Beyond DICE: What the New-Generation Integrated Assessment Models Can Teach Us About the Optimal Abatement Policy or The Importance of Being Prudent

Riccardo Rebonato – Scientific Director ERCII[1]

EDHEC Business School – EDHEC Risk Climate Initiative Instiiute

[1]Work with Dherminder Kainth, Lionel Melin, Dominic O'Kane

Outline of Talk

The Problem - i

- How rapidly should we ramp up our efforts to curb climate change?
- When Integrated Assessment models that broadly share the same modelling assumptions for **the economy and the physics** of the problem have been used to determine the optimal course of abatement actions, radically different answers have been obtained.
- Models of the DICE family have advocated a gradual ramp-up of abatement efforts and relatively low carbon taxes (social cost of carbon).
- Models aligned with the approach taken in the Stern Review have called for a much steeper abatement schedule, and much higher carbon taxes.

The Problem - ii

- At the heart of these differences seems to be the choice of the 'correct' social discount rate, and of the utility discount rate in particular.
- Since the latter is intimately linked to ethical choices, there seems to be little hope to make substantive normative progress in this direction.
- Is there a way out of the impasse?

The Problem - iii

- The debate so far has been framed in the context of the maximization of *separable* utility functions, which unhelpfully entangle aversion to static risk and elasticity of intertemporal substitution.
- There have recently been several attempts to use recursive utility function to break this impasse, but the computational challenges have made the conclusions to date either tentative, or of limited scope, or very approximate – and, in all cases, computationally extremely burdensome.

The Problem - iv

- Making use of a novel, accurate, fast and high-resolution set of approximations, we can use realistically and extensively recursive utility functions to answer a number of fundamental economic questions of direct policy relevance.
- Fig 1 shows one key result: the stark difference between the abatement schedules
  1. produced by DICE;
  2. obtained using a Stern-like utility discount rate (0.0010) with the DICE model;
  3. recursive-utility setting calibrated as in (Bansal and Yaron 2004) *and a Nordhaus-like rate of utility discount (0.015)*.

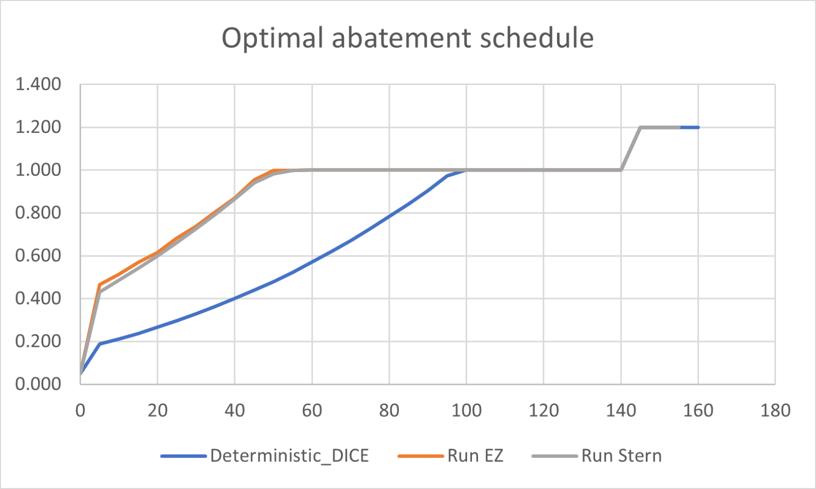

Figure 1: The optimal abatement schedule: Stern and Norhdaus parameters with recursive utility functions.

Key Results - i

First conclusion

- For a very wide range of risk aversion and elasticity of intertemporal substitution parameters consistent with asset prices, *even if we accept the DICE-model 'high' values of the utility discount rate*, recursive utility functions produce abatement schedules much more similar to those obtained in the Stern Review than by the traditional DICE model.

Key Results - ii

- These results are obtained while still centring the parameter uncertainty around the DICE default values for the economy and the physics of the problem.
- The 'traction' in the results comes mainly from unshackling the elasticity of intertemporal substitution from the coefficient of risk aversion.
- The analysis of several simultaneous sources of uncertainty in a recursive-utility framework has so far been deemed beyond computational reach, because of the exponential growth in computational cost of brute-force approaches.
- Our accurate approximation allows us to circumvent this barrier.

Key Results - iii

Second Conclusion

- We find that it is the elasticity of intertemporal substitution rather aversion to static risk that does the heavy lifting.
- However, we also find that the importance of the coefficient of relative risk aversion in determining the pace of the abatement schedule depends significantly on the chosen value for the elasticity of intertemporal substitution.
- The role of relative prudence when returns are stochastic is key.

Key Results - iv

Third Conclusion

- Nordhaus claims that uncertainty in economic growth is key driver of the abatement schedule.
- Ackerman and Stanton (2010) claim that, adding uncertainty to this quantity in isolation has a limited effect, but that it becomes more important when coupled with uncertainty in the damage exponent.
- We find that introducing uncertainty in the rate of growth in the economy often *reduces* the steepness of the optimal abatement schedule.
- This is important, because the slow Nordhaus optimal abatement schedules, and the low carbon taxes they suggest, are explained in terms of 'waiting until we are richer and smarter'.
- But in Nordhaus case we are richer 'for sure'.
- To understand results with uncertain growth **the role of relative prudence** (defined below) **is key.**

Key Results - v

Fourth Conclusion

- We document that negative emissions must play a key role in efficient abatement.
- Public attention and policy has so far been focussed almost exclusively on abatement.
- **This is very inefficient.**
- We should target net negative, not net zero.

Relevance of the Results - i

- Our results are not only economically interesting, but also have important policy implications.
- Cost-benefit analyses such as the one in the DICE model have often been criticized on several grounds:
  1. sometimes for not taking into account of important aspects of welfare to which it is difficult to give a consumption value;
  2. sometimes for focussing too much on the welfare of a representative (and homogeneous) agent, and therefore neglecting issues of inequality and distributive equity;
  3. and at times because the choice of the appropriate discount factor is deemed to be 'impossibly difficult'. [2]
- As a consequence, as important and influential a body as the IPCC has chosen not to make use of cost-benefit analysis ('*Thus standard cost–benefit analyses become difficult to justify (IPCC, 2014a; Dietz et al., 2016) and are not used as an assessment tool in this report*').

[2] '*It is extremely difficult if not impossible to meaningfully estimate discount rates for future costs and benefits*' - IPCC Chapter 1 page 76)

Relevance of the Results - ii

- We show that a cost-benefit analysis carried out with properly calibrated recursive utility functions, **even when used with a high utility discount rate**, points to abatement schedules and carbon taxes which are already very ‘aggressive’ (see again Fig 1).
- Yes, a lower utility discount rate, a richer modelling of welfare that included difficult-to-quantify elements such as loss of biodiversity, or a treatment that accounted for dislike for intragenerational inequality would certainly give rise to even steeper abatement schedules, and even higher optimal carbons taxes.
- However, that the optimal solutions we find are already at the outer boundaries of political implementability, and the policy value of debating whether even more aggressive courses of action would in theory be preferable is therefore moot.

Relevance of the Results - iii

- Subsidies and taxes have so far been targeted to
  1. increase the reach of renewables (very successful);
  2. reducing the use of fossil fuel via carbon taxes (very limited success);
- **Substantial subsidies should be considered for negative emission technologies as well.**

The DICE Model in Brief - i

- The DICE model belongs to the class of neoclassical models of growth theory pioneered by Solow (1970).
- In these models, agents invest in
  1. capital,
  2. education,
  3. technology, ...
- By doing so
  1. they *reduce* consumption today
  2. in the hope to *increase* future consumption.
- The increase in future (expected) consumption must be big enough to compensate for today's sacrifice.

The DICE Model in Brief - ii

- The DICE model extends this approach by introducing the natural capital of the climate system.
- Concentrations of CO2 in the atmosphere can be seen as *negative natural capital*.
- Emission reduction plays the role of an *investment* that increases natural capital – or decreases the deterioration of natural capital.
- If we invest in emission reduction, we consume less today.
- However, we reduce future damaging climate change, and increase consumption in the future.

The DICE Model in Brief - iii

- In the DICE model gross product, $y_{gross}(t)$, is a Cobb-Douglas function of capital, labour and technology.
- The Cobb-Douglas function is of the form

$$y_{gross} = TPF(t) \left(\frac{L(t)}{1,000} \right)^{1-\gamma} K^{\gamma} \tag{1}$$

  where K is the investment capital, L is the labour force, TPF is the total production factor, and 1000 is just a scaling constant.
- In the model, every individual works.

The DICE Model in Brief - iv

- As for the inputs to the production function in 'classic DICE'
  1. the population function, $L(t)$, is deterministic and exogenous;
  2. capital K depreciates at a fixed known rate, and new capital is equal to savings – which is one of the control variables.
  3. the TPF is exogenous and deterministic.
- The (discrete-time) deterministic equation governing the TPF is

$$TPF(t+1) = TPF(t)\left[1 + g_{TPF}(t)\right] \tag{2}$$

- The growth rate of the TPF decreases over time, according to

$$g_{TPF}(t) = \frac{g_{TPF}(t-1)}{1+\delta_{TPF}} \tag{3}$$

The DICE Model in Brief - v

- World economic output after the damage inflicted by climate change, $y_{net}(t)$, is given by

$$y_{net}(t) = y_{gross}\left[1 - damfrac(T)\right] \tag{4}$$

  with $damfrac(T) = a_1 T + a_2 T^2$.

- Investment $inv(t)$ is given by

$$inv(t) = y_{net}(t) \times savrat(t) \tag{5}$$

  and $savrat(t)$ is the savings rate at time t.

- Capital at time t is given by

$$K(t) = K(t-1)\underbrace{(1-dk)^{\Delta t}}_{depreciation} + \Delta t \times inv(t) \tag{6}$$

The DICE Model in Brief- vi

- In the equations above the saving rates at different times are one set of control variables.
- The other control variables are the abatement fractions, ie, the fraction of production devoted to abatement efforts at different times.
- So, the world production after damage and after abatement, $y(t)$ is given by[3]

$$y(t) = y_{gross}(t)\,[1 - damfrac(T_t) - abatecost(t)] \qquad (7)$$

- After subtracting investment, this is the quantity that gives consumption, $c(t)$ (the input to the per-period utility):

$$c(t) = y(t) - inv(t) \qquad (8)$$

[3] The expression is so neatly factorizable because the abatement cost is assumed proportional to y_{gross}.

From DICE to DICE ++ - i

- Some authors (notably Pyndick) have argued that IAMs are too crude and should be ditched.
- We disagree, but we agree that they should be radically revised.
- Huge uncertainties remain, but, as with every good model, we can understand what the ‘pressure points’ are.
- With Pyndick, we agree that models are dangerous if used in an oracular fashion.
- But we think that they are very useful if used critically and ‘suspiciously’.

From DICE to DICE ++ - ii

- We have enriched DICE by
  1. using recursive utility functions;
  2. allowing for *possibility* of tipping points;
  3. allowing for uncertainty in economic growth, in the damage function and in the location of the tipping points;
  4. allowing for Bayesian updating of information;
  5. ensuring that we recover stylized investment facts (equities and fixed income);
  6. allowing for (costly) negative emissions as well.
- We discuss some of these changes in what follows.

Recursive Utility Functions - i

- In the original DICE (and Stern) approach the optimal values of the control functions are obtained by maximizing the sum of discounted welfare for each period. This assumes time-separable utility functions. These are not 'innocuous'.
- If one posits for today's welfare U_0

$$U_0 = \sum_{i=0}^{N} u(c_i) \exp(-\delta t_i) \tag{9}$$

with c_i consumption at time t_i, δ the impatience discount rate for utility and

$$u(c) = \frac{c^{1-\eta} - 1}{1 - \eta} \tag{10}$$

then the parameter η plays a double role:

- It represents both the inverse of the intertemporal elasticity of substitution, and the coefficient of relative risk aversion.

Recursive Utility Functions - ii

- Under conditions of economic growth we expected future generations to be richer, and therefore their marginal utility of consumption lower.
- This means that a future loss in consumption should give rise to a smaller loss in utility than the same consumption loss today.
- The more so, the higher the parameter η: the more we dislike uneven consumption, we more we want to push the abatement burden on our richer (grand)children.

Recursive Utility Functions - iii

- From this one expects that a high η should suggest modest abatement efforts today: high abatement efforts today are seen from this perspective as a tax on the (current) poor to benefit the (future) rich.
- However, if we look at the parameter η from the perspective of relative risk aversion, under conditions of uncertainty a high η should suggest that substantial resources should be spent today because of our aversion to static risk.

Recursive Utility Functions - iv

- As Ackerman (2013) eloquently puts it
  '*these two interpretations of* η *might appear to parallel each other (...) [but] they are actually on a collision course. (...) [With separable utility functions] [t]here is no natural way to model a combination of strong risk aversion and strong future orientation*'.
- Not surprisingly, authors, such as Cline (1992) or Stern (2007), who have chosen a very low value for η, and authors, such as Nordhaus (2007), who have chosen a higher value, have arrived at radically different policy recommendations.

Recursive Utility Functions - v

- One of the more appealing solutions out of this impasse is offered by recursive utility functions, of which those in the Epstein-Zin family are the best known.
- With this formulation, the welfare at time t is now given by

$$U_t = \left[(1-\beta)c_t^\rho + \beta\left(\mu_t[U_{t+1}]\right)^\rho\right]^{\frac{1}{\rho}} \quad (11)$$

- In this expression β is the utility discount factor, linked to the pre time-preference parameter, δ, by $\beta = \exp(-\delta)$, $\mu_t[U_{t+1}]$ is the certainty equivalent of future welfare, given by

$$\mu_t[U_{t+1}] = \left(E_t[U_{t+1}^\alpha]\right)^{\frac{1}{\alpha}} \quad (12)$$

- **This where the risk-aversion bit is localized.**

Recursive Utility Functions - vi

- This formulation allows to disentangle aversion to static risk (the coefficient of relative risk aversion is now equal to $1 - \alpha$), and the elasticity of intertemporal substitution (which is now given by $\psi = \frac{1}{1-\rho}$).
- In asset pricing, as Bansal and Yaron (2004) show, when both coefficients are chosen to be greater than 1 (and consumptions and dividend growth are suitably modelled), one can simultaneously account both for equity returns and the level of interest rates.[4]
- **Discussion of $EIS > 1$ and why equity and bond prices are important – reduced-form models.**
- These desirable features of recursive utility functions however come at a high computational cost.

[4] The Bansal and Yaron model 'can justify the equity premium, the risk-free rate, and the volatility of the market return, risk-free rate, and the price-dividend ratio. As in the data, dividend yields predict returns and the volatility of returns is time-varying.'

Recursive Utility Functions - vii

- In an IAM such as DICE, consumption is an *endogenous*, **path-dependent** quantity.
- Even if the stochastic drivers can be mapped onto a recombining tree, path dependence forces consumption to be mapped onto a non-recombining ('bushy') tree.
- In a bushy tree the number of nodes grows as $O((2^{n_{fac}})^{n_{steps}})$ where n_{fac} is the number of stochastic drivers, and n_{steps} is the number of steps.
- This is why in the recursive-utility-function literature one often finds at most *one* stochastic driver, and time steps as long as 75 years!

Recursive Utility Functions - viii

- We have devised an approximate but very accurate scheme to calculate the expectations needed by the EZ formulation.
- Thanks to this approximation the Bellman (backward induction) equation can be used to calculate the recursive Epstein-Zin welfare.
- We have done this in the presence of as many as four stochastic drivers.[5]

[5] There is no *a priori* reason to limit the number of stochastic drivers to four. We have not found the need to go beyond this value.

The Method - i

- The welfare today can be calculated by backward induction, ie, by application of the Bellman's equation.
- After making the change of variables $W_t = [U_t]^{1-1/\psi}/(1-\beta)$, the function to maximize (for $\psi > 1$) is given by

$$W_t = u(c_t, L_t) + \beta \left[E_t \left\{ W_{t+1}^{\frac{1-\gamma}{1-1/\psi}} \right\} \right]^{\frac{1-1/\psi}{1-\gamma}} \tag{13}$$

and by

$$W_t = u(c_t, L_t) - \beta \left[E_t \left\{ -W_{t+1}^{\frac{1-\gamma}{1-1/\psi}} \right\} \right]^{\frac{1-1/\psi}{1-\gamma}} \tag{14}$$

for $0 < \psi < 1$.

The Method - ii

- Equations 13 and 14 are a function of the exogenously assigned control variables, ie, the abatement fraction, μ_t, the savings rate, sav_t, and the negative emissions, $negem_t$.
- These quantities are then varied by a numerical search in such a way as to maximize Equations 13 and 14, subject to the physical constraints and the specification of the economy afforded by the DICE model.
- In keeping with the set-up of the DICE model we have used $n_{steps} = 100$, a final horizon of 500 years (and hence $\Delta t = 5$).

Modelling Uncertainty – Economic Growth - i

- Following Bansal and Yaron (2004) and Jensen and Traeger (2014) we model consumption growth as made up of two components:
  1. a (small) persistent component; plus
  2. a fluctuating volatility, which describes economic uncertainty.
- This description is appealing for our purposes because Bansal and Yaron (2004) show that this specification accounts well for annual consumption data.
- The time-varying volatility component is needed to account for the empirically observed fluctuating consumption volatility.

Modelling Uncertainty – Economic Growth - ii

- Consumption growth, g_t is modelled as follows.
- The small persistent component, x_t, is described by

$$x_{t+1} = \rho x_t + \phi e_{t+1} \tag{15}$$

- Consumption growth is then given by

$$g_{t+1} = \mu + x_t + \sigma \eta_{t+1} \tag{16}$$

  with e_{t+1} and η_{t+1} independent.
- If $e_{t+1} = \eta_{t+1}$, the process for consumption growth is an ARMA(1,1) process, and if $\phi = \rho$, then the consumption growth follows an AR(1) process.

Modelling Uncertainty – Economic Growth - iii

- In the DICE model, what is directly modelled is the growth of the Total Factor of Production (TFP), TFP_t.
- To model the growth in TFP in a manner consistent with the Bansal and Yaron specification, we follow Jensen and Traeger (2014):

$$TFP_t = TFP_0 \exp[g_t^{TFP} t] \tag{17}$$

with

$$g_t^{TFP} = g_0^{TFP} \exp[-\delta^{TFP} t] + z_t \tag{18}$$

and

$$z_t = x_t + w_t \tag{19}$$

$$w_t = \zeta w_{t-1} + \epsilon_t \tag{20}$$

Modelling Uncertainty – Economic Growth - iv

- Both x_t and ϵ_t are assumed to be normally distributed according to $N(\mu_x, \sigma_x^2)$ and $N(\mu_\epsilon, \sigma_\epsilon^2)$.
- As in Jensen and Traeger 2014, we choose $\zeta = 0.5$, and $\sigma_x = \sigma_\epsilon = 1.9\%$, which gives the same percentage volatility for z_t.
- The means (μ_x and μ_ϵ) are then chosen to match the expected growth path with the DICE deterministic path.

Results - i

We can now move to the results.

- We present the case of $EIS = 1.45$, $\gamma = 5$ (as per Bansal and Yaron). We have looked at many more variations.
- We allow for uncertainty in
  1. the damage exponent, a_3,
  2. the Total Factor of Production, and
  3. the location of the tipping points.
- We also consider the case of no volatility, to disentangle the effect of static risk aversion from the impact of EIS.

Results - ii

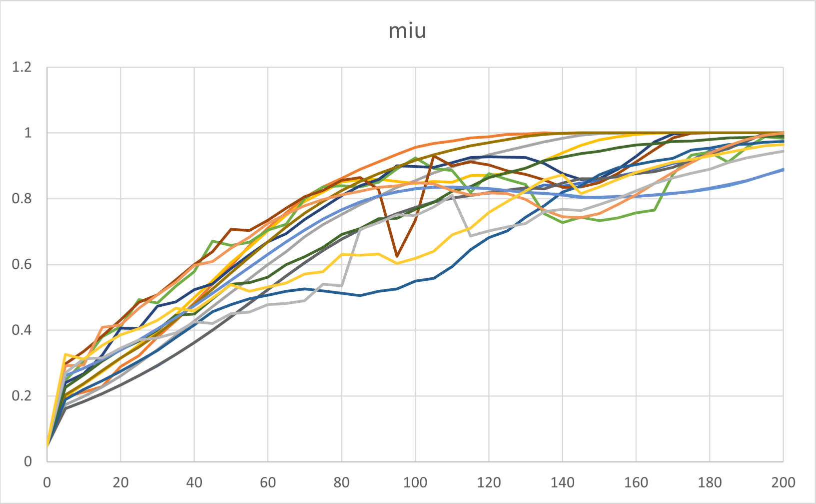

Figure 2: The optimal abatement schedule for different preferences and different stochastic drivers in the presence of negative emissions.

Results - iii

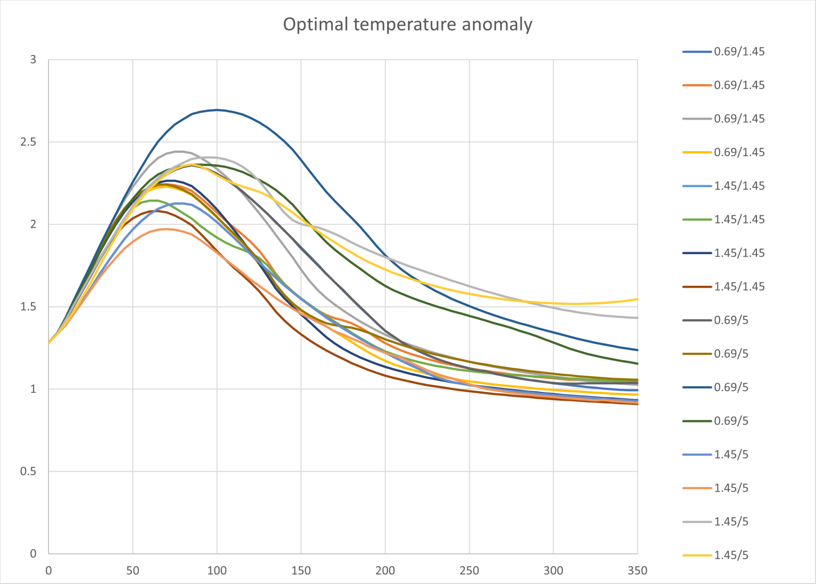

Figure 3: The optimal temperature anomaly for different preferences and different stochastic drivers

Results - iv

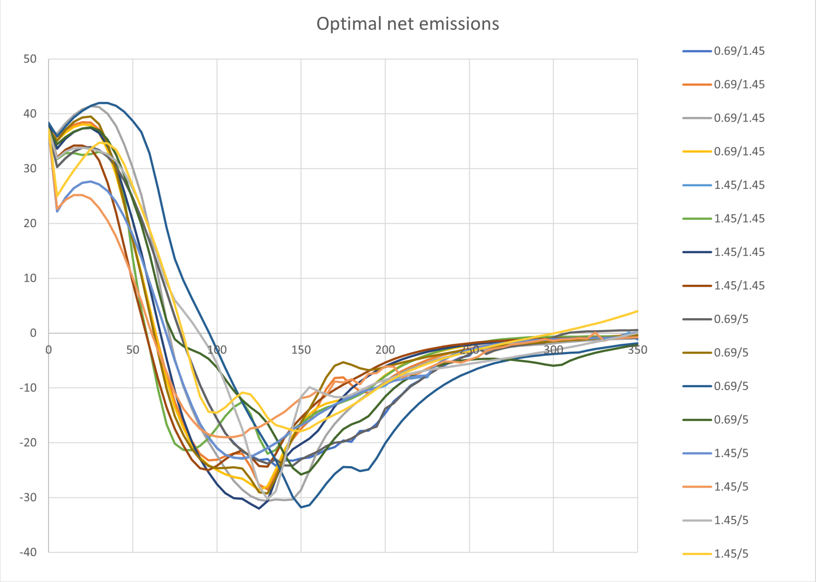

Figure 4: The optimal net emissions for different preferences and different stochastic drivers

Results - v

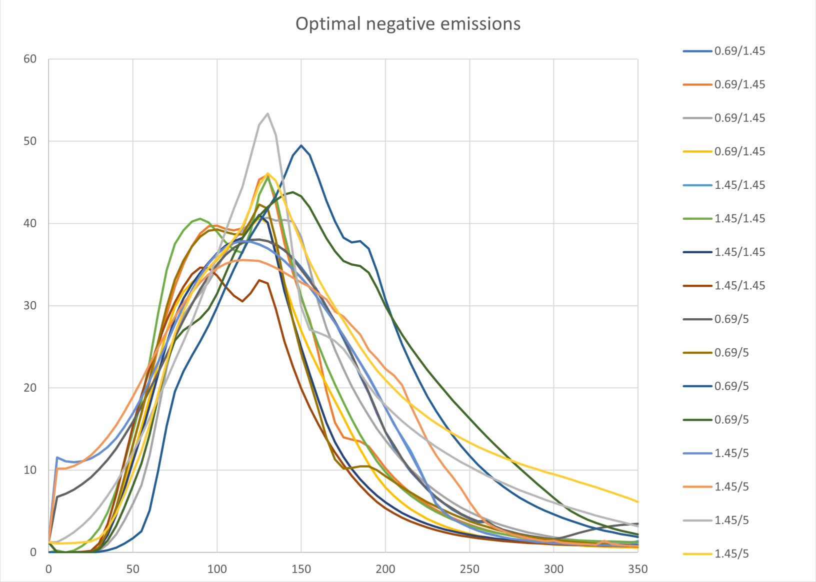

Figure 5: The optimal negative emissions for different preferences and different stochastic drivers

Results - vi

Discussion of Results

- *Ceteris paribus*
  - uncertainty in the damage exponent always makes the abatement schedule more aggressive, while
  - uncertainty in TFP almost always reduces the rate of abatement.
- This may appear surprising, because in both cases, the higher the dislike for uneven consumption (the lower the EIS), the greater the desire to invest to 'ensure the "minimum standard of living" [...] in the face of uncertainty';

Results - vii

Discussion of Results

- The explanation lies in the coefficient of relative prudence.
- Relative prudence in general determines whether the investor should increase or decrease savings (and investments).
- Relative prudence, $p(x)$, is defined for time-separable utility functions as

$$p(x) = -x\frac{u'''(x)}{u''(x)} \tag{21}$$

- In the absence of climate change, the coefficient of relative prudence dictates that uncertainty about growth should only increase investment if the precautionary motive ('ensuring the "minimum standard of living" ') outweighs the aversion for the uncertainty of the associated returns.

Results - viii

- With recursive utility functions prudence becomes a more complex function of EIS and CRRA, but, even in the simpler time-separable utility, one cannot say that increasing uncertainty will always increase savings/investment.
- With recursive utility functions one can obtain

$$p(x) = CRRA[1 + EIS] = CRRA\left[1 + \frac{1}{AIS}\right] \tag{22}$$

- Therefore, *ceteris paribus*, aversion to uneven consumption decreases prudence, and aversion to static risk increases it.

Results - ix

- Consider the simpler case of separable utility functions, and **no climate damages**.
- In this setting, the response to uncertainty in returns is by investing more only if relative prudence, defined in Eq 21, is positive.
- For the isoelastic utility function in the original DICE model $EIS = 1/CRRA$, and $p = 1 + CRRA = 2.45 > 0$.
- So, in this setting the positive of relative prudence points to an increase of investment in production capital under uncertainty.

Results - x

- **When climate damages are taken into account**, for investment to increase in the presence of uncertainty, relative prudence must not only be positive, but greater than the sensitivity of damages to production shocks:

$$p(c) > 2Dam_1(y) - \frac{Dam_1(y)}{AIS(c)} Dam_2(y) \tag{23}$$

where

- $AIS = 1/EIS$,
- $Dam_1 = \frac{d'(y)}{d(y)} y$ (elasticity of marginal damages w/r/t production)
- $Dam_2 = \frac{d''(y)}{d'(y)} y$, and
- $d(y) = -\frac{\partial y_t}{\partial M_t}$ (marginal damage) w/r/t M_t, concentration.

Results - xi

- If damages did not depend on growth, ie, if the term Dam_1 were equal to zero, with positive prudence the asymmetric response to same-size positive and negative damage shocks would call for greater investment under uncertainty.
- If the term $Dam_1(y)$ is greater than zero, the optimal abatement response to uncertainty is positive only if the prudence effect dominates (twice) the damage sensitivity to production.

Results - xii

- The key point is that whether to invest more or less in the presence of return uncertainty and climate damages does not depend on CRRA alone, but on a combination of CRRA and EIS.
- The problem with time-separable utility functions is that EIS and CRRA are 'entangled', and resulting coefficient of relative prudence is likely to be vitiated by this forced link between the two quantities.
- This is another reason why using recursive utility functions, where EIS and CRRA can be assigned independently, is so important.
- **With the Bansal-Yaron-consistent preference parameters, the net effect is a reduction in abatement investment when output is uncertain.**

Conclusions and Future Developments - i

- One of the key messages from using recursive utility functions calibrated to recovering the prices of equities and bonds is that the optimal net reduction in emissions is much steeper than recommended by the standard DICE model.
- This is true even if *the utility discount rate and the physical parameters are kept (or centred) at their DICE levels*.
- Many additional features (tipping points, different damage functions, existence of background unhedgeable risk) make the abatement schedule even steeper.
- **However, we are already hitting (hard) against constraints of political implementability**.

Conclusions and Future Developments - ii

- Not all uncertainty is the same!
- Uncertainty in the damage function *always increases* the optimal abatement schedule (and the SCC).
- Uncertainty in economy growth *often decreases* the optimal abatement schedule (and the SCC): the more returns are uncertain, the more we can become reluctant to 'gamble our investment' – which is *certain* foregone consumption.

Conclusions and Future Developments - iii

- Contrary to the standard DICE results, the social cost of carbon (the cost of one extra ton of CO2) is such that most available abatement initiatives are economically worth undertaking now.[6]
- This should encourage the re-engagement with IAMs of those policy bodies that have turned their backs on them.
- Carbon *removal* should start more slowly, but must play an important role.
- *Going 'cold turkey' on abatement is not only too socially painful, but also inefficient.*
- Apart from political commitment, a big problem is how to combine the abatement initiatives subject to *physical* (not only *cost*) constraints.

[6] The social cost of carbon, SCC, is defined as $SCC_j \equiv -\frac{\partial W/\partial e_j}{\partial W/\partial c_j}$. With our set-up we can study the *state-dependence* of the SCC.

Conclusions and Future Developments - iv

- Depending on the damage function, if we consider tipping points *possible*, the end-of-century Paris-Agreement target of temperature increase between 1.5 and 2 C can be justified as an optimal policy, not an exogenous aspiration to which optimality of means must then be applied.
- This is obtained without changing the central estimate of the damage exponent ($a_3 = 2$) in Dice (and Rudik (2020)).

Thank you!