloan School of Management and the Cowles Foundation at Yale University.

†[a] Department of Economics, Boston University, National Bureau of Economic Research, [b] The Gaidar Institute, [c] Institute for Banking and Finance, University of Zurich, Swiss Finance Institute, [d] The Russian Presidential Academy of National Economy, [e] Department of Finance, University of Lausanne.

1 Introduction

Anthropogenic climate change produces two conceptually distinct negative externalities. The first is a higher expected (average) path of damages. The second is greater volatility in the economy's transition path. This paper focuses exclusively on the second externality. It does so via an overlapping generations (OLG) model with three distinct mean-zero, symmetric shocks. These shocks capture the man-made climate risks highlighted by Weitzman (2012), Golosov et al. (2014), Barnett et al. (2020), and others. The first shock determines whether the economy will use more or less dirty energy. The second shock exacerbates or mitigates the relationship between CO2 emissions and temperature. Finally, the third shock enlarges or shrinks a key parameter in the climate-damage function.

Our OLG model is intentionally bare bones to isolate the cost of each form of risk propagation. In the absence of climate shocks, our model produces minimal climate damage. In consort, the three shocks raise the potential for "climate disasters", which we define as a drop in aggregate consumption by more than one third relative to trend.

In our main calibration, with no carbon policy, but with the three sources of risk activated, the probability, as of time zero, of a climate disaster arising over the next 250 years is more than seven percent, and nine percent over the next 500 years. We find that disasters can arise due to significant skewness in the distribution of damages, reflecting the net impact of the model's non-linear elements. There are two competing supply-side non-linearities in climate models. First, damages are assumed to be a non-linear, convex function of global average surface temperature. Second, the average surface temperature is modeled as a concave, specifically the logarithmic function, of atmospheric CO2. In our model, the supply-side convexity outweighs the supply-side concavity in determining carbon-risk damages. Moreover, when it comes to the welfare effects of carbon risk, the skewed supply-side damages reinforce the skewed demand-side impact arising from risk aversion.

Our paper's primary goal is finding Pareto-efficient paths of carbon-risk taxes – carbon taxes whose raison d'etre is not to limit average carbon damage because average damage is intentionally modeled to be minimal, but to limit downside carbon risk. Unlike prior studies of optimal carbon taxation in the presence of uncertainty, our model features twelve selfish overlapping generations rather than a single, intergenerationally altruistic representative agent. Based on 60 years of adulthood—that is to say, age 20 to age 80, each of our model's periods corresponds to five years. Assuming that current generations are not altruistic toward future generations seems the appropriate framework for modeling carbon policy, given the carbon footprint of current generations. The generationally-selfish life-cycle framework immediately focuses attention on using carbon taxation to effect Pareto improvements—that is, raising the welfare of at least some generations without reducing the welfare of others. Kotlikoff et al. (2019) study such Pareto-improving carbon taxation in a large-scale OlG model, albeit in a deterministic setting. As they show, limiting consideration to efficient policies can materially impact optimal carbon policies.[1]

Ours is hardly the first quantitative analysis of uncertainty's importance to optimal carbon policy. Prior major studies include Brock and Hansen (2017), Gillingham et al. (2015), Jensen and Traeger (2014), Lemoine and Traeger (2016), Cai et al. (2018), Cai et al. (2013), Daniel et al. (2019), and Traeger (2019).[2] Nor is our model the first to posit selfish, life-cycle behavior. Early OLG models that consider resource-extraction and the environment include Howarth and Norgaard (1990), Howarth and Norgaard (1992), Burton (1993), Pecchenino and John (1994), John et al. (1995) and Marini and Scaramozzino (1995), Howarth and Norgaard (1990), Howarth (1991a,b), Burton (1993), Kavuncu and Knabb (2005), and Bovenberg and Heijdra (1998, 2002); Heijdra et al. (2006).[3] Howarth (1991a) considers, in general terms, how to analyze economic efficiency in OLG models in the context of technological shocks.

This said, our model appears to be the first study of optimal carbon-risk tax policy in a large-scale OLG model with shocks both to the climate system and to the macroeconomy. We use our model to identify carbon policies that leave the welfare of current generations unchanged, raise the welfare of future generations by as much as four percent, and lower the probability of a climate disaster from nine to one percent.

[1]Kotlikoff et al. (2019) show that single-agent "optimal" tax solutions may not be Pareto efficient in otherwise identical OLG models. With heterogeneous agents, optimal taxation requires Pareto improvements. Otherwise, all tax schemes, including those that reduce particular generations to starvation, are "optimal" for the right choice of social-planner preferences.

[2]See, e.g., Cai (2020) for a thorough review.

[3] Howarth and Norgaard (1990), positing a pure exchange model, and Howarth (1991b), using a two-period OLG model with capital, point out that policymakers can choose among an infinite number of Pareto efficient paths in the process of correcting negative environmental externalities. Gerlagh and Keyzer (2001); Gerlagh and van der Zwaan (2001) consider the choice among such Pareto paths and the potential use of trust-fund policies that provide future generations with a share of the income derived from the exploitation of natural resources. Gerlagh and van der Zwaan (2001) also point out that demographics can impact the set of efficient policy paths through their impact on the economy's general equilibrium.

What is our precise definition of a generation's welfare? Here we follow Blanchard (2019) in focusing on the generation's expected lifetime utility computed as of time-0 when the policy is initiated. For current generations, the lifetime utility references the remaining lifetime utility. Moreover, as in Blanchard (2019), Pareto improvements are defined with reference to expected lifetime utility.[4]

While achieving a Pareto improvement in terms of expected utility makes sense, there is an infinite number of such improvements to consider, each with its specific configuration of carbon taxation and intergenerational redistribution. Kotlikoff et al. (2019) use Auerbach and Kotlikoff (1987)'s figurative Lump-Sum Redistribution Authority to derive the largest uniform welfare improvement from carbon taxation across all current and future generations. This solution seems of focal interest because of the potential political appeal of uniform treatment. However, in our stochastic model, achieving uniform Pareto gains necessitates state-dependent transfers (see Gottardi and Kubler (2011)). Computing state-dependent policies goes beyond the scope of this paper. Hence, we confine ourselves here to carbon taxes cum state-independent, lump-sum transfer policies that achieve Pareto-improving paths, materially reduce the chances of a climate disaster, and fully compensate current generations for having to pay the carbon taxation. This compensation is financed by the policy's time path of carbon-tax revenue.

Our policy instruments are i) a tax on dirty energy's use that is either time-invariant, or depends on the level of carbon in the atmosphere, and ii) a time-varying sharing-rule of each period's tax revenues among concomitant generations. Revenue sharing is set, period by period, to ensure that current generations achieve the same expected utility as under the business as usual (BAU) scenario. Specifically, a large enough share of revenues is allocated to the oldest generation in period 0 to ensure that its expected utility is unchanged. All other generations receive the identical share. We follow the same procedure in periods 1 through T^*, where T^* is the first period in which an equal share provided to all agents leaves the oldest cohort at time T^* as well off, in terms of ex-anted expected utility, as in BAU. Beyond T^* up to the point that carbon tax revenues are zero, we allocate all revenues equally to all concomitant agents, checking, for each generation after T^* that their cum-policy expected utility is no lower than under BAU.

Our optimal carbon-risk tax policy is set to minimize the chances of reaching a climate-tipping

[4] An alternative Pareto criterion is an ex interim improvement, where an agent is defined by the time and date-event of that individual's birth (see, e.g., Krueger and Kubler (2006)).

point subject to achieving a weak (initial generations remain at their BAU expected-utility levels) Pareto improvement. Reaching a tipping point—that is, experiencing a climate disaster, is defined as experiencing climate damages that result in a drop in aggregate consumption of more than one third. The optimal tax policy raises future generations' expected utilities by up to four percent, depending on their year of birth, and it lowers the probability of a climate disaster from nine to one percent. The policy also reduces worst-case damages from 70 percent to 50 percent of GDP. Moreover, the level of carbon-risk taxation is substantial compared to the average-damage carbon-tax rate.

Our focus on limiting the chance of significant reductions in aggregate consumption is influenced by Barro and Ursua (2008). Their study collects country-specific historical data on significant declines in aggregate consumption. Their findings suggest that aggregate consumption drops of more than one third are very rare, are typically caused by wars, and generally have long-lasting effects on the regions in question.

Our stochastic OLG model is the barest of bones when it comes to concentrating attention to the aforementioned three climate risks. The supply side of our model is similar to Cai et al. (2013) and Nordhaus (2017). Final goods are produced using capital, labor, clean energy, and dirty energy, the use of which emits CO2 into the atmosphere. Clean energy is assumed to be produced by capital and labor and is, thus, subsumed in our production function's capital and labor inputs. Technological progress, captured by dirty energy's coefficient in the model's production function for output, leads, over time, to the crowding-out of dirty energy. Consequently, in the long run, only clean energy is used in production. Unlike Kotlikoff et al. (2019), clean energy is not explicitly modeled. Instead, it is implicitly treated as the economy's ability to produce, through time, with smaller and, ultimately, zero reliance on dirty energy.

We adopt the carbon cycle, and temperature equation posited by Golosov et al. (2014), which simplifies Nordhaus (2017)'s treatment of these elements. Following most of the related literature (see, e.g., Nordhaus (2017), and the references therein), we assume that higher temperatures lower total factor productivity. As Nordhaus (2008) points out, "*the economic impact of climate change ... is the thorniest issue in climate-change economics*". Weitzman (2012) adds a tipping point to the standard Nordhaus damage function, which raises damages dramatically for temperature increases beyond a given level. Below, we follow the specification by Weitzman (2012). Specifically, we posit that damages are the following function of excess temperature (relative to its 1900 value), defined as the global temperature increase T_t^A:

$$D_t = 1 - \frac{1}{1 + \left(\frac{1}{20.46} T_t^A\right)^2 + \left(\frac{1}{6.081} T_t^A\right)^{6.754}}. \tag{1}$$

This specification generates damages that are very similar to those in Nordhaus if the global mean temperature increases by less than 3 degrees Celsius relative to pre-industrial levels. For higher temperature increases, damages are significantly larger. Thus, a 3-degree temperature increase represents our model's tipping point. Such tipping points include losing much of the Amazon rain forest, faster onset of El Niño, the reversal of the Gulf Stream and other ocean circulatory systems, the melting of Greenland's ice sheet, the loss of Siberia's permafrost leading to a massive methane gas release, and the collapse of the West Antarctic ice shelf.

As indicated above, climate change uncertainty comes in three forms. First, dirty energy's share in the production function is subject to symmetric shocks along its path, which trends to zero. Since these shocks can be negative as well as positive, CO2 emissions might remain high, at current or higher levels, for an extended period or might decrease dramatically within a matter of decades. Second, the so-called climate-sensitivity parameter, which determines the increase in temperature arising from increases in CO2, is stochastic, meaning the sensitivity of temperature to carbon concentration can fall as well as rise. Third, the parameter for the quadratic term in the damage function, which governs the probability of crossing the tipping point in damages, is stochastic. We model all three processes as random walks.

We begin by solving the model with no uncertainty, assuming dirty energy is entirely supplanted by clean energy in 120 years. Using Weitzman (2012)'s damage function, damages from climate change are very small. Consequently, efficient, time-invariant dirty energy taxes are, in this context, quite minor as are the efficiency gains from carbon taxation. Stated differently, our model's optimal carbon average-damage tax is de minimis.

We assume throughout that the agents' coefficients of relative risk aversion are 2. This is a moderate-sized coefficient compared to that assumed in the literature (see, e.g., Cai (2020) and references therein). Nevertheless, since damages are skewed, and agents are risk-averse, the average welfare loss is significant. Indeed, as discussed, when all three shocks occur simultaneously, a simple scheme that imposes a carbon tax at $t = 0$, which then increases significantly when temperatures increase, can lead to substantial welfare gains for all generations. Initial generations only gain about 0.1 percent. Welfare gains become larger in about 100 years, with generations born in 150 years gaining about four percent and generations born in 200 years gaining about five percent. Welfare gains then slowly decline, through time, to about three percent.

The paper focuses, at first, exclusively on our dirty-energy usage shock. Adding this shock does not change the expected end date of dirty energy usage. However, it does increase the chances that emissions will remain high over the next 100 years. This, in turn, increases the potential for climate tipping. Our second experiment adds shocks to the climate-sensitivity parameter in addition to the energy-usage shocks. Energy-usage shocks turn out to be a prerequisite for generating climate disasters. The climate-sensitivity parameter plays a crucial role in modern climate modeling. Unfortunately, there is much disagreement about its value (see, e.g., Allen and Frame (2007), Forster et al. (2020), Knutti et al. (2017), or Roe and Baker (2007)). Based on our reading of the climate-science and economics literature, we model this parameter as a random walk with reflecting barriers, where we set the barriers to -30 and +30 percent around the mean. In conjunction with the usage shock, this shock makes climate disasters much more likely, with their probability increasing above five percentage points. The time-invariant, Pareto-improving carbon taxes can help, but not enough to lower the climate-disaster probability below 2.7 percent for these two shocks.

Our final step is to add the damage-function parameter shock to the other two and, thereby, capture climate-damage tipping points. We assume that this parameter shock follows a random walk with an upper reflecting bound and becomes fixed when the temperate reaches the tipping point. In the presence of all three types of uncertainty, the model's disaster probability increases to almost nine percent. With this degree of risk, substantially higher fixed (through time) carbon-risk taxes are Pareto-improving. Moreover, fixed carbon-risk taxes can reduce the probability of climate disasters to around 3 percent. However, as we demonstrate, having CO2-dependent taxes, the likelihood can be reduced to below one percent.

The remainder of the paper is organized as follows: Section 2 presents our model. Section 3 discusses the calibration strategy and reports results for a baseline calibration without uncertainty.

Section 4 considers uncertain dirty energy usage (equivalently, CO2 emissions). Section 5 adds our CO2-temperature sensitivity shock. Section 6 presents the full model with all three shocks, stressing that the carbon-risk tax – the tax needed solely to mitigate carbon risk – can be as large or larger than the carbon average-damage tax – the tax needed solely to mitigate average carbon damage. Section 7 concludes.

2 Model

Time is discrete and indexed by $t = 0, 1, \ldots$. In each period, a cohort of identical agents enters the economy, retires after 10 periods, and dies after 12 periods. Each of our model's periods corresponds to five years. A representative firm produces a single consumption good by using capital, labor, and dirty energy as inputs. Dirty energy is produced using capital and labor.

2.1 Firms

The final goods are produced via

$$Y_t = A_t(D_t) \cdot K_{1t}^{\gamma_t \alpha} \cdot L_{1t}^{\gamma_t(1-\alpha)} \cdot E_t^{1-\gamma_t}, +(1-\bar{\delta})(K_{1t} + K_{2t}) \tag{2}$$

where Y_t is gross output, the price of which is normalized to 1, and A_t, D_t, K_{1t}, L_{1t}, and E_t refer to total factor productivity, climate damage, capital, labor, and dirty energy, respectively. The parameter α represents the capital share. As detailed below in subsection 2.4, the dirty energy's factor share, $1 - \gamma_t$, evolves stochastically as it trends toward 0. Its stochastic path captures our first shock – carbon-usage uncertainty.

Dirty energy is produced with capital and labor—that is, with no fixed or quasi-fixed factors. This is consistent with the modeling in many prior studies (e.g., Cai et al. (2013)). Since financial markets are incomplete and agents are heterogeneous, adding such factors or incorporating adjustment costs would bring the firm's objective function into question—that is, the firm's stochastic discount factor is not uniquely determined and different possible stochastic discount factors lead to different profitmaximizing production plans.[5] Thus:

$$E_t = K_{2t}^{\theta} \cdot L_{2t}^{1-\theta}, \tag{3}$$

where K_{2t} and L_{2t}, respectively refer to capital and labor used in producing dirty energy.[6] Final goods producers purchase dirty energy at the producer price, p_t, plus the carbon tax, τ_t. Each period's revenue from the carbon tax is redistributed among concomitant generations. Capital depreciates at a rate of $\bar{\delta}$, independent of whether it is used in the dirty energy or final goods sectors. CO2 emissions in period t emissions are proportional to E_t with a proportionality factor calibrated to roughly match the industrial emissions of 2015.

Our formulation assumes that elasticities of substitution between dirty energy, capital, and labor equal 1. As Hassler et al. (2012) point out, this is unrealistic when the time period is short. In fact, they find that a Leontief-specification provides a good fit to annual data. However, Hassler et al. (2018) make this Cobb-Douglas (CD) assumption in a model, which, like ours, has periods lasting for 10 years. However, our time period, which, as indicated, corresponds to five years, is relatively short, suggesting the CD assumption is not ideal. In a sensitivity analysis in Section 5, we consider the case where output is produced by a Leontief function of energy and an intermediate input. Energy is Cobb-Douglas in dirty energy and clean energy where both of these energy sources are, themselves, Cobb-Douglas in capital and labor. The intermediate input is produced by a Cobb-Douglas function of capital and labor. In this very simple setup the possibilities of Pareto-improving carbon-taxes turn out to be very limited.[7] The last supply-side point concerns capital's share in producing energy,

[5]The incompleteness reflects the inability of the current generations to trade, in this context pool risk, with proximate, let alone distant future generations. This, in turn, means that differently aged current owners of assets, which are not valued solely based on their current marginal productivity, will view those assets as having different risk-sharing properties in future states of nature. Hence, they will not agree either on how such assets should be valued or on how they should be used or augmented via investment or disinvestment.

[6]We normalize the TFP in dirty energy production at 1 with no loss of generality.

[7]We intend to explore this issue in future work by positing a CES function in two arguments with an intermediate elasitcity of substitution. The outer CES function could incorporate limited ability to substitute between energy and the other CD input, where the inner CES could accommodate relatively high substitutability between clean and dirty energy (see e.g. Papageorgiou et al. (2017) for estimates of the key parameters). Whether these elements would militate toward higher or lower carbon taxation is unclear. An even more realistic treatment would make the dirty energy function depend, not just on capital and labor, but also on the costly extraction of fossil fuels. Furthermore, clean energy would depend on capital and labor as well as a fixed factor, e.g., land, which, as in Kotlikoff et al. (2019), would proxies for a physical limitation, such as the availability of sunshine, on producing clean energy. However, due to the above-mentioned incompleteness of financial markets, entertaining such a more realistic supply-side specification, would raise the aforementioned intractable valuation problem.

which we assume is the same as that for final goods production – an assumption at odds with Barrage (2020). However, relaxing this assumption would muddy the notation without significantly altering our quantitative results.

2.2 Households

Households live for A periods. Those born at time t maximize lifetime expected utility, given by

$$U_t = \mathbb{E}_t \sum_{j=1}^{A} \beta^j \frac{C_{t+j-1,j}^{1-\sigma} - 1}{1-\sigma}, \tag{4}$$

subject to

$$C_{t,j} + a_{t+1,j+1} = (1 + r_t)a_{t,j} + w_t - \theta_{t,j}, \tag{5}$$

where β is the time preference factor, $C_{t,j}$, $a_{t,j}$, $w_{t,j}$ correspond to consumption, assets, and wages of generation j at time t, respectively, and labor supply is normalized at 1. The parameter $\theta_{t,j}$ denotes the possibly state-specific net tax paid by the agent j at time t. The allocation of capital between dirty energy, and goods production is given by

$$\sum_{j=1}^{A} a_{t,j}^{K} = K_t = K_{1t} + K_{2t}. \tag{6}$$

Households born prior to $t = 0$ maximize their remaining lifetime utilities.

2.3 Modeling climate change as a negative externality

We model the carbon cycle as in Golosov et al. (2014). The temperature T_t in period t is determined by the stock of carbon in the atmosphere, S_t,

$$T_t = \lambda_t \frac{\log(S_t/S)}{\log(2)}, \tag{7}$$

where S is the pre-industrial carbon stock. We model λ_t in a stochastic manner.[8] Thus, λ_t is our model's second key shock.

[8] We specify the exact stochastic process below in section 2.4.

Following Golosov et al. (2014), we assume that the CO2 stock in the atmosphere has two components—that is,

$$S_t = S_{1t} + S_{2t}, \tag{8}$$

where

$$S_{1t} = \phi \cdot \xi \cdot E_t + \delta_{S1} \cdot S_{1,t-1}, \tag{9}$$

and where

$$S_{2t} = (1-\phi) \cdot \xi \cdot E_t + \delta_{S2} \cdot S_{2,t-1}. \tag{10}$$

The depreciation parameters satisfy $\delta_{S2} < \delta_{S1}$. We calibrate the former at a low value and the later at a high value again following Golosov et al. (2014). Hence, S_{1t} is a slowly depreciating stock of carbon, whereas S_{2t} is a rapidly depreciating stock. The parameters ϕ and ξ control the fraction of CO2 emissions entering the atmosphere. We take ϕ, ξ, and the depreciation parameters as fixed. As Golosov et al. (2014) point out, there is no consensus in the literature concerning the values of these parameters. In any case, our main results are robust to moderate differences in these parameters as well as to time-varying shocks they may experience.

Next, we slightly modify Weitzman (2012)'s formulation of the temperature damage function (cf. equation (1)) such that it can vary over time—that is,

$$D_t = 1 - \frac{1}{1 + \left(\frac{1}{20.46} T_t^A\right)^2 + \left(\frac{1}{2 \cdot TP_t} T_t^A\right)^{6.754}}, \tag{11}$$

where the term T_t^A refers to global mean surface temperature relative to its 1900 value. Note that the term TP_t in the denominator of equation (11) is our third shock. Weitzman calibrates $2 \cdot TP_t = 6.081$ to be constant over time. This corresponds to a climate tipping point occurring at about 3 degrees excess temperature. Below 3 degrees of excess temperature, the Weitzman tipping-point term in the damage function makes a trivial difference to climate damages. Beyond 3 degrees, it begins to dominate the function.

Climate change reduces output productivity according to

$$A_t = (1 - D_t) \cdot Z, \tag{12}$$

where Z is the constant, non-stochastic production efficiency coefficient—that is, we ignore secular growth for the sake of simplicity.

There is an active debate on the specification of the damage function. Hänsel et al. (2020) and Glanemann et al. (2020) strongly criticize the damage function posited by Nordhaus in his DICE model (see Nordhaus (2017)) and show that alternative, arguably more realistic specifications, lead to much larger damages for 3 degrees or larger increases in temperature and thus to very different optimal carbon taxes. Glanemann et al. (2020), building on Burke et al. (2015), derive a damage function that is similar to our formulation. As Botzen and van den Bergh (2012) work suggests, our damage function specified in equation (11) is very similar to Nordhaus's for temperature increases below 3 degrees. Hänsel et al. (2020) follow Howard and Sterner (2017) and use Nordhaus' functional form, simply changing parameters so that damages reach 6.7 percent of output for a 3-degree temperature increase, as opposed to only 2.1 percent in Nordhaus's calibration. Below, we entertain their damage function in checking the sensitivity of our findings. Optimal taxes are higher, but the probability of disaster is lower under this alternative specification. Simply using the Nordhaus (2017) damages formulation rules out climate disasters and leads to much lower optimal taxes in our model. In 6 we allow for the tipping point to be stochastic. This can be viewed as a generalization of Nordhaus's formulation. With some probability, only Nordhaus's damage function is relevant. If this is not the case, catastrophic damages can arise.

2.4 Stochastic processes

We now specify the stochastic processes for i) the dirty energy share in the production function, $1-\gamma$, ii) the climate sensitivity parameter, λ, and iii) the tipping point, TP. We assume that all three shocks follow random walks:

1. The first shock considers the innovation/emissions uncertainty, and is specified as

$$\gamma_t = \gamma_{t-1} + \epsilon_{\gamma t}, \tag{13}$$

for $\gamma_t \leq 1$, with $\epsilon_{\gamma t} \geq 0$, $\mathbb{E}\epsilon_{\gamma t} = 0.04$, and $\gamma_0 = 0.9$. $\gamma_t = 1$ is an absorbing state, which is reached, on average, in 120 years. To ensure a solution to the model, we assume $\gamma_{60} = 1$—that is, after 300 years, all dirty energy usage ends for certain. This formulation of the γ process captures the gradual decline in the use of dirty energy, punctuated by periodic new dirty energy discoveries that temporarily reverse this trend.

2. The second shock involves the degree to which higher CO2 translates into higher temperatures. Here we assume that the climate sensitivity parameter, λ, follows a random walk,

$$\lambda_t = \lambda_{t-1} + \epsilon_{\lambda t}, \tag{14}$$

with ϵ_{λ_t} being i.i.d. with mean zero and reflecting barriers, namely lower and upper bounds, $\underline{\lambda}$ and $\overline{\lambda}$.

3. The third form of uncertainty concerns damages, the uncertainty of which arises due to uncertainty in the model's tipping-point parameter, TP. We assume that as long as the actual temperature is below the tipping point, the latter follows a random walk with innovation ϵ_{TP}—that is,

$$TP_t = TP_{t-1} + \epsilon_{TP,t}, \tag{15}$$

where $\epsilon_{TP,t}$ is i.i.d. with mean zero, but with a stopping criterion. If, at some period t, atmospheric temperature, T_t^A, reaches the tipping point, TP_t, the tipping point remains fixed. Moreover, we assume that there are reflecting barriers—that is, there exists a lower bound, $\underline{TP}$, as well as an upper bound, $\overline{TP}$. These bounds ensure that the tipping point cannot be too low and that if temperature increases are extreme, a tipping point is reached eventually.

Note that we model emissions uncertainty in a very reduced-form model. Clearly, technological change is, in good part, endogenous, and the speed of green innovations would undoubtedly react to the cost of dirty energy and hence to carbon-taxes (see, e.g., Aghion et al. (2016) and Acemoglu et al. (2012)). If carbon taxes led to more significant and faster green innovation, a given carbon tax would represent a more powerful climate cleaner as assumed in our model. Moreover, by construction, optimal taxes are always higher than in the model where innovations are entirely exogenous. For a given optimal tax-rate for the exogenous innovations model, welfare losses of all current generations (except the 12 periods old) are strictly smaller with endogenous innovations. Hence larger taxes can be sustained, and since they (mechanically) reduce the probability of climate disasters, they will be better. A realistic calibration of this effect is beyond the scope of our analysis.

Our assumption that the climate-sensitivity parameter, which controls the relationship between carbon and temperature, is uncertain, aligns with Hassler et al. (2018)'s modeling. However, recent evidence, including Adams and Dessler (2019), suggests scientists have obtained a more precise understanding of this key parameter's value. Section 5 below examines the sensitivity of our results to this uncertainty. Specifically, we consider the case where the variance of innovation tends to zero over time and, hence, eventually the climate sensitivity becomes deterministic.

Treating the damage function as uncertain reflects two realities. First, the modern world has yet to experience major sustained temperature increases. Hence, we cannot treat the past relationship between relatively small temperature changes and damages as necessarily indicative of the sizes of damages arising with significant temperature changes. Second, there is ample evidence that the temperature-damage relationship is highly non-linear (see Burke et al. (2015)). Hence, it seems realistic to assume that the "tipping point" at which this happens varies stochastically over time.[9]

2.5 Government

In each period, the government imposes taxes on dirty energy use and distributes the revenues among extant generations. As indicated, carbon-risk taxation as well as the distribution of carbon-risk revenues cease when γ_t reaches 1.

We consider both time-invariant as well as carbon-dependent carbon-risk taxes.[10] To find the optimal tax rate in the time-invariant case, we compute different tax rates on a grid. As explained in the introduction, for each tax rate, we implement transfers that use the time path of tax revenues to ensure that each current year enjoys at least their BAU level of expected remaining lifetime utility and that all subsequent generations experience a higher-than-BAU level of expected remaining lifetime utility. In the case of a CO2-dependent tax policy, we add a linear-in-CO2 component to the tax function that delivers a significantly better path of ex-ante expected utility than the fixed tax.

In both the time-invariant and carbon-dependent carbon-risk tax cases, we couple carbon taxation

[9]This statement references mankind's, if not nature's, perspective.

[10]Permitting the carbon to grow at a fixed rate did not materially alter our results.

with time-changing, but state-invariant, generation-specific sharing of carbon-tax revenue. Recall, that generation-specific revenue shares are adjusted in the first 12 periods to ensure that each initial cohort has the same expected remaining lifetime utility under the policy as without it. Furthermore, after the first 12 periods, revenue shares are held fixed at their values in period 12.

We choose the transfers and taxes to guarantee a Pareto-improvement for all current and future generations and to minimize the probability of climate disasters. This is equivalent to maximizing the tax rate under the constraint that no current or future generation loses. In practice, this means keeping current generations at their status-quo welfare (expected remaining lifetime utility) levels and improving the welfare levels (expected lifetime utility) of all future generations.

2.6 Recursive formulation of the OLG model

The aggregate state variables in our model are S_{1t}, S_{2t}, γ_t, λ_t, TP_t, and the aggregate capital stock is K_t. Optimal policies are functions of the aggregate state variables, including the cross-generational distribution of cash on hand. Our computation technique is the projection method developed in Marcet (1988), Marcet and Marshall (1994), Marcet and Lorenzoni (2001), and Judd et al. (2011). Our implementation follows Krusell and Smith (1998), in general, and Kubler and Scheidegger (2019), in particular, by condensing the distribution of assets across agents into one state variable.

In handling the short-term non-stationary policy, we modify the method by Maliar et al. (2015). Stationary dynamic OLG models are generally solved by projecting the economy forward over a long period of time, using the model's Euler conditions to determine optimal choices conditional on the guessed functions, and then updating guessed policy functions using the associated time-series data on optimal choices and the period-specific projected state vector. When a policy is temporarily non-stationary, as is our case, the data projected over the periods of non-stationarity captures only one possible path of the non-stationary policy. To handle this problem, we generate a large number of projections and use the data projected for each time period to update the guessed policy functions for that period.

Agents forecast their future consumption as a function of this condensed state variable and the aggregate states, and they choose an optimal investment based on these forecasts. In equilibrium, the forecasts are almost accurate (within a relative consumption error of at most 10^{-3}, which translates directly into the relative error in consumption-equivalent Euler equations). We approximate the forecasts numerically using Gaussian processes (see, e.g., Scheidegger and Bilionis (2019), and references therein), and we solve for the forecasts using a simulation-based method. For more details, we refer to Kubler and Scheidegger (2019).

Government policies are a function of time and, potentially, the amount of carbon. For the computation of taxes and transfers, we, therefore, include calendar time, $t = 0, 1, \ldots$, as a state variable. Since the economy is non-stationary until only clean energy is used (i.e., until $\gamma_t = 1$), it is crucial to simulate the first 60 periods (i.e., 300 years until we enforce $\gamma = 1$) often to generate good approximations, particularly given the potential for rare events. This is the just mentioned cross-sectional data requirement. We find that 100 simulations of the first 60 periods typically suffice for good numerical results.

We fix the initial conditions at $t = 0$ by assuming that agents live in a deterministic economy with $\gamma = 0.9$ but a constant level of CO2 in the atmosphere through $t = -5$ (i.e., in 25 years) and then suddenly discover the potential for climate change.

Note that in our numerical results, we report probabilities of climate disasters obtained via Monte Carlo simulations. Specifically, we simulate the economy for 100 model-periods (corresponding to 500 years) starting at the initial condition. We repeat this simulation 5000 times, reporting the relative frequency of paths where a climate disaster occurs as the "probability" of a climate disaster.

2.7 Welfare analysis

The welfare of current and future generations is measured as their expected utility at time zero. We report welfare changes as consumption equivalents—that is, given our CRRA specification for the utility function with a coefficient of relative risk aversion of σ, let U_t denote expected utility of generation t in the BAU scenario, and let $\tilde{U}_t$ denote the expected utility with carbon taxes in place. We then compute the consumption-equivalent factor,

$$\left(\frac{\tilde{U}_t}{U_t}\right)^{\frac{1}{1-\sigma}} - 1. \tag{16}$$

The consumption-equivalent factor tells us the higher or lower percentage level of consumption an agent would require, in all states arising under BAU, to achieve the same expected utility as under the policy in question.

3 Calibration

Households live for twelve periods—that is, $A = 12$, where each period corresponds to five years as in Nordhaus (2015). We assume that the capital shares in producing output and energy, α and θ, respectively, both equal 0.3. We set the capital depreciation rate, $\bar{\delta}$ to 0.2. The coefficient of relative risk aversion, σ, is set to 2.0, and the time preference factor, β, is 0.99, leading to an average annual return to capital of about 3 percent. Moreover, through each agent's tenth period, we use a 12-period version of the age-earnings profile as in Kotlikoff et al. (2019) and, after that, assume that the agents work on a 35 percent basis in periods 11 and 12. Having agents continue to work at the end of life proxies for a state-pension system.

In addition, following Golosov et al. (2014), we fix ξ at 0.4, and ϕ at 0.5. Recall from equations (8), (9), and (10) that ξ determines, in part, the degree to which dirty-energy production generates slowly as well as rapidly depleting atmospheric CO2. The coefficient, ϕ, in turn, determines the shares of dirty energy emissions that end up as slowly or rapidly depreciating atmospheric CO2. The slow and fast depreciation rates, δ_{S1} and δ_{S2}, are set at 1.0 and 0.99, respectively. These parameters also coincide, on a period-adjusted basis, with those in Golosov et al. (2014). As in Golosov et al. (2014), we set λ_0 at 3.0, and following Weitzman (2012), we set TP_0 at 3.04. Finally, we calibrate ι, the emissions-proportionality factor, to match initial emissions at 30 GtCO2/year.[11]

3.1 Results for the deterministic benchmark case

Our benchmark model is the OLG model with no uncertainty. As mentioned, we assume a downward trend in γ. Since, to date, there is no clear long-run decline in the dirty-energy share of global energy production (see Hassler et al. (2018)), this assumption is admittedly optimistic.

[11] Nordhaus (2017) states that 2015 CO2 emissions from industrial activity were around 35.85 GtCO2/year). Details on ι's calculation are available from the authors.

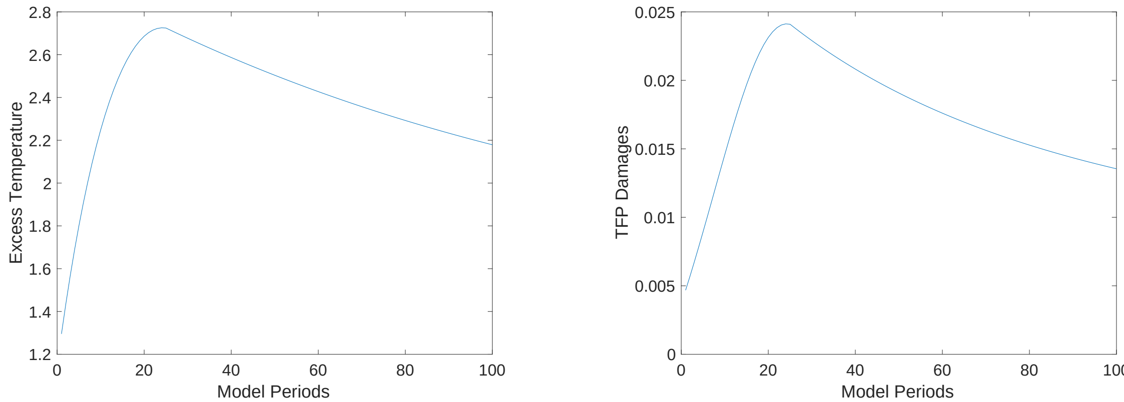

Figure 1: Excess temperature (left panel) and damages (right panel) over the next 500 years—that is, 100 model periods.

Thanks to the downward trend in γ, the solution entails $CO2$ emissions decreasing monotonically and reaching zero after 24 periods—that is, 125 years. The maximum (excess) temperature is slightly below 2.8 degrees, and the maximum damages are less than 2.5 percent of the final output. Figure 1 depicts how temperature and damages evolve over time.

The average return to capital in this baseline deterministic calibration is about 3 percent per year. Using a grid search and lump-sum inter-generational transfers as in Kotlikoff et al. (2019), we find an optimal uniform, welfare-improving carbon-tax of roughly 10 percent that increases the expected utility of all current and future generations by roughly 0.1 percent. Hence, in the absence of risk, there is essentially no scope/need for a carbon-tax policy.

If we restrict the Pareto policy to a maximum uniform increase in the expected utilities of those born in the future, leaving unchanged expected remaining lifetime utilities of initial ($t = 0$) generations, the optimal tax is less than 5 percent. It barely increases the expected utility of all future generations born after 100 years.

This paper's goal is to understand the importance of carbon risk. One means of making this assessment is to compare the sizes of optimal carbon taxes with and without such risk. As just indicated, absent shocks, our model suggests a quite limited role for carbon taxation. This is to be expected since all three of the model's key emission-generating parameters have zero means.

To compare the optimal carbon tax – the carbon average-damage tax – needed to deal solely with average emissions with that needed to deal solely with risky emissions – the carbon-risk tax, we need to specify a reasonable average path of emissions in our deterministic model. To do so, we set ϵ_γ to 0.02. The transition to clean energy now takes 250 years (50 model periods) instead of 125 years above. Consequently, the temperature will increase by 4.5 degrees. This results in damages of about 10 percent of GDP after 150 years, peaking at about 18 percent of GDP after 250 years, and then slowly decreasing after that. Welfare losses to future generations are now substantial – the welfare of generations born in 150 years will be about 6.5 percent lower than the welfare of agents born today. However, the highest tax rate compatible with not over-compensating current generations is only about 15 percent, which will produce welfare gains of up to 3.5 percent for future generations born after 250 years[12].

4 Uncertain CO2 emissions

The crucial elements of uncertainty in our setup run through CO2 emissions and the possibility that the share of dirty energy will remain large over the next 50 years or so. To model this issue, we assume, as explained above, that $\epsilon_{\gamma,t}$, the i.i.d. shock to dirty energy's share, has a positive variance. Moreover, after 60 periods—that is, after some 300 years, we set γ_{60} to 1 if it has not reached this value already. The crucial aspect of this form of uncertainty is that it makes potential long-run damages much more significant. Once CO2 is emitted, about 20 percent remains a permanent feature of the atmosphere. This, in turn, means permanent planetary heating, which implies permanent damages. We assume the following values for this shock:

$$\epsilon_{\gamma t} = \begin{cases} -0.05 & \text{with probability } 0.4, \\ 0.05 & \text{with probability } 0.4, \\ 0.2 & \text{with probability } 0.2. \end{cases}$$

[12]Keeping initial generations whole and ending transfers when dirty energy usage ends rules out higher tax-rate policies that would more than compensate initial generations and, on balance, further help future generations by further improving their climate despite making them pay a larger fiscal bill.

Moreover, we set $\underline{\gamma} = 0.88$ as the reflecting lower boundary. This rules out the dirty energy's production share rising significantly relative to the current status quo. Note that, on average, γ increases by 0.04 in every period, as in the deterministic calibration above. However, it can vary substantially over time. Indeed, in the worst-case scenario, it goes down to 0.88 and stays there for 60 model periods. However, this worst-case scenario has a de minimis probability.

While our formulation of the energy-usage shock lacks micro-foundations, it succinctly captures the qualitative features of this form of uncertainty. Further research, such as Acemoglu et al. (2016), may provide such foundations. Nevertheless, as we now describe, our exercise shows that a slow transition to clean energy can dramatically raise the probability of a climate disaster.

4.1 Business as usual with shocks to dirty energy usage

We now consider shocks to dirty energy usage, starting with the BAU equilibrium. Figures 2, 3, and 4 show the distribution of TFP-damages and temperatures for periods 20 (100 years), 40 (200 years) and 100 (500 years). All histograms show results for 1500 simulations. [13] Note that major economic

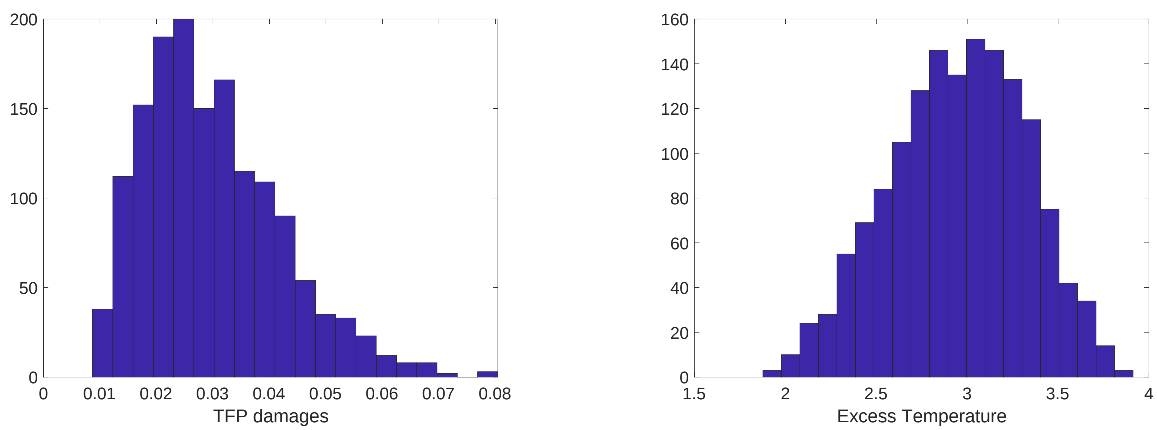

Figure 2: Histogram of damages (left panel) and temperatures (right panel) after 100 years for BAU case with only shocks to γ.

[13]We report the probabilities of climate disasters using 5000 simulations, but use only 1500 simulations for the histograms.

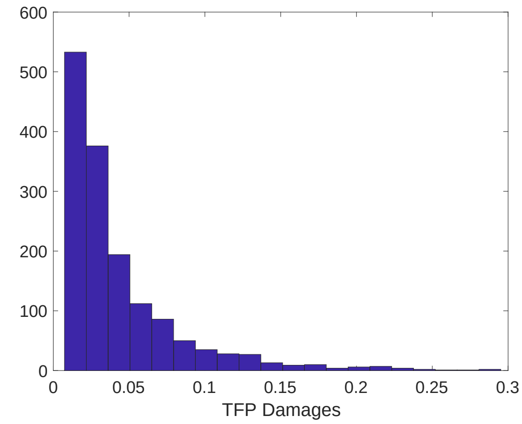

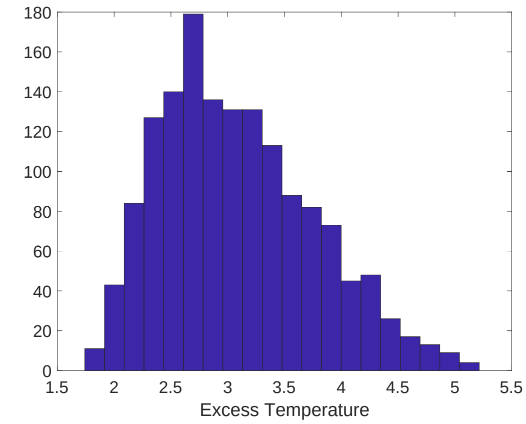

Figure 3: Distribution of damages (left panel) and temperatures (right panel) after 200 years for the BAU case with only shocks to γ.

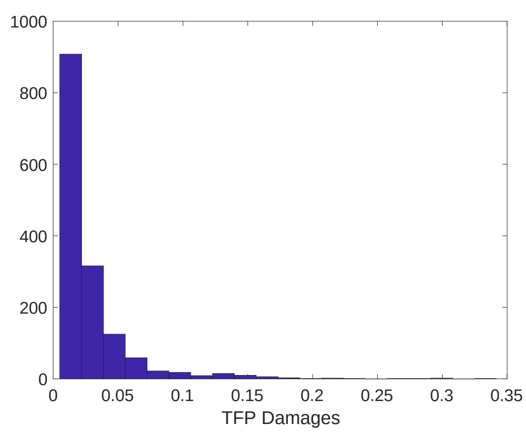

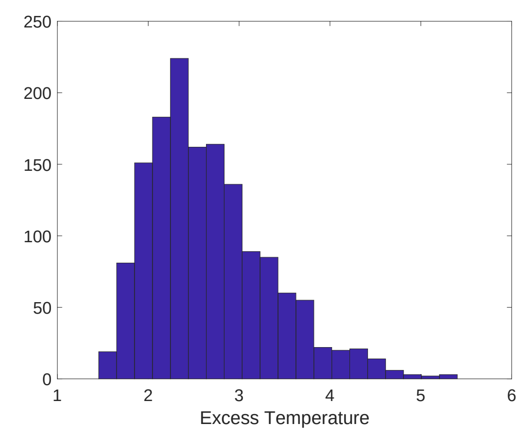

Figure 4: Histogram of damages (left panel) and temperatures (right panel) after 500 years for the case of the BAU case with only shocks to γ.

damages can arise if dirty energy is used for an extended period. Since the transition to green energy can be delayed for many periods, catastrophic damages are possible after 40 periods—that is, 200 years. Figure 3 shows that the worst-case scenario, after 200 years, corresponds to an almost 30 percent reduction in GDP. Figure 4 indicates that after 500 years, there is visible recovery and the damages, with a high probability, are very small. However, some of the worst-case scenarios still lead to significant damages as well as large long-run increases in temperature. When consumption disasters arise, which happens roughly 1 percent of the time, they typically last several decades and reduce aggregate consumption by much more than one third. However, in this specification, none of these consumption disasters turn out to be permanent. This reflects our assumption that only about 60 percent of CO2 in the atmosphere is permanent, with the remaining 40 percent depreciating slowly over time.

4.2 Optimal carbon policy with energy usage shocks

This subsection considers optimal carbon policy assuming just one shock, namely the shock to the decline in dirty energy usage specified above. Our optimal, fixed carbon tax-rate calculation generates a rate of $\tau = 0.25$—that is, a 25 percent tax on dirty energy. With this tax rate, generations born 200 years into the future gain 1.5 percent in consumption equivalence. Current generations, as well as generations in the near future, gain about 0.1 to 0.5 percent. After roughly 125 years—that is, in 25 model periods, the welfare gains become larger, reaching their maximum of 1.5 percent in about 200 years. In about 175 years, the welfare gains are roughly 1.3 percent. Climate disasters become very unlikely with a probability below 0.5 percent.[14]

Figures 5, 6 and 7 show that temperatures after 100 years, and, in particular, after 200 and 500 years are clearly lower than in the BAU case, leading to smaller damages. The shift in temperature appears small at first but, given our form of the damage function, in the high-temperature ranges, small temperature differences make a significant difference with regard to damage.

Clearly, Pareto-improving carbon-risk taxes do not, in this case, entirely prevent deleterious climate change. In particular, significant losses after 200 years and even after 500 years are possible, with temperature increases that are often far beyond that agreed in the Paris Accord. Even higher carbon taxes would limit this problem. However, levying them requires compensating current generations by extracting higher payments from future generations. Unless such payments were state-contingent, they would be imposed on generations born into states with significant climate damages. We leave for

[14] In 5000 simulated paths, only one climate disaster occurred.

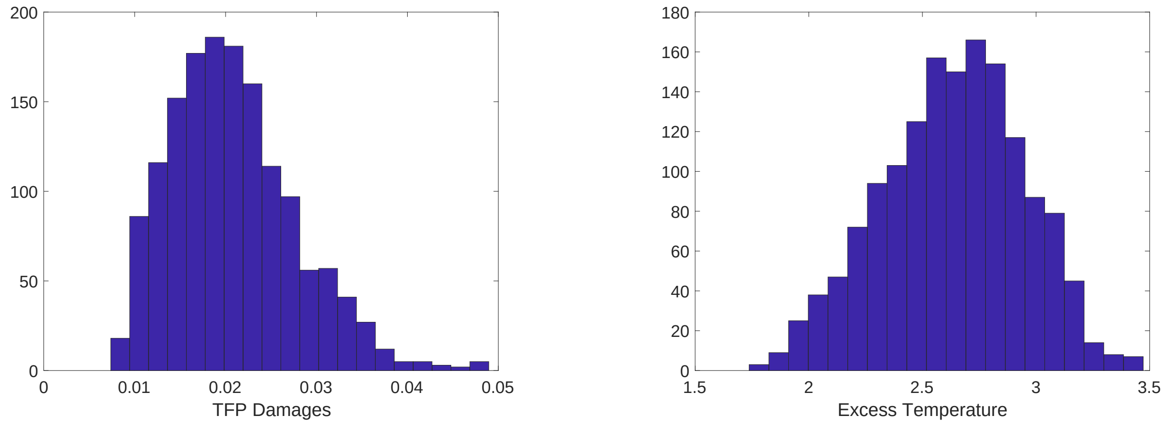

Figure 5: Distribution of damages (left panel) and temperature (right panel) after 100 years for the scenario with only shocks to γ. A tax rate of 25 percent on dirty energy is assumed.

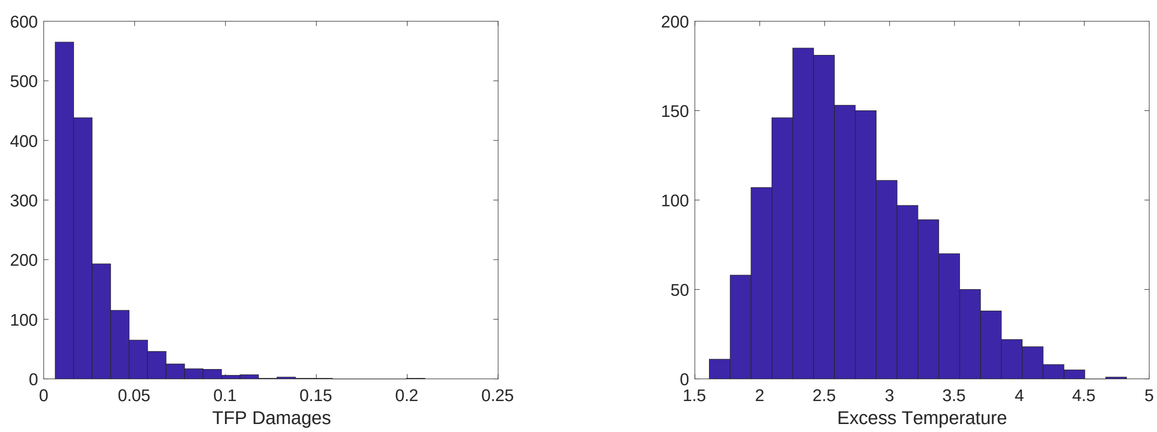

Figure 6: Histogram of damages (left panel) and temperature (right panel) after 200 years for the scenario with energy usage shocks. A tax rate of 25 percent on dirty energy is assumed.

further research the degree to which state- and generation-contingent net taxes can further mitigate climate change.

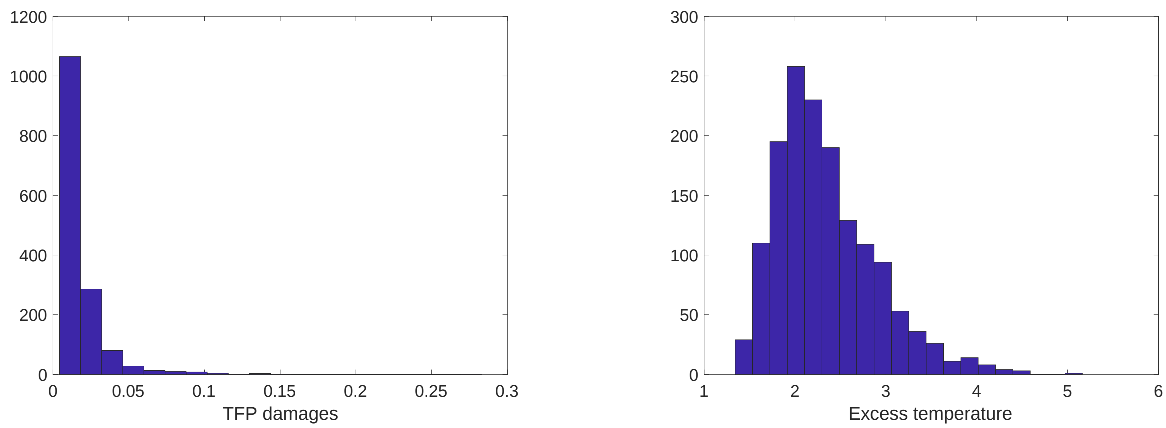

Figure 7: Distribution of damages (left panel) and temperature (right panel) after 500 years for the scenario with only shocks to γ. A tax rate of 25 percent on dirty energy is assumed.

5 Shocks to both dirty energy usage and temperature

In this section, we combine uncertainty about future dirty energy usage and, thus, emissions with uncertainty about the degree to which CO2 emissions translates into higher temperatures. While the climate-sensitivity coefficient plays a crucial role in understanding the effects of CO2 emissions on global warming, there is, unfortunately, little scientific consensus on its value. As Knutti et al. (2017) put it, "*Equilibrium climate sensitivity characterizes the Earth's long-term global temperature response to increased atmospheric CO2 concentration. It has reached almost iconic status as the single number that describes how severe climate change will be. The consensus on the likely range for climate sensitivity of 1.5 to 4.5 degrees Celsius today is the same as given by Jule Charney in 1979.*"

The ratio of the highest estimate to the lowest is roughly three! A large part of this parameter uncertainty is typically attributed to model uncertainty. In other words, researchers believe that the time variation in the true parameter is much smaller than this range, but that our current climate models cannot determine this true parameter. However, there is some evidence that even beyond model-uncertainty, the equilibrium climate sensitivity varies stochastically (see, e.g., Roe and Baker (2007)) through time.

Unfortunately, the climate model that we adopted from Golosov et al. (2014) is too simplistic to model this realistically. In order to get an idea of the impact of stochastic climate sensitivity, we model its variation by assuming that λ_t follows a random walk. Let us consider the following specification for shocks to λ, which translates CO2 levels into the forcing variable that raises the average global temperature. We assume that each period the shock $\epsilon_{\lambda,t}$ can take two values—that is,

$$\epsilon_{\lambda,t} = \begin{cases} -\frac{1}{30} & \text{with prob } \frac{1}{2}, \\ +\frac{1}{30} & \text{with prob } \frac{1}{2}. \end{cases}$$

We assume that $\underline{\lambda} = 0.7\lambda_0$, $\overline{\lambda} = 1.3\lambda_0$ are reflecting barriers for the random walk. This specification does not capture all recognized uncertainty about the climate sensitivity parameter. As Forster et al. (2020) point out, this is much larger and would put the reflecting barriers well below 0.5 and well above 1.5. Hence, our calibration is highly conservative—that is, the actual uncertainty is likely much larger.[15] Our results therefore constitute a lower bound on the possibility of climate disasters.

5.1 Business as usual with dirty energy usage and temperature shocks

As in the previous section, we first compute the equilibrium for an economy with no carbon tax. Figures 8, 9, and 10 show the distribution of TFP-damages and temperature for the periods 20 (100 years), 40 (200 years), and 100 (500 years).

Comparing figures 8, 9, and 10 with figures 2, 3, and 4 from above, it becomes evident that the tails become larger. After 200 years, extreme temperatures and extreme damages now occur much more frequently[16]. Furthermore, after 500 years—that is, in the very long run, substantial damages can persist.

Note that the histograms do not show the fact that with a time-varying climate sensitivity parameter, output and aggregate consumption become very volatile when the CO2 concentration in the atmosphere is high. Given our specification of the damage functions, an increase in the climate sensitivity from $\underline{\lambda}$ to $\overline{\lambda}$ changes the temperature by more than 60 percent, which can change the damages from relatively small to quite large when the initial excess temperature is above 3 degrees.

[15] However, much of it is model uncertainty which we cannot describe without having an explicit model of how learning about this parameter occurs over time.

[16] It is important to note that the scales differ across the examples.

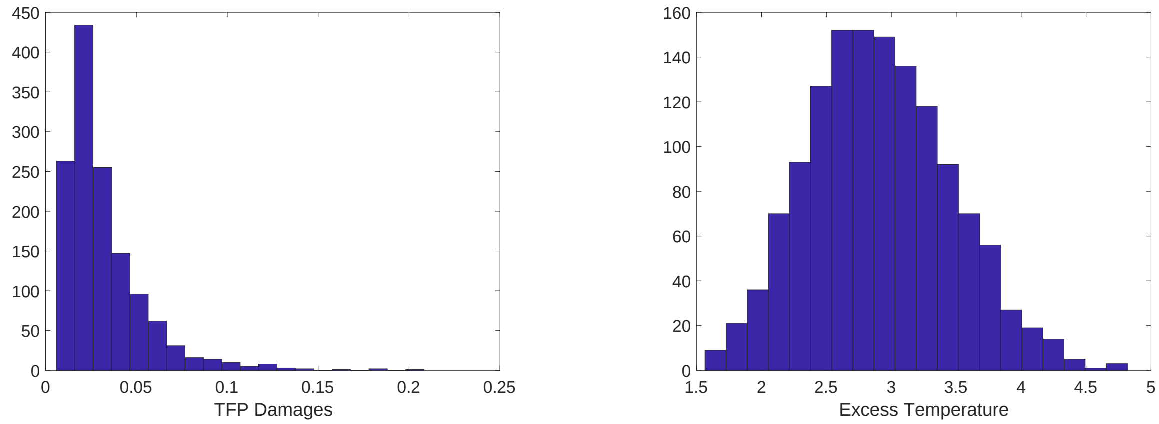

Figure 8: Histogram of damages (left panel) and temperature (right panel) after 100 years for the BAU case with shocks to γ and λ.

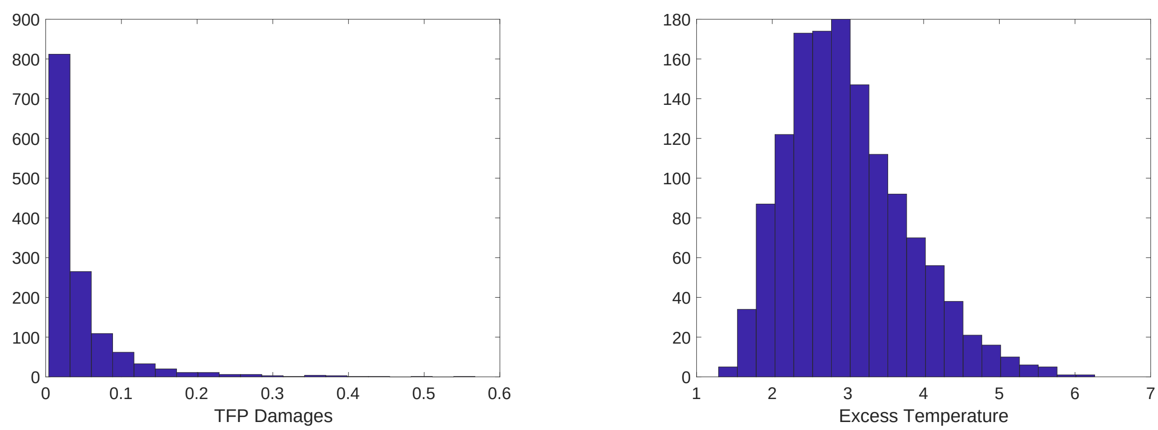

Figure 9: Distribution of damages (left panel) and temperature (right panel) after 200 years for the BAU case with shocks to γ and λ.

Using our specifications here, climate disasters occur frequently. In about five percent of the cases, aggregate consumption drops by more than one third. These drops typically last more than 100 years, in some cases, up to 350-400 years, and in some cases forever.

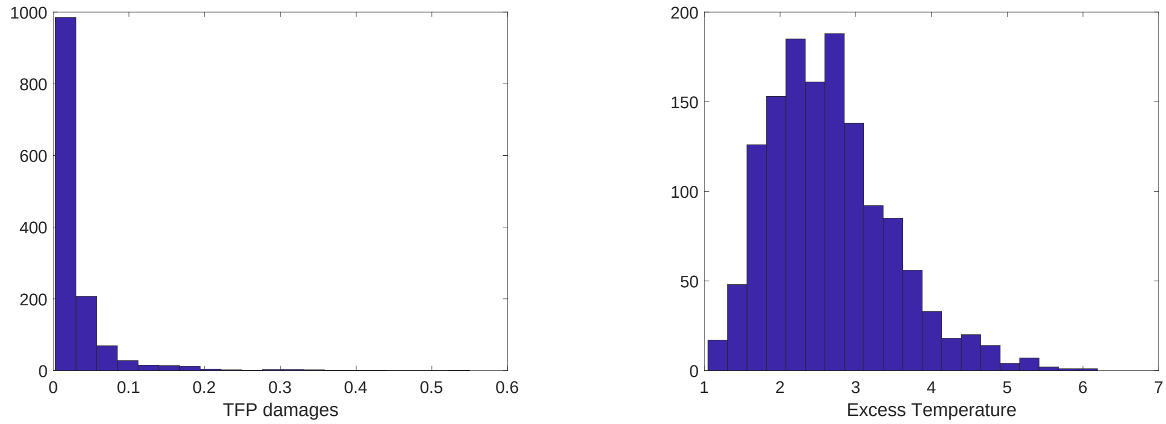

Figure 10: Distribution of damages (left panel) and temperature (right panel) after 500 years for the BAU case with shocks to γ and λ.

Weitzman (2009) and Weitzman (2012) describes scenarios where the said uncertainties can lead to catastrophic outcomes of climate change. However, unlike Weitzman (2009), the events in our model are not climate catastrophes that push expected utility to minus infinity. In our calibration, aggregate consumption never drops by more than two thirds. Since this happens with relatively low probability, the effect on average/expected utility is modest—that is, it drops, measured in certainty equivalents, by around 5 percent for some generations in the future.

5.2 Optimal carbon policy with energy usage and temperature shocks

As in the previous numerical example, the fact that aggregate consumption can decrease substantially in the future does not necessarily mean that higher taxes today are feasible if current generations need to be compensated. The fact that in 500 or even 1000 years, economic damages can be dramatic offers little guidance on how to achieve a Pareto improvement today.

In this two-shock case, the "optimal" fixed tax rate on dirty energy is 20 percent tax. Generations born 150 years from now or even later gain about 2 percent in certainty equivalence, and the probability of a climate disaster is reduced from more than 5 percent to around 2.5 percent.

Note that this carbon tax is 5 percent lower than the one we found above for the case of no climate uncertainty. The reason for this is that for generations born within the first five periods, an increase in precautionary savings that is caused by climate uncertainty counteracts the negative effects of the extra uncertainty, and they can no longer be compensated for facing a tax rate of 25 percent.

5.3 Sensitivity analysis

To examine the robustness of the result we conduct three different sensitivity analyses. First, we alter the stochastic process for the climate sensitivity parameter, λ, and assume that eventually this parameter becomes deterministic. Second, we consider Hänsel et al. (2020)'s damage function. Third, we assume that energy and an intermediate good are perfect complements in a Leontief production function.

5.3.1 Decreasing uncertainty about climate sensitivity

We assume here that the variance of the innovation in (14) decreases over time and becomes deterministic in 30 periods (150 years). In particular, for $t \geq 30$ we take

$$\epsilon_{\lambda,t} = \begin{cases} -\frac{30-t}{30} \cdot (30-t) & \text{with prob } \frac{1}{2}, \\ +\frac{30-t}{30} & \text{with prob } \frac{1}{2}, \end{cases}$$

and we take $\epsilon_{t,} = 0$ for $t > 30$. As before, $\underline{\lambda} = 0.7$ and $\overline{\lambda} = 1.3$.

The following table shows the impact of a long-run decline in climate sensitivity uncertainty. The table's results are quite robust to the precise specification of the stochastic process for λ. The

	const. vol.	decreasing vol.
Probability of disaster, BAU	0.05	0.038
Probability of disaster, opt. tax	0.025	0.012
Optimal fixed tax	0.2	0.2
Welfare gains	0.02	0.02

Table 1: Sensitivity to stochastic climate sensitivity

probability of disaster decreases significantly in the BAU scenario. This is due to the fact that in the base-case specification disaster can happen well after the full transition to clean energy. Once γ reaches 1, the temperature starts dropping with certainty as soon as λ becomes deterministic. This is obviously not what most climate models predict, therefore we chose to keep λ random in our main calibration.

5.3.2 Alternative damage function

Following Hänsel et al. (2020), we use the "preferred specification" of damages from Howard and Sterner (2017). We assume that $D(T_t^A) = 0.007438(T_t^A)^2$. Unlike in our base-case specification, there are no tipping points in this formulation. However, damages are much larger than in Nordhaus's specification, and clearly, the damage function is only well defined for $T_t^A < 11.56$ since otherwise damages are larger than 100 percent of GDP.

The following table 2 compares the key results for the two damage functions. This table shows

	Weitzman	Hansel et al.
Probability of disaster, BAU	0.05	0.016
Probability of disaster, opt. tax	0.025	0.003
Optimal fixed tax	0.2	0.35
Welfare gains	0.02	0.06

Table 2: Sensitivity with respect to the form of the damage function.

that the results depend crucially on the specification of the damage function. With the alternative damage function, climate disasters are much less common, yet welfare gains from carbon taxation and the optimal tax rate are much higher. This is easily explained by the fact that average damages are much higher with this damage function, but extreme damages only come with very high temperature-increases.

The analysis shows that the proper specification of the damage function remains an extremely important open question for this strand of research.

5.3.3 Alternative production function

As in Hassler et al. (2012), we consider now the case where total output is produced by a Leontief production function—that is,

$$Y_t = A(D_t) \cdot \min[\zeta \cdot K_{1t}^{\alpha} \cdot L_{1T}^{1-\alpha}, (1-\zeta) \cdot E_t], \tag{17}$$

and we assume that energy is produced from dirty and clean energy, i.e.,

$$E_t = E_{tc}^{\gamma_t} \cdot E_{td}^{1-\gamma_t}. \tag{18}$$

Clean and dirty energy, E_{tc} and E_{td}, are produced in identical Cobb-Douglas production functions with capital share α.

We set $\zeta = 0.1$, $\gamma_0 = 0.08$, and calibrate the stochastic process for γ in such a a way that in the BAU scenario damages are the same as in our main calibration above.

Table 3 compares the key results for the two production functions. It shows that the optimal

	Cobb-Douglas	Leontief
Probability of disaster, BAU	0.05	0.05
Probability of disaster, opt. tax	0.025	0.049
Optimal fixed tax	0.2	0.1
Welfare gains	0.02	0.005

Table 3: Sensitivity with respect to the production function.

tax as well as welfare gains depend crucially on the specification of the production function. When the clean energy share in the production of energy is low, and energy cannot readily substitute for the intermediate input, carbon taxes have, as expected, a minimal impact on carbon emission. Furthermore, the level of Pareto-improving carbon-taxes that can be sustained is very small.

6 Three sources of risk

This section incorporates all three types of shocks mentioned above. To model the stochastic damage function, we assume that each period the shock $\epsilon_{TP,t}$ can realize two values—that is,

$$\epsilon_{TP,t} = \begin{cases} -\frac{1}{10} & \text{with prob } \frac{1}{2}, \\ +\frac{1}{10} & \text{with prob } \frac{1}{2}. \end{cases}$$

Moreover, we assume that $\overline{TP} = 2.5$ and $\underline{TP} = 3.5$ are reflecting barriers for the random walk. As explained above, when the current excess temperature reaches TP, the random walk stops at that value. If the current tipping point is at 2.5 when the excess temperature reaches 2.5 degrees, very

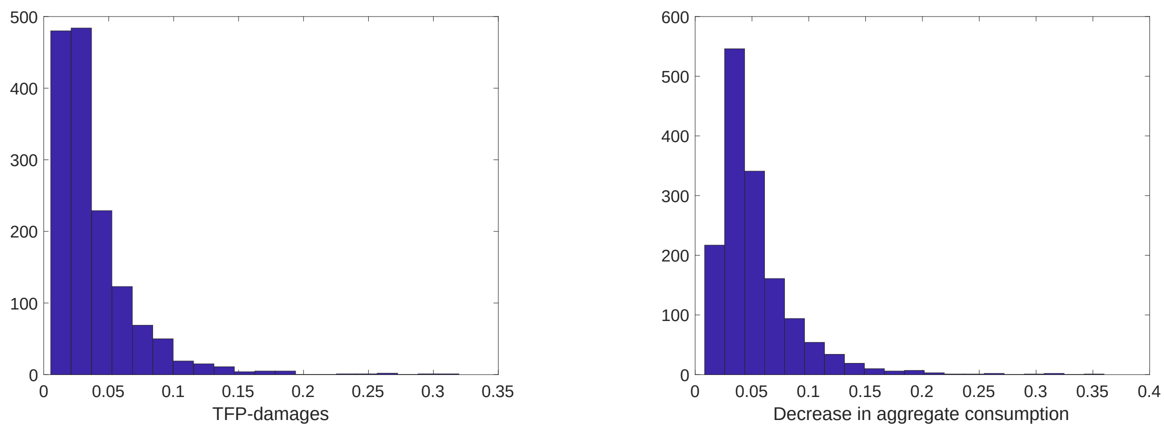

Figure 11: Distribution of damages (left panel) and aggregate consumption losses (right panel) after 100 years in BAU scenario for full model.

significant damage is likely to result. If the current temperature never reaches the tipping point, the damages are quite small.

While it seems clear that there are large uncertainties about economic damages from higher temperatures, there is no unambiguous way to model them. As mentioned in the introduction, Golosov et al. (2014) follow a somewhat similar strategy to ours and assume uncertainty about a coefficient in the damage function. However, in their approach, all uncertainty is resolved at some given future date, $\hat{T}$. At $\hat{T}$, their coefficient randomly takes one of two values and before $\hat{T}$, their coefficient is set to the average of these two values. It seems more realistic to have the values evolve stochastically over time.

Without carbon taxation, our three-shock model results in significant damages with fairly high probability. Figures 11, 12, and 13 show the distribution of TFP-damages and aggregate consumption losses for the 20 (corresponding to a duration of 100 years), 40 (200 years), and 100 (500 years) periods. Since the temperature itself is not as meaningful, given the uncertainty in the damage function, we depict the distribution of aggregate consumption losses instead. The figures are normalized in such a way that aggregate consumption at t=0 is 1.

Compared to the figures 2 and 10 from above, we see significant differences in damages. With

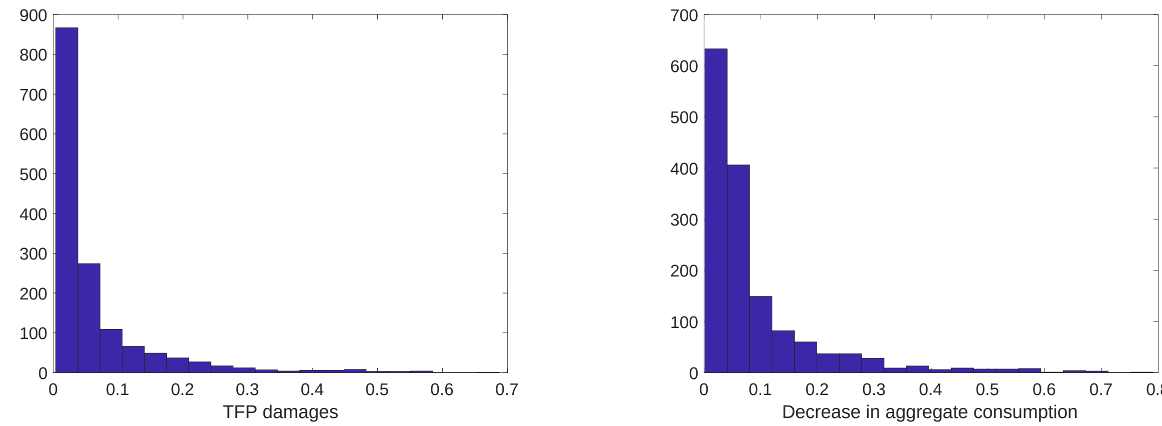

Figure 12: Distribution of damages (left panel) and aggregate consumption losses (right panel) after 200 years in the BAU scenario for full model.

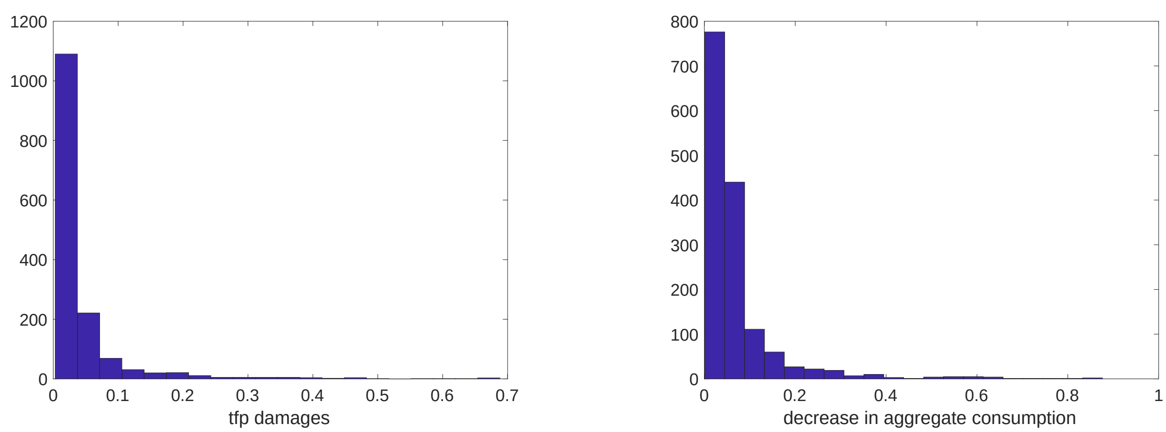

Figure 13: Distribution of damages (left panel) and aggregate consumption losses (right panel) after 500 years in the BAU scenario for full model.

this third source of uncertainty, the probability of climate disasters increases to about nine percent[17]

[17]Note that in the figures presented here, the frequency of a climate disaster seems significantly lower than nine percent. This is due to the fact that the climate sensitivity parameter follows a random walk and along some paths the economy experiences transitory climate disasters, i.e., many climate disasters only occur before year 200 and many only after year 200.

compared to five percent in the model without a stochastic tipping point in the damage function. As in the case in section 5, most disaster periods last for more than 100 years and many for more than 500 years. In this model, declines in aggregate consumption of almost 80 percent become possible, even though they are unlikely. In the BAU scenario, some future generations are born into catastrophic economies.

The level of climate risk we model in this section now also decreases average future utilities significantly. Generations that are born 200 years from now have, in consumption equivalents, a seven percent lower expected utility than generations born today.

6.1 Optimal carbon policy with all three shocks

As lower tipping points and larger damages become possible, substantial damage from climate change can occur earlier in the model, making a larger initial tax feasible. With a fixed tax of 35 percent, we find that the average welfare of generations born 200 years in the future increases (in consumption equivalents) by around 4 percent. The probability of climate disasters decreases to around 3.5 percent. Hence there is still substantial climate disaster risk that cannot be prevented by a constant tax rate.

It is useful to compare the taxes in the specification of the model with three sources of uncertainty to optimal taxes in a deterministic model. In this thought experiment, we focus on the model with much more significant damages to illustrate the role of uncertainty in the damage function. To do this, we consider the calibration from section 3.1 (the 'deterministic benchmark case'), but multiply damages by a factor of 4—that is, the damage function becomes

$$D_t = 1 - \frac{\mathbf{4}}{1 + \left(\frac{1}{20.46} T_t^A\right)^2 + \left(\frac{1}{2TP_t} T_t^A\right)^{6.754}}. \tag{19}$$

This shifts up the damages depicted in Figure 1 by a factor of roughly 4 (not exactly because more significant damages lead to a decrease in production and CO2 emissions) resulting in substantial damages after 20 model periods. The Pareto-improving fixed tax rate that leads to the largest welfare gains for generations born after 120 years is roughly about 20 percent.[18] This results in maximal

[18]Note that this is still a much lower carbon tax than the optimal tax arising in realistically calibrated models with growth and intergenerational transfers.

welfare gains of about 1.5 percent for generations born after 110 years. This thought experiment illustrates the significant quantitative effects of uncertainty. In a world of certainty, even if damages are four times larger than in typical calibrations, taxes and welfare gains are still relatively small.

6.2 Carbon dependent taxes

If we allow taxes to vary over time and with the level of atmospheric carbon, we reduce the risk of climate disasters even further by increasing the tax when CO2 concentration in the atmosphere increases substantially. In order to keep the tax Pareto-improving, one needs to start with a substantial initial tax and start increasing the tax relatively late when future generations already benefit from the initial tax.

Note that in the presence of a stochastic damage function, it might not be sufficient to control the concentration of CO2 in the atmosphere since a low tipping point might result in significant damages. Moreover, imposing high carbon taxes once the CO2 concentration has reached a certain level might itself lead to consumption disasters. In the presence of significant damages, these taxes push aggregate consumption down further.

The "best"specification we could find imposes additional taxes earlier than above and increases taxes very steeply once CO2 concentration has reached a given level. We set

$$\tau_t = 0.3 + \frac{1}{10}\max(0, \frac{\log(S_t/S)}{\log(2)} - 2)^3, \tag{20}$$

where S_t denotes the total amount of $CO2$ in the atmosphere.

Observe that taxes start increasing if the excess temperature (at a climate sensitivity of 3) is 2 degrees. Moreover, at a temperature increase (again with $\lambda = 3$) of more than 3 degrees, taxes start rising very rapidly. Along some paths, carbon-risk taxes increase to over 100 percent.

In Figure 14 we depict the carbon tax rates for the periods where they are relevant (i.e. $\gamma_t < 1$) for 600 simulated periods. The histogram shows that in the vast majority of periods taxes are 30 percent and therefore slightly lower than the optimal fixed tax of 35 percent above. However, it often increases to more than 35 percent and in some cases well beyond that.

Using this specification, the probability of climate disaster is reduced to around 1.3 percent – a

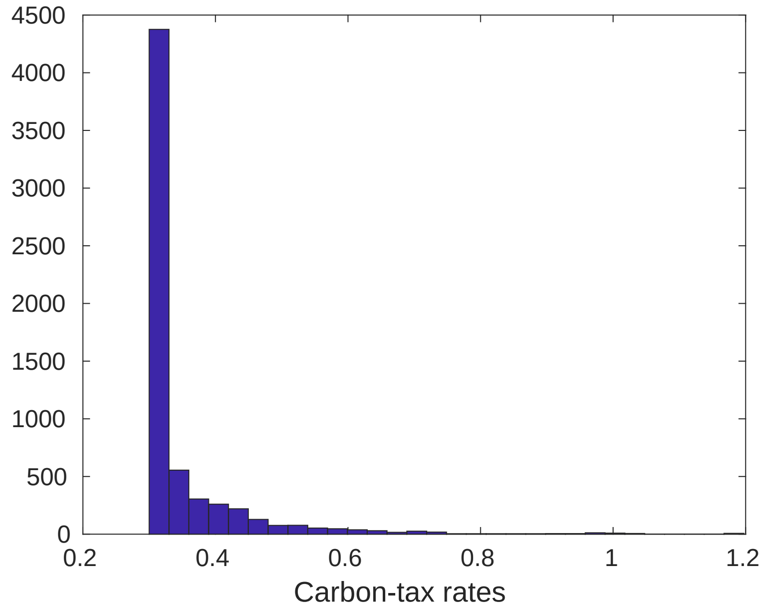

Figure 14: Distribution of Carbon tax-rates.

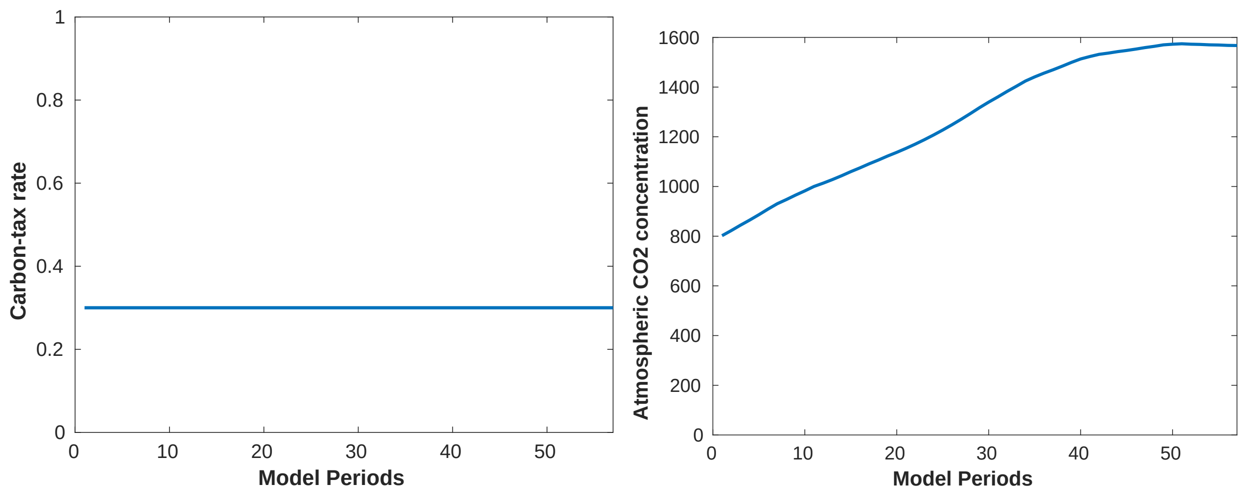

Figure 15: Tax rate and CO2 in atmosphere for fixed taxes.

much lower level than with a fixed tax. However, the welfare gains for future generations are very similar to the case of a fixed tax. Generations born in around 150 years gain slightly less (because some agents are born into economies that still use significant amounts of dirty energy and hence suffer from higher taxes), whereas generations born after 250 years gain slightly more, but the differences are below 0.5 percent.

Figures 15 and 16 show taxes and CO2 concentration in the atmosphere for the two cases of constant taxes and risk taxes for a given path. On this path, the transition to clean energy takes 57 model periods, so this is a path for a high emission, high damage scenario. One can see that risk-taxes increase dramatically after around 20 to 25 model periods, further reducing CO2 emissions

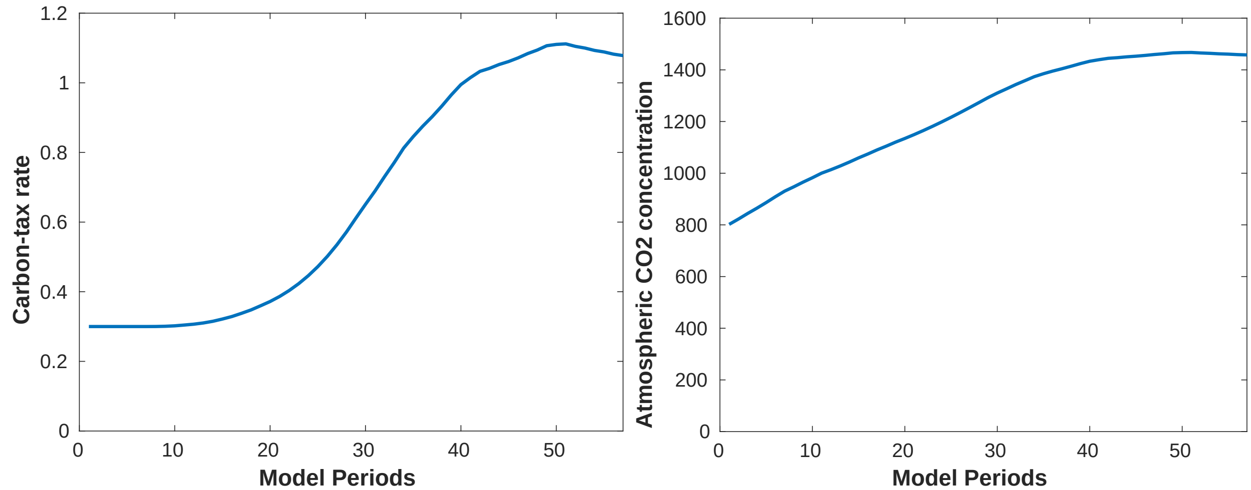

Figure 16: Tax rate and CO2 in atmosphere for risk taxes.

significantly and leading to large welfare gains in the future. On average, a generation born after 40 model periods is still better off despite the fact that carbon taxes are 100 percent and the agents born into this scenario are certainly worse off than in the constant tax scenario.

To decrease this likelihood even further, one could consider a tax rate that changes according to variables other than CO2 concentration. However, it seems unrealistic to make the tax rate dependent on the tipping point - after all, our modeling strategy is a simplified version of a model where agents do not know the tipping point and learn it when it happens. After the tipping point is reached, it is undoubtedly too late to increase taxes.

7 Conclusion

This paper examines the role of uncertainty in optimal carbon taxation. We argue that starting from a model without uncertainty, the introduction of mean-preserving shocks can lead to significant and long-lasting damages. Even with modest degrees of risk aversion, the welfare loss of future generations can be significant and climate disasters can be quite frequent. Carbon-risk taxes can achieve Pareto improvements and if they increase at a sufficient rate with CO2 concentration, they can certainly help prevent climate disasters.

Although many aspects of the model are roughly calibrated, the parameters that we choose for the relevant stochastic processes are very "conservative" in the sense that the true uncertainty is likely to be much larger. Hence, our results provide lower bounds on possible damages and on improving carbon taxes. Finally, we show that carbon-risk tax rates can be as large, if not larger, than carbon average-damage tax rates depending on which intergenerational policies are feasible.

References

Acemoglu, D., Aghion, P., Bursztyn, L., and Hemous, D. (2012). The environment and directed technical change. *American economic review*, 102(1):131–66.

Acemoglu, D., Akcigit, U., Hanley, D., and Kerr, W. (2016). Transition to clean technology. *Journal of Political Economy*, 124(1):52–104.

Adams, B. and Dessler, A. (2019). Estimating transient climate response in a large-ensemble global climate model simulation. *Geophysical Research Letters*, 46(1):311–317.

Aghion, P., Dechezleprêtre, A., Hemous, D., Martin, R., and Van Reenen, J. (2016). Carbon taxes, path dependency, and directed technical change: Evidence from the auto industry. *Journal of Political Economy*, 124(1):1–51.

Allen, M. R. and Frame, D. J. (2007). Call off the quest. *Science*, 318(5850):582–583.

Auerbach, A. J. and Kotlikoff, L. J. (1987). *Dynamic fiscal policy*. Cambridge University Press.

Barnett, M., Brock, W., and Hansen, L. P. (2020). Pricing uncertainty induced by climate change. *The Review of Financial Studies*, 33(3):1024–1066.

Barrage, L. (2020). Optimal dynamic carbon taxes in a climate–economy model with distortionary fiscal policy. *The Review of Economic Studies*, 87(1):1–39.

Barro, R. J. and Ursua, J. F. (2008). Consumption disasters in the twentieth century. *American Economic Review*, 98(2):58–63.

Blanchard, O. (2019). Public debt and low interest rates. *American Economic Review*, 109(4):1197–1229.

Botzen, W. W. and van den Bergh, J. C. (2012). How sensitive is nordhaus to weitzman? climate policy in dice with an alternative damage function. *Economics Letters*, 117(1):372–374.

Bovenberg, A. and Heijdra, B. J. (1998). Environmental tax policy and intergenerational distribution. *Journal of Public Economics*, 67(1):1–24.

Bovenberg, A. L. and Heijdra, B. J. (2002). Environmental abatement and intergenerational distribution. *Environmental and Resource Economics*, 23(1):45–84.

Brock, W. A. and Hansen, L. P. (2017). Wrestling with uncertainty in climate economic models. *working paper*.

Burke, M., Hsiang, S. M., and Miguel, E. (2015). Global non-linear effect of temperature on economic production. *Nature*, 527(7577):235–239.

Burton, P. S. (1993). Intertemporal preferences and intergenerational equity considerations in optimal resource harvesting. *Journal of Environmental Economics and Management*, 24(2):119–132.

Cai, Y. (2020). The role of uncertainty in controlling climate change. *arXiv preprint arXiv:2003.01615*.

Cai, Y., Brock, W., Xepapadeas, A., and Judd, K. (2018). Climate policy under cooperation and competition between regions with spatial heat transport. Technical report, National Bureau of Economic Research.

Cai, Y., Judd, K. L., and Lontzek, T. S. (2013). The social cost of stochastic and irreversible climate change. Technical report, National Bureau of Economic Research.

Daniel, K. D., Litterman, R. B., and Wagner, G. (2019). Declining co2 price paths. *Proceedings of the National Academy of Sciences*, page 201817444.

Forster, P. M., Maycock, A. C., McKenna, C. M., and Smith, C. J. (2020). Latest climate models confirm need for urgent mitigation. *Nature Climate Change*, 10(1):7–10.

Gerlagh, R. and Keyzer, M. A. (2001). Sustainability and the intergenerational distribution of natural resource entitlements. *Journal of Public Economics*, 79(2):315–341.

Gerlagh, R. and van der Zwaan, B. (2001). The effects of ageing and an environmental trust fund in an overlapping generations model on carbon emission reductions. *Ecological Economics*, 36(2):311–326.

Gillingham, K., Nordhaus, W. D., Anthoff, D., Blanford, G., Bosetti, V., Christensen, P., McJeon, H., Reilly, J., and Sztorc, P. (2015). Modeling uncertainty in climate change: A multi-model comparison. Technical report, National Bureau of Economic Research.

Glanemann, N., Willner, S. N., and Levermann, A. (2020). Paris climate agreement passes the cost-benefit test. *Nature communications*, 11(1):1–11.

Golosov, M., Hassler, J., Krusell, P., and Tsyvinski, A. (2014). Optimal taxes on fossil fuel in general equilibrium. *Econometrica*, 82(1):41–88.

Gottardi, P. and Kubler, F. (2011). Social security and risk sharing. *Journal of Economic Theory*, 146(3):1078–1106.

Hänsel, M. C., Drupp, M. A., Johansson, D. J., Nesje, F., Azar, C., Freeman, M. C., Groom, B., and Sterner, T. (2020). Climate economics support for the un climate targets. *Nature Climate Change*, pages 1–9.

Hassler, J., Krusell, P., and Olovsson, C. (2012). Energy-saving technical change. Technical report, National Bureau of Economic Research.

Hassler, J., Krusell, P., and Olovsson, C. (2018). The consequences of uncertainty: climate sensitivity and economic sensitivity to the climate. *Annual Review of Economics*, 10:189–205.

Heijdra, B. J., Kooiman, J. P., and Ligthart, J. E. (2006). Environmental quality, the macroeconomy, and intergenerational distribution. *Resource and Energy Economics*, 28(1):74–104.

Howard, P. H. and Sterner, T. (2017). Few and not so far between: a meta-analysis of climate damage estimates. *Environmental and Resource Economics*, 68(1):197–225.

Howarth, R. B. (1991a). Intergenerational competitive equilibria under technological uncertainty and an exhaustible resource constraint. *Journal of Environmental Economics and Management*, 21(3):225–243.

Howarth, R. B. (1991b). Intertemporal equilibria and exhaustible resources: an overlapping generations approach. *Ecological Economics*, 4(3):237–252.

Howarth, R. B. and Norgaard, R. B. (1990). Intergenerational resource rights, efficiency, and social optimality. *Land Economics*, 66(1):1–11.

Howarth, R. B. and Norgaard, R. B. (1992). Environmental valuation under sustainable development. *American Economic Review*, 82(2):473–477.

Jensen, S. and Traeger, C. P. (2014). Optimal climate change mitigation under long-term growth uncertainty: Stochastic integrated assessment and analytic findings. *European Economic Review*, 69:104–125.

John, A., Pecchenino, R., Schimmelpfennig, D., and Schreft, S. (1995). Short-lived agents and the long-lived environment. *Journal of Public Economics*, 58(1):127–141.

Judd, K., Maliar, L., and Maliar, S. (2011). Numerically stable stochastic simulation approaches for solving dynamic economic models. *Quantitative Economics*, 2:173–210.

Kavuncu, Y. O. and Knabb, S. D. (2005). Stabilizing greenhouse gas emissions: assessing the intergenerational costs and benefits of the Kyoto Protocol. *Energy Economics*, 27(3):369–386.

Knutti, R., Rugenstein, M. A., and Hegerl, G. C. (2017). Beyond equilibrium climate sensitivity. *Nature Geoscience*, 10(10):727–736.

Kotlikoff, L. J., Kubler, F., Polbin, A., Sachs, J. D., and Scheidegger, S. (2019). Making carbon taxation a generational win win. Technical report, National Bureau of Economic Research.

Krueger, D. and Kubler, F. (2006). Pareto-improving social security reform when financial markets are incomplete!? *American Economic Review*, 96(3):737–755.

Krusell, P. and Smith, Jr, A. A. (1998). Income and wealth heterogeneity in the macroeconomy. *Journal of Political Economy*, 106(5):867–896.

Kubler, F. and Scheidegger, S. (2019). Self-justified equilibria: Existence and computation. *Available at SSRN 3494876*.

Lemoine, D. and Traeger, C. P. (2016). Economics of tipping the climate dominoes. *Nature Climate Change*, 6(5):514.

Maliar, L., Maliar, S., Taylor, J., and Tsener, I. (2015). A tractable framework for analyzing a class of nonstationary markov models. Technical report, National Bureau of Economic Research.

Marcet, A. (1988). Solution of nonlinear models by parameterizing expectations. Technical report, Carnegie Mellon University.

Marcet, A. and Lorenzoni, G. (2001). The parameterized expectations approach: Some practical issues. In Marimon, R. and Scott, A., editors, *Computational Methods for the Study of Dynamic Economies*. Oxford University Press, Chicago.

Marcet, A. and Marshall, D. A. (1994). Solving nonlinear rational expectations models by parametrized expectations: Convergence to Stationary Solutions. Discussion paper 91, Federal Reserve Bank of Minneapolis.

Marini, G. and Scaramozzino, P. (1995). Overlapping generations and environmental control. *Journal of Environmental Economics and Management*, 29(1):64–77.

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

Nordhaus, W. D. (2015). Climate clubs: overcoming free-riding in international climate policy. *The American Economic Review*, 105(4):1339–1370.

Nordhaus, W. D. (2017). Revisiting the social cost of carbon. *Proceedings of the National Academy of Sciences*, page 201609244.

Papageorgiou, C., Saam, M., and Schulte, P. (2017). Substitution between clean and dirty energy inputs: A macroeconomic perspective. *Review of Economics and Statistics*, 99(2):281–290.

Pecchenino, R. and John, A. (1994). An overlapping generations model of growth and the environment. *The Economic Journal*, 104(427):1393–1410.

Roe, G. H. and Baker, M. B. (2007). Why is climate sensitivity so unpredictable? *Science*, 318(5850):629–632.

Scheidegger, S. and Bilionis, I. (2019). Machine learning for high-dimensional dynamic stochastic economies. *Journal of Computational Science*, 33:68 – 82.

Traeger, C. P. (2019). Ace–analytic climate economy (with temperature and uncertainty). *Working paper*.

Weitzman, M. L. (2009). On modeling and interpreting the economics of catastrophic climate change. *The Review of Economics and Statistics*, 91(1):1–19.

Weitzman, M. L. (2012). Ghg targets as insurance against catastrophic climate damages. *Journal of Public Economic Theory*, 14(2):221–244.