FONDAZIONE ENI ENRICO MATTEI

NOTA DI LAVORO

16.2017

Optimal Clean Energy R&D Investments Under Uncertainty

Giacomo Marangoni, FEEM, CMCC and Politecnico di Milano
Gauthier De Maere, FEEM
Valentina Bosetti, FEEM, CMCC and Bocconi University

Mitigation, Innovation and Transformation Pathways
Series Editor: Massimo Tavoni

Optimal Clean Energy R&D Investments Under Uncertainty

By Giacomo Marangoni, FEEM, CMCC and Politecnico di Milano
Gauthier De Maere, FEEM
Valentina Bosetti, FEEM, CMCC and Bocconi University

Summary

The availability of technology plays a major role in the feasibility and costs of climate policy. Nonetheless, technological change is highly uncertain and capital intensive, requiring risky efforts in research and development of clean energy technologies. In this paper, we introduce a two-track method that makes it possible to maintain the rich set of information produced by climate-economy models while introducing the dimension of uncertainty in innovation efforts, without succumbing to computation complexity. In particular, we solve the problem of an optimal R&D portfolio by employing Approximate Dynamic Programming, through multiple runs of an integrated assessment model (IAM) for the purpose of computing the value function, and expert elicitation data to quantify the relevant uncertainties. We exemplify the methodology with the problem of evaluating optimal near-term innovation investment portfolios in four key clean energy technologies (solar, biofuels, bioelectricity and personal electric vehicle batteries), taking into account the uncertainty surrounding the effectiveness of innovation to improve the performance of these technologies. We employ an IAM (WITCH) which has a fairly rich description of the energy technologies and experts' beliefs on future costs for the above-mentioned technologies. Focusing on Europe and its short-term climate policy commitments, we find that batteries in personal transportation dominate the optimal public R&D portfolio. The resulting ranking across technologies is robust to changes in risk-aversion, R&D budget limitation and assump- tions on crowding out of other investments. These results suggest an important upscaling of R&D efforts compared to the recent past.

Keywords: Energy, Innovation, Technological Change, Uncertainty, Climate Policy

JEL Classification: O30, O33, Q40, Q41, Q50, Q55

Address for correspondence:
Giacomo Marangoni
Fondazione Eni Enrico Mattei
C.so Magenta 63
20123 Milan
Italy
E-mail: giacomo.marangoni@feem.it

Optimal Clean Energy R&D Investments Under Uncertainty

Giacomo Marangoni[1,2,3] Gauthier De Maere[1]
Valentina Bosetti[1,2,4]

March 15, 2017

[1]Fondazione Eni Enrico Mattei (FEEM), Italy.
[2]Centro Euro-Mediterrano sui Cambiamenti Climatici (CMCC), Italy.
[3]Politecnico di Milano, Department of Management, Economics and Industrial Engineering, Italy.
[4]Bocconi University, Via Roentgen 1, 20136 Milan, Italy.

Abstract

The availability of technology plays a major role in the feasibility and costs of climate policy. Nonetheless, technological change is highly uncertain and capital intensive, requiring risky efforts in research and development of clean energy technologies.

In this paper, we introduce a two-track method that makes it possible to maintain the rich set of information produced by climate-economy models while introducing the dimension of uncertainty in innovation efforts, without succumbing to computation complexity. In particular, we solve the problem of an optimal R&D portfolio by employing Approximate Dynamic Programming, through multiple runs of an integrated assessment model (IAM) for the purpose of computing the value function, and expert elicitation data to quantify the relevant uncertainties. We exemplify the methodology with the problem of evaluating optimal near-term innovation investment portfolios in four key clean energy technologies (solar, biofuels, bioelectricity and personal electric vehicle batteries), taking into account the uncertainty surrounding the effectiveness of innovation to improve the performance of these technologies. We employ an IAM (WITCH) which has a fairly rich description of the energy technologies and experts' beliefs on future costs for the abovementioned technologies.

Focusing on Europe and its short-term climate policy commitments, we find that batteries in personal transportation dominate the optimal public R&D portfolio. The resulting ranking across technologies is robust to changes in risk-aversion, R&D budget limitation and assumptions on crowding out of other investments. These results suggest an important upscaling of R&D efforts compared to the recent past.

1 Introduction

A successful climate change mitigation strategy will require significant improvements of existing technologies, and the introduction in the market of alternatives currently available only in the lab, to reduce energy consumption and energy carbon intensity at acceptable costs. Applying deterministic approaches, researchers have been developing complex Integrated Assessment Models (IAMs) to inform these types of decisions, and to identify ideal decarbonization pathways (Clarke et al., 2014; Marangoni and Tavoni, 2014). Experience with these models has shown that the availability of technology plays a major role in the feasibility and costs of facing the challenge of stringent climatic targets (Kriegler et al., 2014). Nonetheless, technological change is highly uncertain and capital intensive, requiring risky efforts in research and development (R&D). Thus a more appropriate approach should account for this source of uncertainty. The fact that these efforts may or may not lead to technological breakthroughs has important implications when considering energy R&D investment strategies for a low-carbon future.

Previous literature has already highlighted how the joint modelling of endogenous technical change and uncertainty has important quantitative and qualitative impacts on optimal technological policies for climate change (Baker and Shittu, 2008). When accounting for the uncertain effectiveness of R&D investments, results can drastically differ from their deterministic counterparts, in terms of both magnitude and composition. In addition to this, an extensive literature on expert elicitation of energy technologies (Baker et al., 2015a; Chan et al., 2011; Anadón et al., 2012; Nemet et al., 2016) emphasizes the significant uncertainty that experts attach to R&D investments, as well as the huge disagreement across experts. This uncertainty cannot be neglected, especially if we consider the significant impact of future technological costs on the implementation of stringent climate policies (Bosetti et al., 2015).

So far, the most common approach has been to include uncertainty within a simple analytical framework, with inputs derived from the output of IAMs. Blanford (2009) captures the essential elements underlying the relationship between R&D investment and research outcomes, where the latter are assessed by running an IAM (MERGE, Manne et al. (1995)) in a variety of technology scenarios representing outcomes of alternative R&D programs. Optimal portfolios are then calculated by linking R&D investments to a probability distribution over alternative outcomes. Likewise, Bosetti and Tavoni (2009) develop a simple analytical model, with two time periods and two technologies, which mimics a social planner who minimizes costs by choosing optimal abatement and innovation efforts consistent with a given environmental target. Uncertainty is introduced by modeling the R&D outcome on the abatement cost of a carbon-free breakthrough technology (backstop) as uncertain. A stochastic version of WITCH is devised to account for such uncertainty, but only two states of nature for the effectiveness of R&D in a single technology can be introduced while maintaining the model computationally tractable.

The need to consider problems with multiple technologies, coupled with more complex IAMs, has spurred researchers to find new ways to overcome the resulting "curse of dimensionality." In a recent study, Baker et al. (2015b) consider four sets of probabilistic distributions, related to different elicitation teams of experts, conditional to three funding levels per set, and five technologies. Furthermore, the economic interactions of technologies are estimated through a large IAM. To alleviate the computational burden, the authors use importance sampling. A single sampling distribution is derived from all those available, and the resulting set of samples is run through the IAM. The sampled output is then used to evaluate alternative portfolios.

In the different context of assessing the social cost of carbon, Cai et al. (2013) show the application of an alternative approach, Approximate Dynamic Programming, to account for uncertainty in a IAM. The authors jointly model the uncertain elements of catastrophic climate change damages and annual economic productivity within a dynamic stochastic general equilibrium version of DICE, a widely accepted IAM. The problem is solved within the framework of dynamic programming, where the value function given by the solution of the original model is approximated with a finitely parameterized collection of functions. A similar approach allows Lemoine and Traeger (2014) to calculate optimal climate policies in DICE, including the endogenous possibility of climatic tipping points, their welfare implications, and learning about their trigger temperatures. Different variants of the Approximate Dynamic Programming (ADP) framework have been successfully applied to studying the policy impact of other decision-dependent cost uncertainties in DICE, as in Webster et al. (2012), providing a rough approximation of the technology R&D case we tackle here.

In this paper, we employ the ADP algorithm proposed by Cai et al. (2013), which allows us to keep the complexity of the IAM results intact and, at the same time, to account for the uncertain effectiveness of R&D efforts in four innovative low-carbon technologies. To compute the value function we use a fairly complex IAM (WITCH). In order to quantify the intrinsic uncertainty concerning learning rates, costs and efficiency parameters, we resort to recent expert elicitations (Bosetti et al., 2011). To provide novel and robust insights on the optimal portfolio of clean energy R&D investments, we perform a set of experiments.

First, we define the optimal level and composition of a public R&D portfolio of investments in four key clean energy technologies, given experts' judgments on their future probabilistic costs, the potential economic and technological implications of such costs (as modelled by a complex IAM), and the uncertain effect of R&D investments on costs.

Second, we study how the portfolio composition changes when considering different limits on the RD&D budget, different risk-aversion preferences, and different assumptions about the characteristics of the R&D program.

While providing policy relevant results, the main goal of the paper is to introduce a fairly general method that can be easily adapted to using alternative IAMs, or a collection of them, to approximate the value function, as well as different expert elicitations or historically based data to inform the probabilistic relationship between R&D and the future evolution of technological costs.

2 Optimal R&D Portfolio Problem

The problem of finding the optimal near-term investments in R&D during a first period T_1 that have an uncertain effectiveness for a period T_2 can be formulated as a two-stage stochastic program. This can be schematically described as:

$$U_{T_1}(\boldsymbol{I}_{T_1}) + \mathbb{E}\left[U_{T_2}(\boldsymbol{I}_{T_1}, \boldsymbol{\lambda})\right]. \quad (1)$$

where $\boldsymbol{I}_{T_1} = \boldsymbol{I} = [I_1, \ldots, I_n]$ represents the vector of R&D investments in the different n technologies; $U_{T_1}(\boldsymbol{I}_{T_1})$ is the objective function over the first period T_1 and $U_{T_2}(\boldsymbol{I}_{T_1}, \boldsymbol{\lambda})$ is the objective function over the second period T_2. In the context of R&D policy making, the objective is to maximize the sum of these two terms. The first term decreases deterministically with the expenditure in R&D, $\boldsymbol{I}$, as the resulting benefits are assumed to materialize only after the first stage. The second term models the welfare associated with the stream of consumption in the second period, given the realization of future technological costs. This term increases with the expenditures in R&D $\boldsymbol{I}$ done in the first period, as costs are reduced with the cumulated level of knowledge, and more wealth can be generated. Nonetheless, benefits depend on the realization of the effectiveness of the R&D investments. This uncertain effectiveness can be conceptualized by introducing stochastic learning rates $\boldsymbol{\lambda} = [\lambda_1, \ldots, \lambda_n]$, which capture how knowledge turns into cost reductions (Gritsevskyi and Nakićenovi, 2000). Optimal R&D investments should then be chosen by the policy maker by trading off the benefits of shifting the distributions of future technological costs towards the lower end, with the burden of sustaining those investments today.

In order to approach the problem of Eq. (1), two elements are necessary (Baker et al., 2015b). The first is the quantification of the stochastic relationship between R&D investments and their effect on future technological performance (i.e. the distribution of learning rate $\boldsymbol{\lambda}$); the second relates to the calculation of the stream of benefits associated with the various realizations of technological costs and efficiency parameters.

This paper shows how the framework just described can be implemented to support today public R&D investment decisions in clean low-carbon energy technologies. Both empirical analyses of past R&D program effectiveness (Wiesenthal et al., 2012) and expert elicitations (Morgan et al., 1992) can be used to quantify the probabilistic relationship between R&D investments and their effectiveness. Here, we use data from an expert elicitation regarding the future cost and efficiency of key energy technologies and how these parameters might be affected by changes in R&D efforts (Bosetti et al., 2011). We focus on the impact of R&D investments on future costs of four technologies related to low-carbon energy supply:

- the cost of electricity produced with solar technologies;
- the production cost of liquid biofuels;
- the module cost of batteries for light-duty vehicles;
- the cost of electricity produced with biomass.

To quantify future welfare implications of different techno-economic-climatic scenarios, the scientific community typically resorts to integrated assessment models (IAMs). Here, we employ WITCH (World Induced Technical Change Hybrid Model), a general equilibrium IAM that can project forward the societal implications of energy technology costs (Emmerling et al., 2016; Bosetti et al., 2009, 2006). The space of possible future cost realizations is sampled to provide a discrete set of input configurations for WITCH. The simulations yield a set of economic outcomes which are then used to approximate a continuous value function.

The choice of WITCH among the many existing IAMs was mainly driven by its fairly rich technological and economic description, and by its numerous previous applications to the study of endogenous innovation in the energy system (Marangoni and Tavoni, 2014; Bosetti and Tavoni, 2009). The model divides the worldwide economy into 13 regions, whose main macroeconomic variables are represented through a top-down inter-temporal optimal growth structure, while the energy sector is detailed in a bottom-up fashion. The different regions behave as forward-looking agents optimizing their welfare in a non-cooperative game-theoretic set-up. Actions of each agent interrelate through several externalities, such as dependence on exhaustible natural resources and trade of oil. While the focus of this paper is on the EU15+EFTA region, economic scenarios involve assumptions and optimization for all the other regions. For our application to be policy relevant, we assume that countries represented in the model optimize welfare while obeying a lenient climate change policy, in line with recent UNFCCC negotiated targets[1].

3 Two-stage stochastic program

The first stage decision concerns which technologies to promote by investing a dedicated amount in R&D during the 2010-2030 period. This happens before the realization of the effectiveness of this R&D investment in 2030. In this first

[1]We assume that countries will pursue similar levels of climate mitigation stringency for the rest of the century. This target entails an expected increase in global average temperature of around 2.8°C above pre-industrial levels by the end of the century. The probability that the temperature increase will exceed 2°C, the threshold commonly considered as safe by the scientific community to avoid irreversible climate changes, is very high (88%-97%, Kriegler et al. (2013)). While climate policy makers will hopefully commit to stronger and more coordinated actions in the future, the stringent pledges we consider here reflect realistic national near-term efforts, also in line with the commitment shown by the latest Intended Nationally Determined Contributions (INDCs) circulated in the COP21 in Paris.

stage, the decision-maker allocates a time-invariant R&D budget I_j yearly between 2010 and 2030, for each modelled technology $j \in \mathcal{J} = \{j_1, ..., j_J\}$. R&D expenditures reduce available consumption of final good Q_t in period $t \in T_1 = \{2010, 2015, 2020, 2025\}$ by subtraction from a counterfactual consumption level $\bar{Q}_t$ without R&D:

$$Q_t = \bar{Q}_t - r \sum_j I_j \qquad t \in T_1 \tag{2}$$

Public investments in energy R&D crowd out investments in other R&D's, which have a social rate of return r times higher than that of private investments. It is assumed that 1 dollar of I_j costs $r = 4$ dollars otherwise usable for direct consumption or private investments (Popp, 2004). Reference consumption $\bar{Q}_t$ is a baseline counterfactual calculated by the integrated assessment model WITCH, if we assume that there will be no explicit R&D expenditures and median cost realizations after 2030 for the technologies in $\mathcal{J}$. Any repercussion of I on the economy other than the one in Eq. (2) is assumed to be negligible. This is reasonable as R&D investments constitute a tiny fraction of overall GDP, and benefits from R&D may take some time to materialize, in this case a couple of decades. t is defined on a discretized time horizon of periods of $T_\Delta = 5$ years. Over this time horizon, knowledge in each technology j builds up according to the usual capital law of motion. Starting from an initial value $K_{0,j}$, the R&D stock in period t is given by a fraction $(1 - \delta_R)^{T_\Delta}$ of the stock in the previous period, plus the flow of new ideas due to T_Δ years of investments:

$$K_{j,t+T_\Delta} = K_{j,t}(1 - \delta_R)^{T_\Delta} + T_\Delta I_j. \tag{3}$$

The initial R&D stock $K_{0,j}$ in 2010 is estimated from IEA (2015) by applying the same accumulation dynamics to historical R&D budgets, separately for each region. Yearly obsolescence of R&D δ_R is assumed to be 5% for all technologies.

The link between knowledge accumulated in the first stage $\boldsymbol{K} = [K_1, ..., K_J]$ by 2030, and the cost $\boldsymbol{C} = [C_1, ..., C_J]$ in 2030, follows a one-factor learning curve. In particular, the cost $\boldsymbol{C}$ can be thought of as the sum of a floor cost $\boldsymbol{C_f}$ with the initial cost $\boldsymbol{C_0}$ scaled according to the upscaling in stock of knowledge $\boldsymbol{K}$ over time, and to a learning rate $\boldsymbol{\lambda} = [\lambda_1, ..., \lambda_J]$:

$$\boldsymbol{C}(\boldsymbol{I}, \boldsymbol{\lambda}(\boldsymbol{\omega})) = \boldsymbol{C_f} + \boldsymbol{C_0} \left(\frac{\boldsymbol{K}(\boldsymbol{I})}{\boldsymbol{K_0}} \right)^{-\boldsymbol{\lambda}(\boldsymbol{\omega})}. \tag{4}$$

$\boldsymbol{\lambda}$ is assumed to be revealed in 2030 and depends on the realization $\boldsymbol{\omega} = [\omega_1, ..., \omega_J]$. This parameter drives the uncertainty of the problem of choosing investments today. If we omit the dependency of $\boldsymbol{\lambda}$ on $\boldsymbol{\omega}$ for brevity, optimal investments $\boldsymbol{I}$ maximize the following utility function:

$$\begin{aligned} U(\boldsymbol{I},\boldsymbol{\lambda}) &:= U_{T_1}(\boldsymbol{I}) + \mathbb{E}\left[U_{T_2}(\boldsymbol{I},\boldsymbol{\lambda})\right] \\ U_{T_1}(\boldsymbol{I}) &:= \sum_{t\in T_1} u_t(Q_t(\boldsymbol{I})) \\ U_{T_2}(\boldsymbol{I},\boldsymbol{\lambda}) &:= \sum_{t\in T_2} u_t(Q_t(\boldsymbol{C}(\boldsymbol{I},\boldsymbol{\lambda}))) = V(\boldsymbol{C}(\boldsymbol{I},\boldsymbol{\lambda})) \end{aligned} \tag{5}$$

In the first part of the objective function $U_{T_1}(\boldsymbol{I})$, utility at time t is a concave function u_t increasing with consumption Q_t, and decreasing with $|\boldsymbol{I}|$. A constant relative risk aversion equal to the inverse of the elasticity of intertemporal substitution η is implied. The dependency of u_t on time is due to changing levels of population L_t and discount factor β_t[2].

The second part of the objective function $U_{T_2}(\boldsymbol{I},\boldsymbol{\lambda})$ represents future welfare from 2030 till the end of the time horizon (2150 in WITCH). It depends on the sum of future consumption levels Q_t over period T_2, according to the same temporal utility function u_t as before. This time consumption depends explicitly on realized technological costs $\boldsymbol{C}$, which in turn depend on the knowledge previously accumulated and on the actual sample outcomes $\boldsymbol{\omega}$ of the stochastic learning parameter $\boldsymbol{\lambda} = [\lambda_1, ..., \lambda_J]$. Function V represents the overall functional dependency of utility on technological costs, and is implemented via the WITCH computational model by observing the economic and energy-related consequences of changing cost assumptions $\boldsymbol{C}$.

Multiple technologies and continuous distributions for $\boldsymbol{\lambda}$ lead to an intractable problem if any change in $\boldsymbol{I}$ or $\boldsymbol{\lambda}$ (hence in $\boldsymbol{C}$) needs to be propagated into the solution of a new instance of a complex integrated assessment model like WITCH. We adopt an ADP approach (Cai et al., 2013) by substituting $V(\cdot)$ with an approximating function much easier to calculate. This replacement, along with the assumption that uncertainty about technological change in the second stage does not impact the first stage objective function $U_{T_1}(\boldsymbol{I})$, makes the R&D portfolio problem numerically solvable in reasonable times. Finally, the expectation in Eq. (5) is translated into an average of over 1000 scenarios for $\boldsymbol{\lambda}$, obtained via latin hypercube sampling of its distribution.

To show that the problem is well formulated and has a unique solution, we would need to prove that the utility is strictly concave. In section A.1 in the appendix we study and provide proof of concavity in the case of one technology. The case for multiple technologies is treated only numerically, as the intuition behind the well-posedness of the problem would be obscured by the required cumbersome analytical derivation.

The optimization problem changes slightly to accommodate for different risk-aversion preferences. As section A.2 in the appendix illustrates, the problem turns into the following minimization:

[2]In period t, utility is multiplied by a standard geometric discount factor $\beta_t = 1/(1+\rho)^{t-2005}$, with a pure rate of time preference ρ equal to 1%. The specific expression we consider for u_t is $L_t\beta_t/(1-\eta)((Q_t/L_t)^{1-\eta}-1)$

$$\min \sum_{t \in T_1} G_t(Q_t(\boldsymbol{I})) + \mathbb{E}\left[(H(\boldsymbol{C}(\boldsymbol{I}, \boldsymbol{\lambda})))^{\frac{1-\alpha}{1-\eta}} \right]^{\frac{1-\eta}{1-\alpha}} \tag{6}$$

for some appropriate functions G_t and H (related to utility in T_1 and T_2 respectively). This formulation allows the risk-aversion parameter α to be made explicit and independent from the inverse of the elasticity of intertemporal substitution η.

3.1 Learning rate estimation

Since we are considering relatively novel technologies, with scarce historical data and high potential for improvement, extrapolating historical trends may be inadequate to represent their future behaviors. We rely instead on expert elicitation to assess probabilistic distributions of future costs. Distributions are estimated from the data collected by the ICARUS project (Bosetti et al., 2011). For this project, leading experts from the academic world, the private sector and international institutions took part in a survey designed to collect probabilistic information on the role of R&D investments in lowering costs and increasing penetration of 8 carbon-free technologies. The survey could focus only on a limited amount of R&D budgets for each technology. In our framework we want to be able to search for the optimal budget allocation in a continuum of possibilities. This is possible by introducing and calibrating a particular model for technical change, such as the single-factor learning model considered here (Eq. (4)). Then, uncertainty is assumed to lie exclusively in the learning rate λ.

The goal is to translate cost distributions, given by the experts, into distributions of learning rates. Let the random variable C_s be the cost in 2030 of one technology under an R&D scenario s, chosen among 3 future R&D budget scenarios. The scenarios represent an increase of 0%, 50% and 100% with respect to baseline levels, in line with the protocol used for the expert elicitation. The ensemble of experts' predictions is summarized into N samples of $C_{j,s}$ cumulative distribution function (CDF). The resulting empirical CDFs are illustrated for one technology (battery) in the first panel of Figure 1. By inverting Eq. (4), N samples for λ_s CDF are obtained. If the one-factor model with uncertain learning perfectly represented the experts' aggregated view, these CDF would overlap, as the learning rate (hence its CDF) would not depend on the scenario: while investments do affect costs, they should be independent from the learning rates according to Eq. (4). As shown in the second panel of Figure 1, this modelling assumption fits well in the case of batteries, and remains reasonable for the other technologies as well. Taking the mean across s of λ_s empirical distributions and fitting a Weibull distribution, we identify a parametric description of the uncertain learning rates. Using Eq. (4) again, a parametric description of the uncertain costs as a function of investments can be derived. The third panel of Figure 1 confirms the agreement between fitted and empirical distributions for the original 3 budget scenarios, and demonstrates how novel budget scenarios can now be explored thanks to the parametric description obtained. The whole estimation process is further documented in Procedure 1 in the Appendix, while similar plots to those of Figure 1 are reported for all technologies in Figure 6 in the Appendix.

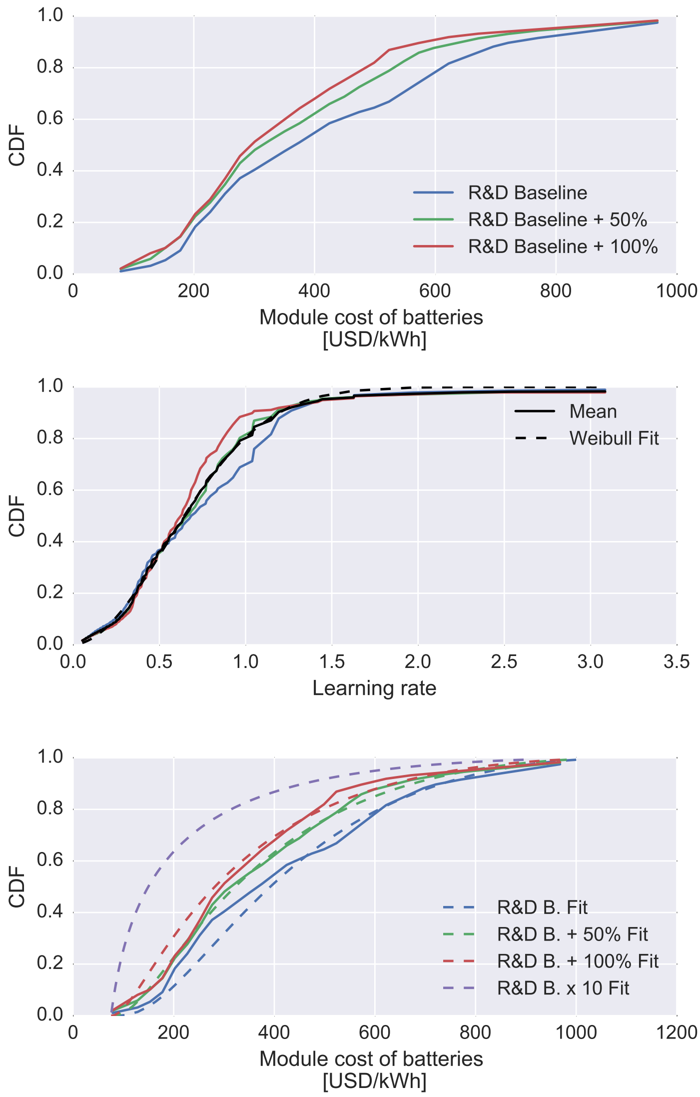

Figure 1: First panel: empirical CDFs of 2030 battery costs as elicited from experts, one R&D budget for each color. Second panel: corresponding learning rate distributions, according to the one-factor learning curve model, plus a Weibull fit of the mean empirical CDF. Third panel: empirical CDFs of costs along with their fitted versions, for the 3 original R&D budgets plus an extra one.

3.2 Value function interpolation

In order to apply the Approximate Dynamic Programming paradigm, we need to construct the value function, i.e. a continuous function approximating the WITCH regional welfare response to technological cost $\boldsymbol{C}$ in the space of interest. We start by evaluating WITCH on a representative discrete subset of costs combinations. The extreme of the ranges of this space are reported in Table 3.2 under the min and max cost columns. Current cost $C_{0,j}$, used until 2030, is just below the maximum value. Costs are assumed to reveal themselves in 2030, and are obtained by sampling a value in the ranges. Afterwards, they decay autonomously a further 20% by 2060, and then they remain constant. Floor cost $C_{f,j}$, appearing in the learning equation (4), is slightly above the minimum value of the range. The margins from extreme values, especially on the lower end, keep the R&D portfolio program away from extreme boundary behaviors.

Table 1: Table summarizing minimum, maximum, current and floor costs of the four technologies considered.

Technology j	Unit	Min cost ($C_{m,j}$)	Floor cost ($C_{f,j}$)	Current cost ($C_{0,j}$)	Max cost ($C_{M,j}$)
Solar	cUSD/kWh	2	3	27.8	28
Biofuels	USD/lge	0.05	0.08	2.98	3
Batteries	USD/kWh	50	75	1019	1025
Bioelectricity	cUSD/kWh	3	4.5	24.9	25

We follow the methodology presented by Cai et al. (2013) and construct the value function based on Hermite approximation, including both the Lagrange data and the slope information (available using the dual variables of the technological cost constraints in the welfare maximization). To do so, ten samples are picked along each of the cost dimensions according to a Chebysev nodal formula[3], which is known to provide greater stability for polynomial interpolation.

As we study in this paper the impact on four different technologies, a total of ten thousands runs of WITCH is required to perform the approximation (which is eventually based on 10,000 welfare and 40,000 directional derivative values). The exact procedure is described in section A.4 in the appendix.

4 Results

With our model we run two scenarios: "Stochastic", where the optimal portfolio program is solved assuming that the learning rate λ is sampled from its fitted Weibull distribution, and "Certainty equivalent", where we compute the

[3]The i-th sample cost for technology j is chosen as: $C_j^{(i)} = \frac{C_{m,j}+C_{M,j}}{2} + \frac{C_{M,j}-C_{m,j}}{2}\cos\left(\pi\frac{i-0.5}{10}\right)$.

portfolio assuming that λ is deterministic and equal to the mean value of its stochastic distribution. By comparing the utility of using a "Stochastic" strategy with the one resulting from a "Certainty equivalent" approach, averaged across a sampled set of learning rate outcomes, we can capture the value of including more information in the solution process, which is also called the Expected Value of Better Information (EVBI) in the literature (Baker and Peng, 2011). As expected, the "Stochastic" solution yields a greater utility. While the implications on average utility in our numerical case are relatively small, the strategies look very different with and without probabilistic information. To further explore the case of a "Stochastic" approach, different values of the crowding out factor r, the risk-aversion parameter α, and the total allowable budget ($\max \sum_i \boldsymbol{I}_i$) are considered.

Figure 2 compares yearly public R&D investments per technology under the two simulated scenarios with actual historical data on energy R&D (labelled "Historical"). The latter come from averaging 2010 to 2014 R&D investment flows reported by IEA in the relevant technological categories (IEA, 2015). Also for the estimation of the initial R&D capital, values in the period 2010-2014 were averaged, as a way to incorporate in the model the most recent available information. Results from both optimization runs imply a stark break with past trends, in terms of both total R&D expenditures (left hand side panel) and composition of the portfolio (right hand side panel). The "Certainty equivalent" case implies an optimal future budget 27 times higher than the historical 2010-2014 average of 485 million USD. Uncertainty seems to duplicate optimal R&D efforts, for a total of 28.5 billion USD. In increasing order, the three budgets constitute 0.003%, 0.10% and 0.20% of GDP in 2010 for the EU15+EFTA region. While GDP is estimated to grow in the first-stage, the fixed yearly "Stochastic" budget would remain above 0.10% of the overall economic output.

A clear break with history happens also in terms of shares. Most recent European R&D investments were devoted to solar technologies, with biofuels in the second place, followed by bioelectricity and lastly batteries for personal transport. Our stochastic solution suggests a quite different scenario, where batteries dominate the portfolio with a share of 80% in the stochastic case, and likewise in the deterministic one. This is equivalent to an upscaling of two to three orders of magnitude in R&D investments in batteries, with respect to current efforts. A minor role is played by other technologies, with different rankings depending on whether or not the full distribution is taken into account. In particular, shares in the "Certainty equivalent" case are 76.8%, 12.4%, 8.3% and 2.5% for batteries, solar, biofuels and bioelectricity respectively. In the "Stochastic" case, the share of solar and bioelectricity budgets (9.5% and 1.8% respectively) are slightly diminished, mostly in favor of battery (80.3%). Biofuels (8.5%) keep their share almost unchanged.

Several reasons may justify these results. First, the expected net present value of future welfare is much more sensitive to changes in future costs of bat-

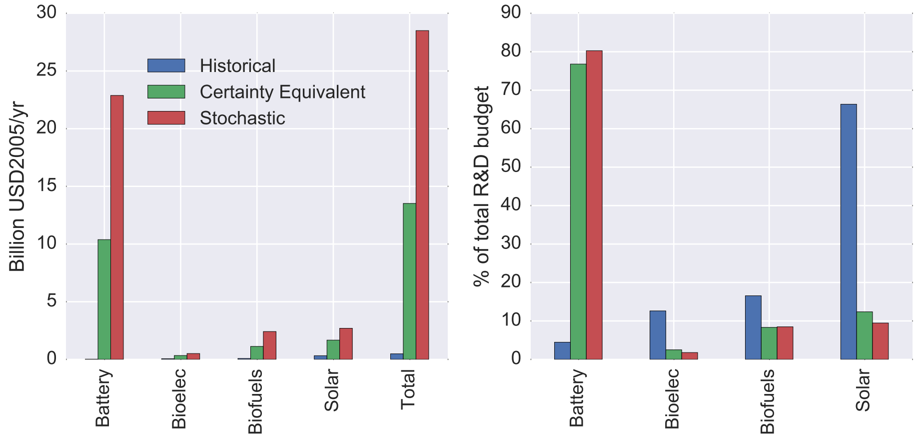

Figure 2: R&D portfolio composition. The left-hand side reports yearly R&D investments per technology for the two scenarios and the historical 2010-2014 average as reported in the IEA RD&D Database. The right-hand side shows the percentage of total R&D investments in each of the four technologies.

teries than in all the other costs (see Figure 7). The non-electric sector, and in particular the personal transport sector, is traditionally considered one of the most difficult and expensive to decarbonize (Luderer et al., 2011; Bosetti et al., 2015). Through electrification of vehicles, Europe can expand to transport emissions the benefits from the already widespread efforts in support of clean power generation. Europe's determined contribution can then be achieved at contained costs. Electrification of this kind is possible only with an adequate battery technology in place, which requires a considerable increase in the R&D from the status quo. Second, according to the experts' judgment, the probabilistic distribution of future battery costs seems to be less sensitive to an upscaling in investments when compared to other technologies. As reported in the last column of Figure 6, a ten-fold increase in battery investments yields the smallest change in the cumulative distribution function of 2030 costs. Finally, batteries have received marginal attention in terms of European R&D funds in the recent past, making the initial cost quite high. It is important to remark that the scope of our analysis is limited to public R&D expenditures in Europe. By doing so, we are potentially neglecting spillovers of innovation from the private sectors, and from R&D efforts in other countries. This might lead to an over-estimation of optimal regional public R&D needs, which could partially justify the deviation from historical trends.

Comparing the "Stochastic" and "Certainty equivalent" solutions, we can grasp the effect of accounting for uncertainty in the optimal response. While uncertainty seems to have a minor effect on the shares of the portfolio, its role is clear in the upscaling of battery investment levels. A precautionary mechanism emerges from the optimization, where the risks of low learning in batteries are strongly hedged against by increased investments. This is done to the detriment of the power sector, which receive fewer percentage points of the budget. Biofuels keep their minor role in complementing batteries for decarbonizing the transport sector.

Next, we report results for different sensitivity tests aimed at understanding the robustness of the optimal stochastic solution to reasonable changes in key components of the model. As discussed in Section 3, we employ the assumption that 1 USD of public investments in energy R&D crowds out r dollars of investments in other R&D's, with a nominal value of 4 (Popp, 2004). Figure 3 shows how results change in response to different assumptions about the crowding out factor. As r decreases, opportunity costs of energy R&D investments are lower, hence greater energy R&D budgets can be optimally allocated, and more so again in the battery sector. Increasing r, on the other hand, increases the social cost of energy R&D investments, which become less attractive. Going from $r = 2$ to $r = 6$ means cutting more than half of the investments. A slight convexity makes this behavior marginally decrease towards greater r values. We remark that the concern about energy R&D crowding out other forms of R&D might be minimal according to some empirical evidence (Popp and Newell, 2009), so that r is more likely gravitating towards 2 rather than 6.

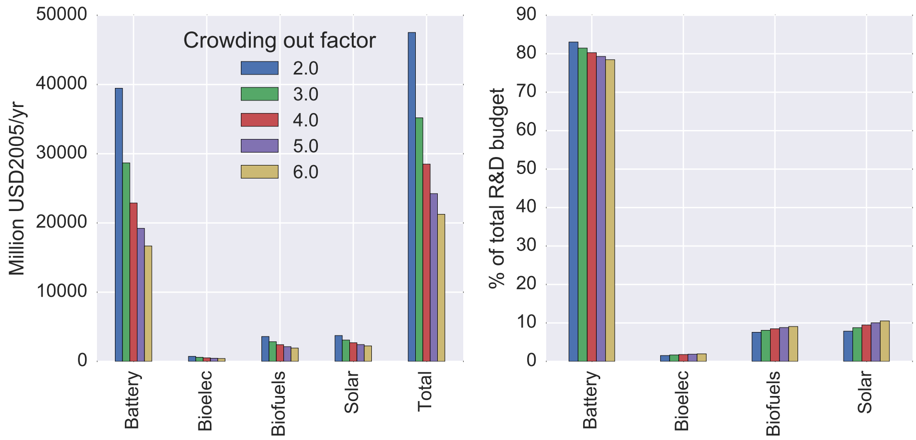

Figure 3: Sensitivity to different crowding-out assumptions of the optimal R&D portfolio under uncertainty. Results are presented in terms of both yearly absolute investments levels (left-hand side) and corresponding shares (right-hand side).

Another important parameter is the one affecting the propensity of decision-makers to avoid risk. In the nominal solution we assumed the relative risk aversion parameter α to be equal to the inverse of the elasticity of intertemporal substitution $\eta = 1.5$. Empirical estimates of α from observed trade behavior can be in the 10-20 range. Figure 4 shows the impact of increasing α on the optimal solution of the stochastic problem, using the modified utility derived in Section A.2. Greater risk aversion justifies greater R&D investments, and the additional budget is directed towards batteries. As discussed previously, the model is particularly sensitive to the uncertain learning rate of batteries, and reducing the uncertainty of the problem involves mostly reducing the uncertainty in batteries' future cost realizations. Total and battery investments increase linearly in α, with an average of 0.4% gained per added unit of α.

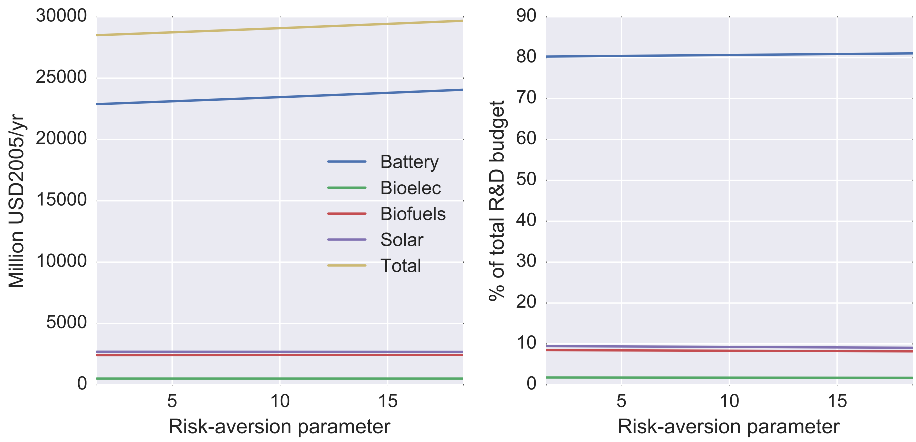

Figure 4: Sensitivity to different risk aversion assumptions of the optimal R&D portfolio under uncertainty. Results are presented in terms of both yearly absolute investments levels (left-hand side) and corresponding shares (right-hand side).

Finally, the actual budget for energy R&D may be limited by other constraints, not explicitly modelled in this work. Current budget deficits and other financial constraints might impose a total energy R&D budget considerably smaller than the optimal one. Figure 5 illustrates the change in R&D allocation when the total R&D budget is constrained. The constraint is expressed as a fraction of the total R&D budget that would be otherwise optimal. Along with total investments, the levels of the individual technology budgets also scale proportionally, except for the smallest fractions. As the available total R&D funds shrink and eventually get to 10% of the unconstrained ones, battery R&D becomes less predominant in the portfolio. Given a constrained total budget, the model tends to reallocate efforts more equally across technologies. Lower investments in other technologies make more likely lower realizations of the learning rates, and the model starts to hedge against these. This is true for solar, biofuels and, to a lesser extent, bioelectricity.

5 Conclusions

This paper copes with the problem of including uncertainty endogenously in the optimization problem of a complex Integrated Assessment Model used

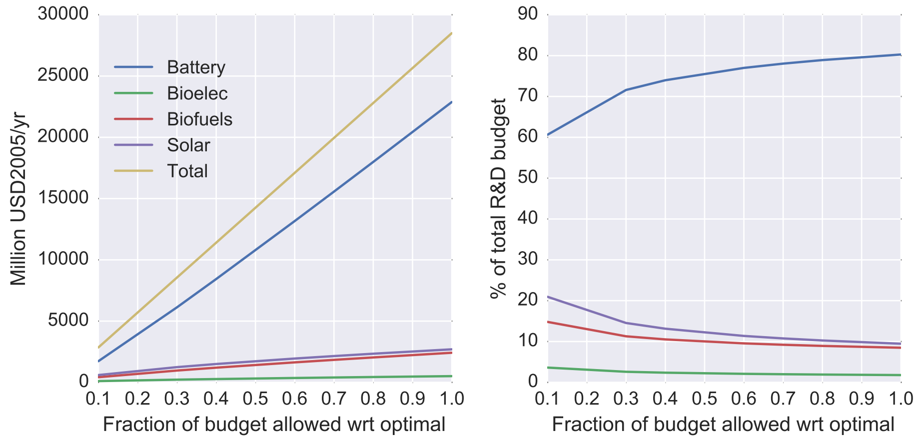

Figure 5: Sensitivity to different total budget limitations of the optimal R&D portfolio under uncertainty. Results are presented both in terms of yearly absolute investments levels (left-hand side) and corresponding shares (right-hand side).

for climate policy assessment. In our case, the decision variables of interest are energy R&D public investments to be allocated across 4 key low-carbon technologies in the short-term. Uncertainty affects the learning rate that will make these investments more or less effective in decreasing 2030 costs of these technologies. The problem is framed as a two-stage stochastic program: in the first stage (before 2030), a stock of knowledge is built through yearly R&D investments, decided for each technology; in the second stage (after 2030), the actual effectiveness of these R&D investments is revealed, and so are the actual costs of the four energy technologies, ultimately depending on the realization of the learning rates and on the accumulated stock of knowledge. The utility resulting from these possible technological futures is evaluated through the WITCH model. To make the problem computationally tractable, we apply the Approximate Dynamic Programming paradigm by replacing actual runs of the WITCH model with a surrogate function interpolated from 10,000 instances of the model, evaluated at different points in the energy technologies cost space.

The scope of this analysis is limited to Europe, as the estimation of learning rate distributions is based on surveys that gather the information of experts on future technological costs in Europe, conditionally to different European R&D scenarios. Nonetheless, WITCH solves the economy and determines the operation of the energy sector at a global scale, so that the European value function in the stochastic program accounts for world-wide strategic interactions and climate policy commitments. We decided to portray a realistic future scenario, where countries adhere to a stringent interpretation of their Copenhagen pledges, given the recent progress in climate negotiations obtained in Paris.

In this framework, we find the optimal R&D portfolio to be dominated by the battery sector for personal electric vehicles. Batteries seem to have a great potential in supporting the required decarbonization, and the low learning rates that need to be hedged against. Compared to the past, a significant upscaling of investments is suggested: 10-fold for the total budget and 100-fold for batteries. The share of batteries is robust to different assumptions of risk-aversion, R&D budget limitation and crowding-out effects. Had we not performed a full stochastic analysis, we would have severely underestimated the required scaling up of batteries investments.

Several further improvements could follow this study. As a first application of the methodology introduced in this paper, the number of technologies is limited to four, and only one climate policy is considered. In the future, these dimensions could be augmented. Investment decisions could be extended to all WITCH regions in the form of a Nash game. Other models beyond WITCH could be involved in a similar exercise to provide a further layer of robustness check to the analysis. Nonetheless, the main merit of this paper lies in having proved that the suggested methodology can be successfully applied to a fairly complex IAM for informing robust optimal clean energy innovation policies, and could similarly be applied in many other contexts.

6 End notes

6.1 Acknowledgements

The research leading to these results has received funding from the European Union's Seventh Framework Programme [FP7/2007-2013] under grant agreement n° 30832 (ADVANCE) and ERC StG 336703 — project RISICO. The authors also thank participants in the INFORMS 2014 conference.

References

Laura D. Anadón, Valentina Bosetti, Matthew Bunn, Michela Catenacci, and Audrey Lee. Expert judgments about RD&d and the future of nuclear energy. *Environmental Science & Technology*, 46(21):11497–11504, nov 2012. doi: 10.1021/es300612c. URL http://dx.doi.org/10.1021/es300612c.

Erin Baker and Yiming Peng. The value of better information on technology r&d programs in response to climate change. *Environmental Modeling & Assessment*, 17(1-2):107–121, jul 2011. doi: 10.1007/s10666-011-9278-y. URL https://doi.org/10.1007%2Fs10666-011-9278-y.

Erin Baker and Ekundayo Shittu. Uncertainty and endogenous technical change in climate policy models. *Energy Economics*, 30(6):2817–2828, November 2008. ISSN 0140-9883. doi: 10.1016/j.eneco.2007.10.001. URL http://www.sciencedirect.com/science/article/pii/S0140988307001211. 00040.

Erin Baker, Valentina Bosetti, Laura Diaz Anadon, Max Henrion, and Lara Aleluia Reis. Future costs of key low-carbon energy technologies: Harmonization and aggregation of energy technology expert elicitation data.

Energy Policy, 80:219–232, may 2015a. doi: 10.1016/j.enpol.2014.10.008. URL http://dx.doi.org/10.1016/j.enpol.2014.10.008.

Erin Baker, Olaitan Olaleye, and Lara Aleluia Reis. Decision frameworks and the investment in r&d. *Energy Policy*, 80:275–285, may 2015b. doi: 10.1016/j.enpol.2015.01.027. URL http://dx.doi.org/10.1016/j.enpol.2015.01.027.

Geoffrey J. Blanford. R&D investment strategy for climate change. *Energy Economics*, 31, Supplement 1:S27–S36, 2009. ISSN 0140-9883. doi: 10.1016/j.eneco.2008.03.010. URL http://www.sciencedirect.com/science/article/pii/S0140988308000534. 00032.

Valentina Bosetti and Massimo Tavoni. Uncertain R&D, backstop technology and GHGs stabilization. *Energy Economics*, 31, Supplement 1:S18–S26, 2009. ISSN 0140-9883. doi: 10.1016/j.eneco.2008.03.002. URL http://www.sciencedirect.com/science/article/pii/S0140988308000455. 00051.

Valentina Bosetti, Carlo Carraro, Marzio Galeotti, Emanuele Massetti, and Massimo Tavoni. WITCH - A World Induced Technical Change Hybrid Model. SSRN Scholarly Paper ID 948382, Social Science Research Network, Rochester, NY, November 2006. URL http://papers.ssrn.com/abstract=948382.

Valentina Bosetti, Massimo Tavoni, Enrica De Cian, and Alessandra Sgobbi. The 2008 WITCH model: New model features and baseline. Technical Report 85.2009, Nota di lavoro // Fondazione Eni Enrico Mattei: Sustainable development, 2009. URL http://www.econstor.eu/handle/10419/53280.

Valentina Bosetti, Michela Catenacci, Giulia Fiorese, and Elena Verdolini. Icarus expert elicitation reports. 2011. URL http://www.icarus-project.org/wp-content/uploads/2011/06/wp-icarus-web.pdf. 00002.

Valentina Bosetti, Giacomo Marangoni, Emanuele Borgonovo, Laura Diaz Anadon, Robert Barron, Haewon C. McJeon, Savvas Politis, and Paul Friley. Sensitivity to energy technology costs: A multi-model comparison analysis. *Energy Policy*, 80:244–263, May 2015. ISSN 0301-4215. doi: 10.1016/j.enpol.2014.12.012. URL http://www.sciencedirect.com/science/article/pii/S0301421514006776.

Yongyang Cai, Kenneth L. Judd, and Thomas S. Lontzek. The Social Cost of Stochastic and Irreversible Climate Change. Working Paper 18704, National Bureau of Economic Research, January 2013. URL http://www.nber.org/papers/w18704.

Gabriel Chan, Laura D. Anadon, Melissa Chan, and Audrey Lee. Expert elicitation of cost, performance, and RD&d budgets for coal power with CCS. *Energy Procedia*, 4:2685–2692, 2011. doi: 10.1016/j.egypro.2011.02.169. URL http://dx.doi.org/10.1016/j.egypro.2011.02.169.

Leon Clarke, Kejun Jiang, Keigo Akimoto, Mustafa Babiker, Geoffrey Blanford, Karen Fisher-Vanden, J-C Hourcade, Volker Krey, Elmar Kriegler, Andreas Löschel, et al. Assessing transformation pathways. *Contribution of Working Group III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, 2014.

Johannes Emmerling, Laurent Drouet, Lara Aleluia Reis, Michela Bevione, Loïc Berger, Valentina Bosetti, Samuel Carrara, De Cian, Enrica, de Maere d'Aertrycke, Gauthier, Thomas Longden, Maurizio Malpede, Giacomo Marangoni, Fabio Sferra, Massimo Tavoni, Jan Witajewski-Baltvilks, and Petr Havlík. The WITCH 2016 Model - Documentation and Implementation of the Shared Socioeconomic Pathways. SSRN Scholarly Paper ID 2800970, Social Science Research Network, Rochester, NY, June 2016. URL http://papers.ssrn.com/abstract=2800970. 00002.

Larry G. Epstein and Stanley E. Zin. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica: Journal of the Econometric Society*, pages 937–969, 1989. URL http://www.jstor.org/stable/1913778. 03182.

Andrii Gritsevskyi and Nebojša Nakićenovi. Modeling uncertainty of induced technological change. *Energy Policy*, 28(13):907–921, nov 2000. doi: 10.1016/s0301-4215(00)00082-3. URL http://dx.doi.org/10.1016/S0301-4215(00)00082-3.

IEA. RD&D Budget (Edition 2015). Technical report, Organisation for Economic Co-operation and Development, Paris, December 2015. URL http://www.oecd-ilibrary.org/content/data/c75cad3b-en. 00000.

Elmar Kriegler, Massimo Tavoni, Tino Aboumahboub, Gunnar Luderer, Katherine Calvin, Gauthier Demaere, Volker Krey, Keywan Riahi, Hilke Rösler, Michiel Schaeffer, and Detlef P. Van Vuuren. What Does The 2°C Target Imply For A Global Climate Agreement In 2020? The Limits Study On Durban Platform Scenarios. *Climate Change Economics*, 04(04):1340008, November 2013. ISSN 2010-0078. doi: 10.1142/S2010007813400083. URL http://www.worldscientific.com/doi/abs/10.1142/S2010007813400083. 00007.

Elmar Kriegler, John P. Weyant, Geoffrey J. Blanford, Volker Krey, Leon Clarke, Jae Edmonds, Allen Fawcett, Gunnar Luderer, Keywan Riahi, Richard Richels, Steven K. Rose, Massimo Tavoni, and Detlef P. van Vuuren. The role of technology for achieving climate policy objectives: overview of the EMF 27 study on global technology and climate policy strategies. *Climatic Change*, 123(3-4):353–367, jan 2014. doi: 10.1007/s10584-013-0953-7. URL http://dx.doi.org/10.1007/s10584-013-0953-7.

Derek Lemoine and Christian Traeger. Watch your step: Optimal policy in a tipping climate. *American Economic Journal: Economic Policy*, 6(1):137–

166, feb 2014. doi: 10.1257/pol.6.1.137. URL https://doi.org/10.1257%2Fpol.6.1.137.

Gunnar Luderer, Valentina Bosetti, Michael Jakob, Marian Leimbach, Jan C. Steckel, Henri Waisman, and Ottmar Edenhofer. The economics of decarbonizing the energy system—results and insights from the RECIPE model intercomparison. *Climatic Change*, 114(1):9–37, jun 2011. doi: 10.1007/s10584-011-0105-x. URL http://dx.doi.org/10.1007/s10584-011-0105-x.

Alan Manne, Robert Mendelsohn, and Richard Richels. MERGE. *Energy Policy*, 23(1):17–34, jan 1995. doi: 10.1016/0301-4215(95)90763-w. URL http://dx.doi.org/10.1016/0301-4215(95)90763-W.

Giacomo Marangoni and Massimo Tavoni. The Clean Energy R&D Strategy for 2°C. *Climate Change Economics*, 05(01):1440003, February 2014. ISSN 2010-0078. doi: 10.1142/S201000781440003X. URL http://www.worldscientific.com/doi/abs/10.1142/S201000781440003X. 00000.

Millett Granger Morgan, Max Henrion, and Mitchell Small. *Uncertainty: a guide to dealing with uncertainty in quantitative risk and policy analysis*. Cambridge university press, 1992.

Gregory F. Nemet, Laura Diaz Anadon, and Elena Verdolini. Quantifying the effects of expert selection and elicitation design on experts' confidence in their judgments about future energy technologies. *Risk Analysis*, mar 2016. doi: 10.1111/risa.12604. URL http://dx.doi.org/10.1111/risa.12604.

David Popp. ENTICE: endogenous technological change in the DICE model of global warming. *Journal of Environmental Economics and Management*, 48(1):742–768, July 2004. ISSN 0095-0696. doi: 10.1016/j.jeem.2003.09.002. URL http://www.sciencedirect.com/science/article/pii/S0095069603001220. 00304.

David Popp and Richard G Newell. Where does energy r&d come from? examining crowding out from environmentally-friendly r&d. Technical report, National Bureau of Economic Research, 2009.

Mort Webster, Nidhi Santen, and Panos Parpas. An approximate dynamic programming framework for modeling global climate policy under decision-dependent uncertainty. *Computational Management Science*, 9(3):339–362, may 2012. doi: 10.1007/s10287-012-0147-1. URL https://doi.org/10.1007%2Fs10287-012-0147-1.

Tobias Wiesenthal, Guillaume Leduc, Karel Haegeman, and Hans-Günther Schwarz. Bottom-up estimation of industrial and public r&d investment by technology in support of policy-making: The case of selected low-carbon energy technologies. *Research Policy*, 41(1):116–131, feb 2012. doi: 10.1016/j.respol.2011.08.007. URL http://dx.doi.org/10.1016/j.respol.2011.08.007.

A Appendix

A.1 Solution existence and uniqueness for one technology

The problem of maximizing Eq. (5) has a unique solution if utility is strictly concave. For simplicity, we consider analytically only the case with one technology, i.e. $J = 1$. Let us consider $U_{T_1}(\boldsymbol{I})$. Its first and second derivatives are:

$$\frac{\partial U_{T_1}(I)}{\partial I} = \frac{\partial}{\partial I}\sum_{t\in T_1} L_t\beta_t/(1-\eta)((Q_{0t}-rI)/L_t)^{1-\eta}-1) \tag{7}$$

$$= \sum_{t\in T_1} -r\beta_t((Q_{0t}-rI)/L_t)^{-\eta} \tag{8}$$

$$\frac{\partial^2 U_{T_1}(I)}{\partial I^2} = \frac{\partial}{\partial I}\sum_{t\in T_1} -r\beta_t((Q_{0t}-rI)/L_t)^{-\eta} \tag{9}$$

$$= \sum_{t\in T_1} -\frac{r^2\eta\beta_t}{L_t}((Q_{0t}-rI)/L_t)^{-\eta-1} \tag{10}$$

All parameters have positive values, so that $U_{T_1}(I)$ is decreasing ($U'_{T_1} < 0$) with a strictly concave behavior ($U''_{T_1} < 0$).

Let us consider $U_{T_2}(\boldsymbol{I})$, assuming λ fixed.

$$\frac{\partial U_{T_2}(I)}{\partial I} = \frac{\partial V(C(I))}{\partial I} = \frac{\partial V(C)}{\partial C}\frac{\partial C(I)}{\partial I} \tag{11}$$

The cost C depends on investments in R&D I via one-factor learning and accumulation of R&D capital equations:

$$C(I) = C_f + C_0\left(\frac{K(I)}{K_0}\right)^{-\lambda} \tag{12}$$

$$K(I) = K_0(1-\delta_R)^{2030-2005} + \Delta_T I \sum_{t=2010,\dots,2030}(1-\delta_R)^{t-2010}. \tag{13}$$

Renaming positive constants to yield a form as compact as possible, Eq. (11) becomes:

$$\frac{\partial U_{T_2}(I)}{\partial I} = V'(C)\frac{\partial}{\partial I}[c_f + c_0(\Delta_a+\Delta_b I)^{-\lambda}] \tag{14}$$

$$= V'(C)[-\Delta_c(\Delta_d + I)^{-\lambda-1}] \tag{15}$$

$V(C)$ is decreasing ($V'(C) < 0$), so that $U_{T_2}(I)$ results increasing in I ($U'_{T_2} > 0$). Regarding the second derivative:

$$\frac{\partial^2 U_{T_2}(I)}{\partial I^2} = \frac{\partial}{\partial I}V'(C)[-\Delta_c(\Delta_d + I)^{-\lambda-1}] \tag{16}$$

$$= V''(C)[-\Delta_c(\Delta_d + I)^{-\lambda-1}] + V'(C)[\Delta_c(\lambda+1)(\Delta_d + I)^{-\lambda-2}] \tag{17}$$

This is the sum of two negative terms, as $V(C)$ is convex ($V''(C) > 0$), and decreasing with C ($V'(C) < 0$). U_{T_2} is thus strictly concave in I. and the maximization problem is well posed.

A.2 Risk aversion derivation

So far, we have made implicit assumptions about the risk aversion of the decision maker. We introduce an explicit parameterization of risk preferences following Epstein and Zin (1989). Let us consider first a recursive utility V_t at time t, defined as a constant elasticity of substitution (CES) production function of Q_t/l_t and of utility at time $t+1$:

$$V_t = \left((1-\gamma_t) \left(\frac{Q_t}{l_t} \right)^{1-\eta} + \gamma_t V_{t+1}^{1-\eta} \right)^{\frac{1}{1-\eta}} \tag{18}$$

If we raise V_t to the power $(1-\eta)$ and unfold the recursion starting from the first period, replacing for brevity years $\{2005, 2010, 2015, \dots\}$ with their integer indices $\{1, 2, 3, \dots\}$, we obtain:

$$\bar{W}_1 = V_1^{1-\eta} = (1-\gamma_1) \left(\frac{Q_1}{l_1} \right)^{1-\eta} + \gamma_1 (1-\gamma_2) \left(\frac{Q_2}{l_2} \right)^{1-\eta} \tag{19}$$

$$+ \gamma_1 \gamma_2 (1-\gamma_3) \left(\frac{Q_3}{l_3} \right)^{1-\eta} + \dots \tag{20}$$

$$= (1-\gamma_1) \left(\frac{Q_1}{l_1} \right)^{1-\eta} + \Gamma_1 (1-\gamma_2) \left(\frac{Q_2}{l_2} \right)^{1-\eta} + \Gamma_2 (1-\gamma_3) \left(\frac{Q_3}{l_3} \right)^{1-\eta} + \dots \tag{21}$$

$$= \theta_1 \left(\frac{Q_1}{l_1} \right)^{1-\eta} + \theta_2 \left(\frac{Q_2}{l_2} \right)^{1-\eta} + \theta_3 \left(\frac{Q_3}{l_3} \right)^{1-\eta} + \dots \tag{22}$$

which is related by an affine transformation to a scaled version $\widetilde{W}_1$ of WITCH utility W_1, obtained by dividing by the sum of population levels over the finite time horizon of the model $L := \sum_t l_t$:

$$\widetilde{W}_1 = \frac{1}{L} W_1 = \frac{\bar{W}_1 - \sum_t \theta_t}{1-\eta} \tag{23}$$

The need to introduce a scaled version of W_1 comes from the CES requirements on γ_t to be less than one. New coefficients are related to each other and to WITCH ones by:

$$\theta_t = \frac{\beta_t l_t}{L} \tag{24}$$

$$\gamma_t = \frac{1 - \sum_{t' \le t} \theta_{t'}}{1 - \sum_{t'' \le t-1} \theta_{t''}} \tag{25}$$

$$\Gamma_t = \prod_{t' \le t} \gamma_{t'} = 1 - \sum_{t' \le t} \theta_{t'} \tag{26}$$

with the last equation due to the telescoping nature of the product of γ_t. In particular, the following holds. For $\eta = 1.5$, maximizing W_1 (i.e. solving WITCH) is equivalent to maximizing $\widetilde{W}_1$, or minimizing $\bar{W}_1$, or maximizing $V_1 = \bar{W}_1^{1/(1-\eta)}$. The advantage of thinking in terms of V_t is that a simple transform can be applied to future utility in Eq. (18) to make the Epstein-Zin risk preference parameter α explicit:

$$V_t = \left((1-\gamma_t)\left(\frac{Q_t}{l_t}\right)^{1-\eta} + \gamma_t \mathbb{E}\left[V_{t+1}^{1-\alpha}\right]^{\frac{1-\eta}{1-\alpha}}\right)^{\frac{1}{1-\eta}} \tag{27}$$

$$= \left((1-\gamma_t)\left(\frac{Q_t}{l_t}\right)^{1-\eta} + \gamma_t \mathbb{E}\left[\left(V_{t+1}^{1-\eta}\right)^{\frac{1-\alpha}{1-\eta}}\right]^{\frac{1-\eta}{1-\alpha}}\right)^{\frac{1}{1-\eta}} \tag{28}$$

$$= \left((1-\gamma_t)\left(\frac{Q_t}{l_t}\right)^{1-\eta} + \gamma_t \mathbb{E}\left[\left(\bar{W}_{t+1}\right)^{\frac{1-\alpha}{1-\eta}}\right]^{\frac{1-\eta}{1-\alpha}}\right)^{\frac{1}{1-\eta}} \tag{29}$$

If we unfold the recursion above for our stochastic program, the expectation operator appears only after the first 5 periods, when utility starts to be affected by uncertainty:

$$V_1^{1-\eta} = \sum_{t\in\{1,\dots,5\}} \theta_t \left(\frac{Q_t}{l_t}\right)^{1-\eta} + \Gamma_5 \mathbb{E}\left[\left(\bar{W}_6\right)^{\frac{1-\alpha}{1-\eta}}\right]^{\frac{1-\eta}{1-\alpha}} \tag{30}$$

We want to link $\bar{W}_6$ with WITCH utility after-2030 W_6:

$$\frac{W_6}{L} = \frac{1}{1-\eta}\left[\left(\theta_6\left(\frac{Q_6}{l_6}\right)^{1-\eta} + \theta_7\left(\frac{Q_7}{l_7}\right)^{1-\eta} + \dots\right) - \sum_{t\geq 6}\theta_t\right] \tag{31}$$

$$\bar{W}_6 = (1-\gamma_6)\left(\frac{Q_6}{l_6}\right)^{1-\eta} + \gamma_6(1-\gamma_7)\left(\frac{Q_7}{l_7}\right)^{1-\eta} + \dots \tag{32}$$

$$= \frac{1}{\Gamma_5}\left(\theta_6\left(\frac{Q_6}{l_6}\right)^{1-\eta} + \theta_7\left(\frac{Q_7}{l_7}\right)^{1-\eta} + \dots\right) \tag{33}$$

$$= \frac{1}{\Gamma_5}\left((1-\eta)\frac{W_6}{L} + \sum_{t\geq 6}\theta_t\right) \tag{34}$$

Maximizing W_1 in the risk aversion formulation is thus equivalent to the problem of:

$$\min \sum_{t\in\{1,\dots,5\}} \theta_t \left(\frac{Q_t}{l_t}\right)^{1-\eta} + \Gamma_5 \mathbb{E}\left[\left(\frac{1}{\Gamma_5}\left((1-\eta)\frac{W_6}{L} + \sum_{t\geq 6}\theta_t\right)\right)^{\frac{1-\alpha}{1-\eta}}\right]^{\frac{1-\eta}{1-\alpha}} \tag{35}$$

A.3 Fitting learning rate distributions

Procedure 1

procedure $\text{FIT}(k_0, \delta, I, c_F, C_{i,s}, \text{EMPCDF}(C_{i,s}), \text{FITCDF}_x)$

$k_0 \leftarrow$ R&D capital in start year

$\delta \leftarrow$ Yearly depreciation rate of R&D capital

$I \leftarrow$ Yearly baseline R&D investment

$c_F \leftarrow$ Floor cost

$C_{i,s} \leftarrow$ Cost sample i under R&D scenario s

$\text{EMPCDF}(C_{i,s}) \leftarrow$ Empirical CDF of $C_{i,s}$ according to experts

$\text{FITCDF}_x \leftarrow$ CDF w/ parameters x for LbR rates, to be fit to EMPCDF

    for $s \leftarrow \{1, 1.5, 2\}$ **do** ▷ 3 R&D baseline investment multiplier scenarios

        $K_{2010,s} \leftarrow k_0$

        for $t \leftarrow \{2011, ..., 2030\}$ **do**

            $K_{t,s} \leftarrow K_{t-1,s}^{1-\delta} + sI$ ▷ Capital accumulation

        end for

        $R_s \leftarrow K_{2030,s}/K_{2010,s}$ ▷ Ratio over initial capital

        for $i \leftarrow$ index sample in empirical CDF from experts **do**

            $L_{i,s} \leftarrow -\dfrac{\log\left(\left(C_{i,s} - c_F\right)/c_0\right)}{\log R_s}$ ▷ Invert 1-factor learning curve to obtain sample LbR rates

        end for

        $\text{EMPCDF}(L_{i,s}) \leftarrow 1 - \text{EMPCDF}(C_{i,s})$ ▷ $\text{EMPCDF}(L_{i,s}) =$ Prob(LbR rate $\leq$ given LbR rate)= = Prob(Cost $\geq$ given Cost)= = 1 - Prob(Cost $\leq$ given Cost) = = 1 - $\text{EMPCDF}(C_{i,s})$

    end for

    $x \leftarrow \arg\min_x \sum_{i,s} \left(\text{FITCDF}_x(L_{i,s}) - \text{EMPCDF}(L_{i,s})\right)^2$

    return x

end procedure

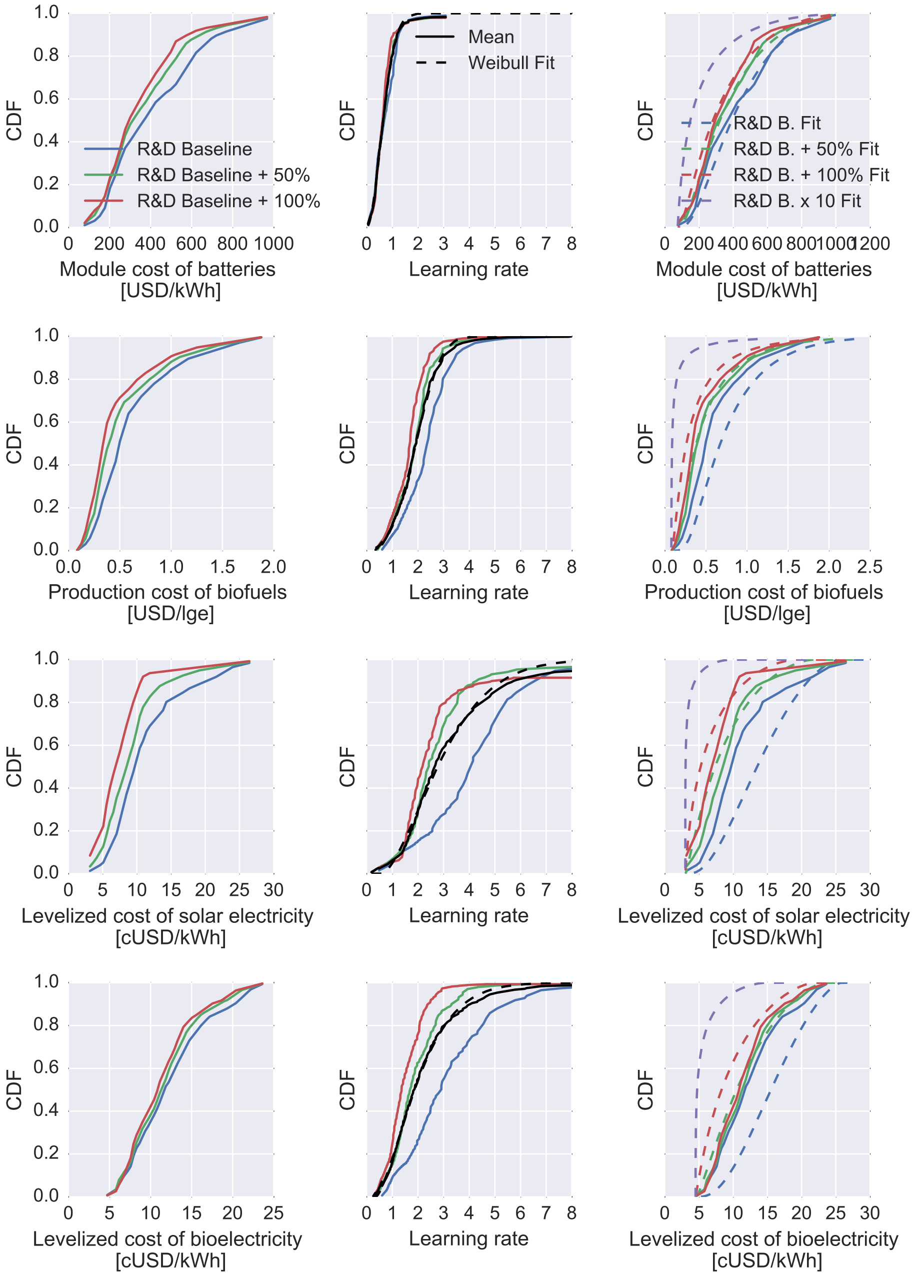

Figure 6: Left column: empirical CDFs of 2030 costs as elicited from experts, one technology for each row and one R&D budget for each color. Middle column: corresponding learning rate distributions, according to the one-factor learning curve model, plus a Weibull fit of the mean empirical CDF. Right column: empirical CDFs of costs along with their fitted versions, for the 3 original R&D budgets plus an extra one.

Table 2: Table summarizing statistics of fitted learning rate distributions for the four considered technologies.

	Solar	Biofuels	Battery	Bioelectricity
25th Quantile	1.83	1.40	0.41	1.19
Median	2.77	1.92	0.64	1.87
Mean	3.06	1.93	0.69	2.12
75th Quantile	3.98	2.45	0.92	2.78

A.4 Hermite interpolation of the value function

Let $\boldsymbol{i} = [i_1, ..., i_J]$ be the index of a single cost/welfare sample. After running the model, we collect all welfare samples $W^{(\boldsymbol{i})}$:

$$W^{(\boldsymbol{i})} := V(\boldsymbol{C}^{\boldsymbol{i}}) = \sum_{t\in\{2030,2035,...\}} L_t\beta_t/(1-\eta)((Q_t(\boldsymbol{C}^{(\boldsymbol{i})})/L_t)^{1-\eta} - 1) \tag{36}$$

along with first derivatives $\partial W/\partial C_j|_{\boldsymbol{C}^{(i)}}$, obtained from the marginals of the equations that input 2030 costs into the model. Next, we choose a polynomial of degree m defined in the normalized cost space $\mathcal{U} = [-1,1]^J$ to fit the welfare data points. Let $\boldsymbol{l} = [l_1, ..., l_J]$ be a vector of integers such that $\sum_j l_j = m$, and $\tilde{V}(\boldsymbol{X}) = a_{\boldsymbol{l}} \prod_j (X_j)^{l_j}, \boldsymbol{X} = [X_1, ..., X_J] \in \mathcal{U}$ our polynomial function. The interpolation problem is cast in the form:

$$\min_{a_{\boldsymbol{l}}} \sum_{\boldsymbol{i}} \left(\tilde{V}(\boldsymbol{X}^{(i)}) - W^{(\boldsymbol{i})}\right)^2 + \sum_j \gamma_j \left(\frac{\partial \tilde{V}(\boldsymbol{X})}{\partial X_j}\Big|_{\boldsymbol{X}^{(i)}} - \frac{\partial W}{\partial C_j}\frac{\partial C_j}{\partial X_j}\Big|_{\boldsymbol{X}^{(i)},\boldsymbol{C}^{(i)}} \right)^2 \tag{37}$$

$$\text{s.t.} \qquad X_j^{(\boldsymbol{i})} = -1 + 2\frac{C_j^{(\boldsymbol{i})} - C_{m,j}}{C_{M,j} - C_{m,j}} \qquad j \in \mathcal{J} \tag{38}$$

The function to be minimized is a weighted sum of the error both in absolute and derivative terms, the latter being done with respect to all technologies. An affine transformation maps $\boldsymbol{X}$ into $\boldsymbol{C}$, while the γ_j balances differences in units.

Both absolute levels and curvatures of the actual WITCH value function can be well replicated with a polynomial of degree 4 (Figure 7). If we look at bidimensional slices of the function, it turns out that the greatest impact is given by changing the cost of batteries, followed by biofuels. Changing the cost of solar and bioelectricity seems to yield minor effects. These results are consistent with the expected challenges of decarbonizing the non-electrical sector, and underline the crucial impact of technical change supporting this decarbonization.

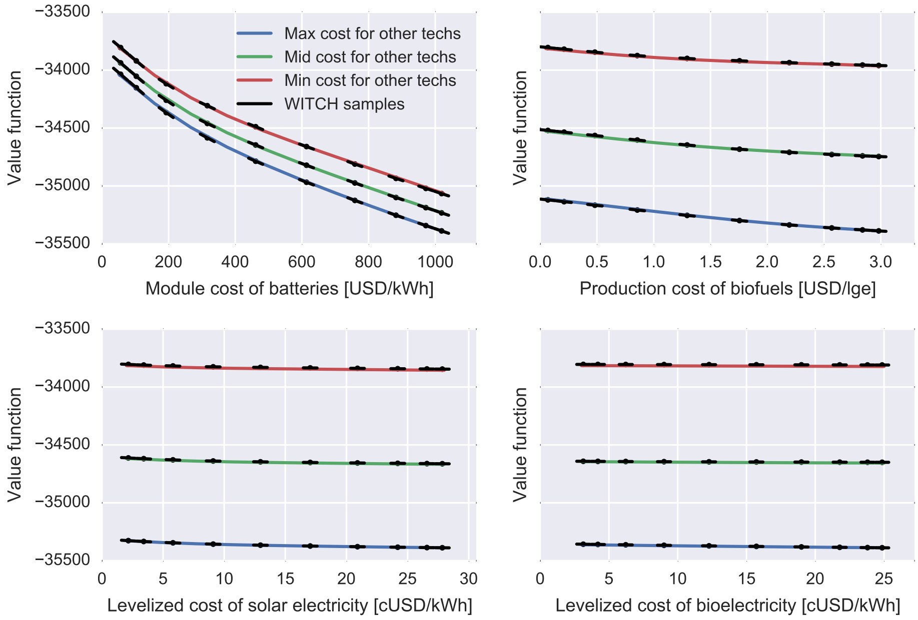

Figure 7: Bidimensional slices of WITCH 2030-onwards utility as a function of the 4 costs considered. A fitted polynomial of degree 4 is evaluated and plotted in color along each cost dimension, with the other costs either at maximum, minimum, or middle level. Actual values from WITCH are plotted in black. The segment around each sample brings information about actual partial derivatives.

NOTE DI LAVORO DELLA FONDAZIONE ENI ENRICO MATTEI

Fondazione Eni Enrico Mattei Working Paper Series

Our Note di Lavoro are available on the Internet at the following addresses:
http://www.feem.it/getpage.aspx?id=73&sez=Publications&padre=20&tab=1
http://papers.ssrn.com/sol3/JELJOUR_Results.cfm?form_name=journalbrowse&journal_id=266659
http://ideas.repec.org/s/fem/femwpa.html
http://www.econis.eu/LNG=EN/FAM?PPN=505954494
http://ageconsearch.umn.edu/handle/35978
http://www.bepress.com/feem/
http://labs.jstor.org/sustainability/

NOTE DI LAVORO PUBLISHED IN 2017

SAS	1.2017	Anna Alberini, Milan Ščasný: The Benefits of Avoiding Cancer (or Dying from Cancer): Evidence from a Four-country Study
ET	2.2017	Cesare Dosi, Michele Moretto: Cost Uncertainty and Time Overruns in Public Procurement: a Scoring Auction for a Contract with Delay Penalties
SAS	3.2017	Gianni Guastella, Stefano Pareglio, Paolo Sckokai: A Spatial Econometric Analysis of Land Use Efficiency in Large and Small Municipalities
ESP	4.2017	Sara Brzuszkiewicz: The Social Contract in the MENA Region and the Energy Sector Reforms
ET	5.2017	Berno Buechel, Lydia Mechtenberg: The Swing Voter's Curse in Social Networks
ET	6.2017	Andrea Bastianin, Marzio Galeotti, Matteo Manera: Statistical and Economic Evaluation of Time Series. Models for Forecasting Arrivals at Call Centers
MITP	7.2017	Robert C. Pietzcker, Falko Ueckerdt, Samuel Carrara, Harmen Sytze de Boer, Jacques Després, Shinichiro Fujimori, Nils Johnson, Alban Kitous, Yvonne Scholz, Patrick Sullivan, Gunnar Luderer: System Integration of Wind and Solar Power in Integrated Assessment Models: a Cross-model Evaluation of New Approaches
MITP	8.2017	Samuel Carrara, Thomas Longden: Freight Futures: The Potential Impact of Road Freight on Climate Policy
ET	9.2017	Claudio Morana, Giacomo Sbrana: Temperature Anomalies, Radiative Forcing and ENSO
ESP	10.2017	Valeria Di Cosmo, Laura Malaguzzi Valeri: Wind, Storage, Interconnection and the Cost of Electricity Generation
EIA	11.2017	Elisa Delpiazzo, Ramiro Parrado, Gabriele Standardi: Extending the Public Sector in the ICES Model with an Explicit Government Institution
MITP	12.2017	Bai-Chen Xie, Jie Gao, Shuang Zhang, ZhongXiang Zhang: What Factors Affect the Competiveness of Power Generation Sector in China? An Analysis Based on Game Cross-efficiency
MITP	13.2017	Stergios Athanasoglou, Valentina Bosetti, Laurent Drouet: A Simple Framework for Climate-Change Policy under Model Uncertainty
MITP	14.2017	Loïc Berger and Johannes Emmerling: Welfare as Simple(x) Equity Equivalents
ET	15.2017	Christoph M. Rheinberger, Felix Schläpfer, Michael Lobsiger: A Novel Approach to Estimating the Demand Value of Road Safety
MITP	16.2017	Giacomo Marangoni, Gauthier De Maere, Valentina Bosetti: Optimal Clean Energy R&D Investments Under Uncertainty