Climate Change Risk Is Systematic

Bob Litterman
November 19, 2010

Based on “Optimal Climate Policy” joint work with Kent Daniel

Two reasons to talk about climate change at this conference:

1. Financial engineering provides tools and insights necessary to develop appropriate climate policy

   - **Risk management**
     - tipping points
     - fat tails
     - positive feedbacks
     - non-linear responses

   - **Bayesian analysis**
     - realistic characterization of uncertain distributions of outcomes

   - Use of **stochastic discount rates**
     - to present value future damages

   - Incorporating realistic **risk aversion**
     - calibration to financial market risk premia

2. Carbon emissions prices are not in equilibrium, and this disequilibrium creates an opportunity for investors

In the literature on climate policy the impact of risk being systematic is often missed

- Historical average annual real returns:
  - t-bills: 2.0%
  - long equities: 7.5% per year **above** t-bills
  - short equities: 7.5% per year **below** t-bills
- Risk premia depend on **covariances** not volatilities
- More fundamentally premia depend on whether payoffs are obtained in **good** times or **bad** times
- Shorting equities pays off in bad times just as investments in climate mitigation would
- Based on market prices (i.e. the large equity risk premium) people are willing to pay a lot to reduce systematic risk
- Thus, like shorting equities, the hurdle rate of return on climate mitigation should be **a significant negative rate**
- Many economists have **incorrectly** argued that because such investments are risky they require a high rate of return

This analysis builds directly on a number of previous works

1. Summers & Zeckhauser, ‘Policymaking for posterity’ (2008)
   - A simple two-period model; the impact of increased uncertainty on policy depends on parameters of technology and preferences
2. Weitzman, ‘GHG targets as insurance against catastrophic climate damages’ (2010)
   - The essence of the emissions externality is the risk of catastrophic damages
3. Ackerman, et al, ‘Limitations of integrated assessment models…’ (2009)
   - Standard CRRA utility used in climate models match neither market interest rates nor the equity risk premium
4. Epstein & Zin, ‘Substitution, risk aversion…’ (1989)
   - A utility function that separates the elasticity of intertemporal substitution from the degree of risk aversion

Climate change science IS uncertain: that uncertainty implies increased risk

- The earth has a reservoir of capacity to safely absorb greenhouse gas emissions

- The capacity of that reservoir is unknown

- Scientists can speculate about what will happen if the capacity is exceeded

  - There may be a tipping point
  - There may be positive feedbacks and nonlinearities
  - There may be a global collapse of ecosystems
  - There may be a very significant decrease in human utility

- Scientists cannot be highly confident about the level of damage to the earth’s very complex, and highly integrated ecosystems

- Economists should not be overly confident about the level of impact to human welfare

The key issue facing climate policy is where to price that systematic risk

- Science makes clear that there is an uncertain distribution of outcomes associated with the increased GHG emissions

- This distribution includes a tail with climate catastrophes

- Thus there is a systematic risk that needs to be priced

- Scientists don't have the tools to determine the appropriate price

- Financial economists do have those tools

Not pricing risk leads to disasters

- It is easy to point to many historical examples:
  - New Orleans neglect of it's levee system before Katrina
  - Ignoring clear warning signs while drilling in deep water in the gulf of Mexico on the Macondo well
  - Getting drunk while responsible for piloting the Exxon Valdiz
  - One of my personal favorite examples: the Johnstown flood of 1889
- Not pricing systematic risk leads to global catastrophes
  - One might point to the two world wars
  - The current financial crisis is another example
  - Not pricing carbon emissions is potentially the most significant example in history

Consider the Johnstown Flood of 1889
The Courts Ruled it: “An Act of God”

There was a reservoir with limited capacity

There was a tipping point

Positive feedback and nonlinearities
led to an “unimaginable” disaster

…2,200 Lives Were Lost

Of course the calamity was totally forecastable
Risk was not priced appropriately

- Many mistakes were made:
  - A drain for the reservoir was disabled
  - The dam was lowered to allow a road to run across the top
  - A screen to prevent loss of fish covered the causeway and became an obstruction

- The reservoir filled during 4 days of rain

- Warning was given too late for the population to react

- Even at its peak, the water level increased at a rate of only 1 inch per hour

Carbon emissions also create risk and this **risk should be priced**

- This risk is an **externality** that should be priced, not simply in the sense of insurance – being systematic it deserves an additional risk premium

- Yet today the current global effective price of emissions is significantly **negative**

  - Total global fossil fuel subsidies were $312 bn in 2009 according to the International Energy Agency
    - (Compare this to renewable energy subsidies of $57 bn)

  - Thus, the current subsidy of fossil fuel use equates to a **negative** carbon emissions tax on the order of $10 per ton of CO_2

- This reality is **a significant deviation from equilibrium**

There is no serious disagreement with the need to price this externality

- The appropriate **price of carbon emissions** equals the expected present value of the damages created – the so-called **'social cost of carbon'** or SCC
  - The appropriate current price clearly depends on future emissions
  - Future emissions depend on future prices
- Thus, the SCC is **the first point** in the context of a specified **dynamic price path** for emissions prices
- Of interest is not only the appropriate current price, but the **shape of the optimal price path**, in particular its **slope**
- The **slope** depends critically on the **degree of risk aversion**
- In particular, **higher risk aversion implies a decreasing slope**

Higher risk aversion impacts the slope of the optimal path of emissions prices

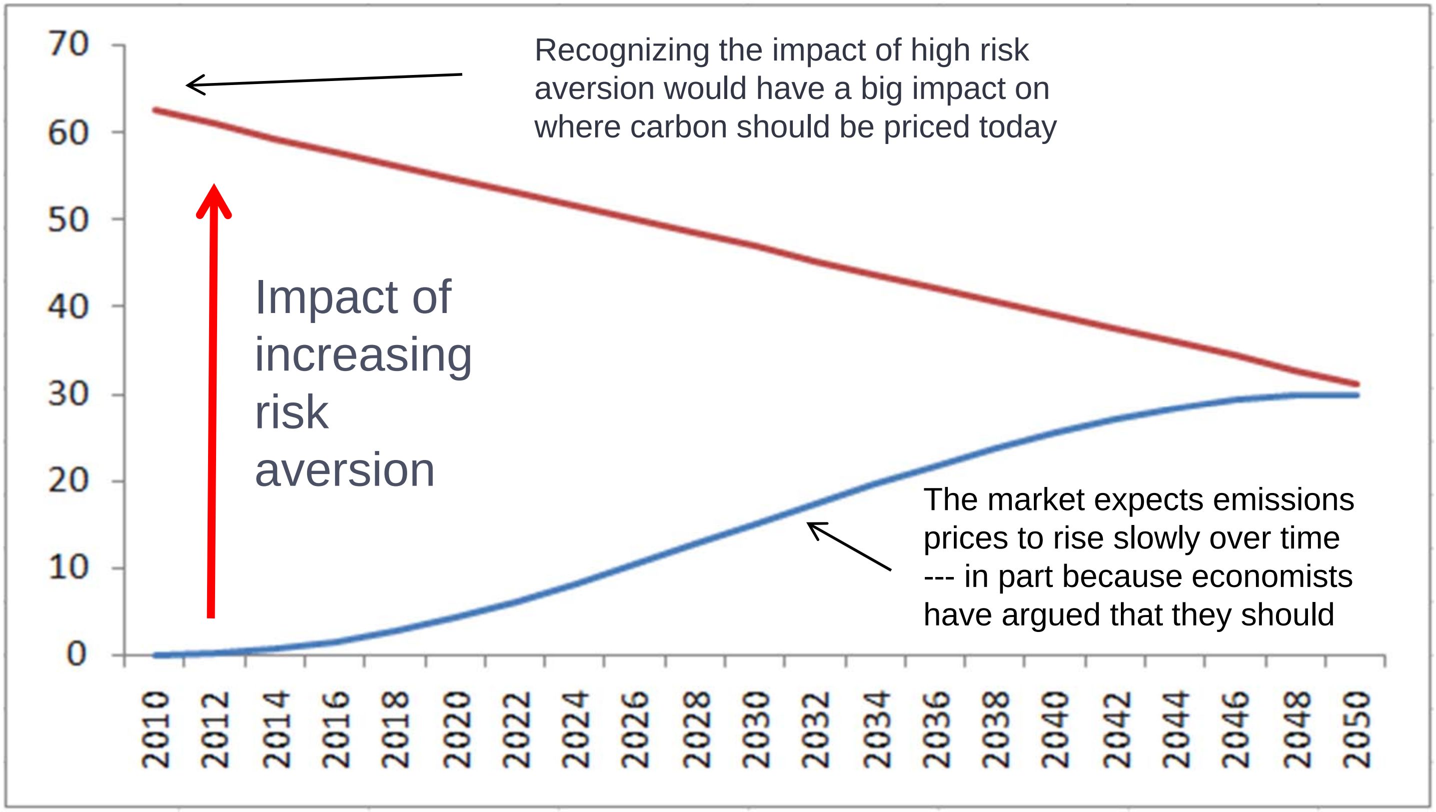

Why? Think about dynamic optimization

…with uncertainty, tipping points and nonlinear responses

A risk averse rider, when recognizing that he may be out of control, brakes hard and expects to ease off in the future as uncertainty is resolved

Slowly easing onto the brake makes sense when the road is well known and control is certain

The price of carbon emissions is our brake as we confront the uncertainty of climate

The appropriate price depends on:

risk

and

risk aversion

- The impact of today's emissions on:
  the **uncertain distribution** of future damages
- The **stochastic discount rates**
  used to calculate the present value of future damages
- The current and future **costs of emissions reductions**

Calculating the discounted present value of an investment in mitigating emissions

Welfare is a sum of discounted utility

$$W = \int_0^\infty u(c)e^{-\delta t}\,dt$$

where utility of future consumption is discounted at a rate of pure time preference, δ

The Stern report, which used an unusually low value for δ of .1% per year, has been highly criticized

Consumption in a **random state** of nature in the future must be discounted not only as a function of time, but also by:

- the **probability** of that state, and
- the **marginal utility** of that state

Thus damages are not discounted at a fixed rate, but rather at a **stochastic discount rate** conditional on the state in which they occur

The standard approach is to specify constant relative risk aversion utility

CRRA Utility $$u(c) = \frac{c^{1-\eta}}{1-\eta}$$

where the curvature paramater, η, the elasticity of marginal utility of consumption, determines **both** risk aversion and the elasticity of intertemporal substitution

In the absence of uncertainty
\$1 at time t is worth \$1 / $(1 + r)^t$ today
where $r \approx \delta + \eta g$

The discount rate r is a function of:

1. the rate of pure time preference: δ
2. the curvature of the utility function: η
3. and the growth of consumption: g

But CRRA utility has a well known problem

Risk aversion

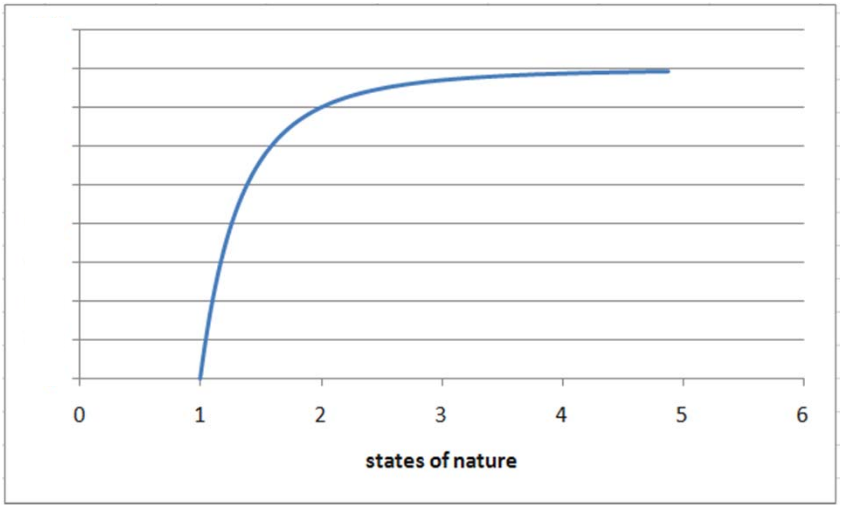

Intertemporal substitution

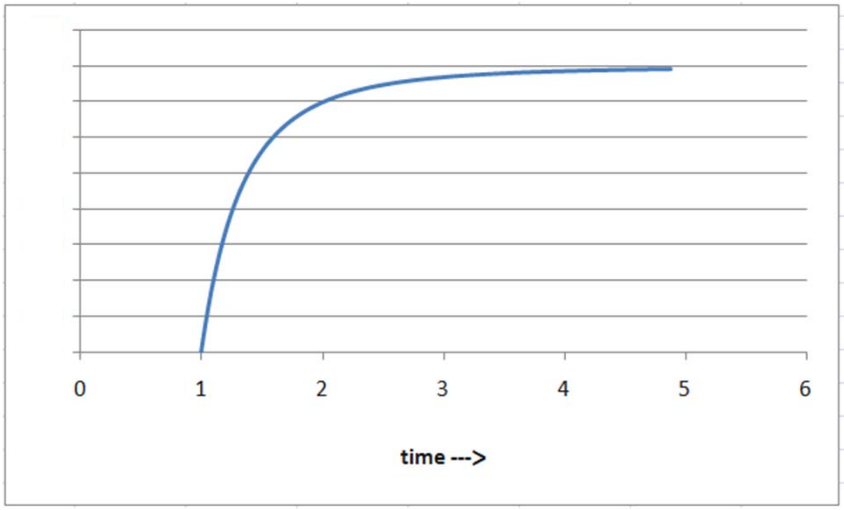

The curvature across states of nature required to generate a significant equity risk premium is much too large to be consistent with interest rates that induce people to postpone current consumption in favor of savings

These trade-offs are at the heart of evaluating optimal climate policies

Climate modelers generally choose a low value of η in the context of CRRA utility

Stern, for example, set $\eta = 1$ or "log utility"

This value implies an equity risk premium of around 12 basis points around 60 times too low relative to observed risk premia

Something is very wrong here

Counter to intuition **increasing the risk aversion** in the context of CRRA utility makes curbing emissions **less** urgent

Higher values of η **increase the discount rate** and imply people will be less willing to postpone current consumption

The higher discount rate lowers the weight put on climate induced damages expected in the distant future

Typical integrated assessment model results give risk and risk aversion virtually no role!

Estimates of the social cost of carbon from Anthoff, Tol, and Yohe (2009)

no uncertainty

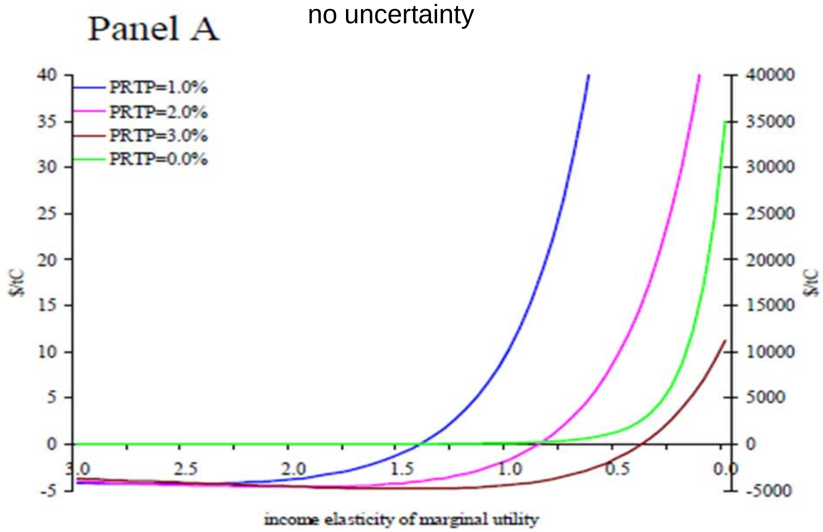

← Increasing risk aversion

with parameter uncertainty

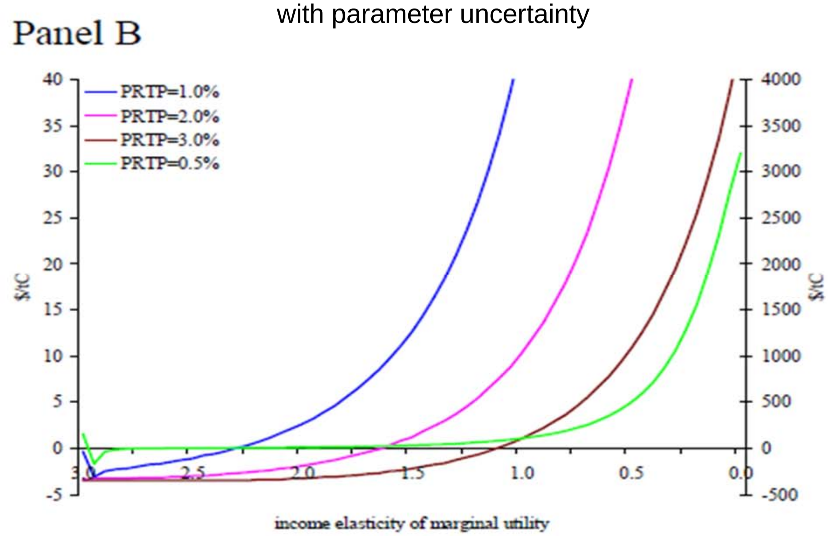

← Increasing risk aversion

Increased risk aversion combined with increased uncertainty
does not increase the social cost of emissions in their analysis

Carbon emissions deserve a risk premium: The issue is “how big?”

- Financial market practitioners are familiar with how to price uncertain future cash flows (such as damages caused by carbon emissions)
- We rely on stochastic discount factors
- Cash flows in **bad states of nature, where marginal utility is high,** are worth more than cash flows in good states of nature
- The benefits of emissions reductions today will accrue primarily in bad states of nature, therefore they are more valuable than risk free cash flows
- How big should the risk premium be?

A key issue in the economics of climate policy is how to **calibrate** intertemporal substitution and risk aversion

- The issue is **how to reconcile**
  - the high elasticity of intertemporal substitution evident in low **real interest rates**

  with
  - the high risk aversion evident in the very significant **equity risk premium**

- In the CRRA utility function the one parameter, η, determines both, but cannot be made consistent with market data.

- Epstein-Zin utility allows separate calibration to both interest rates and the equity risk premium

Epstein-Zin utility allows intertemporal substitution to be separated from risk aversion

Risk aversion

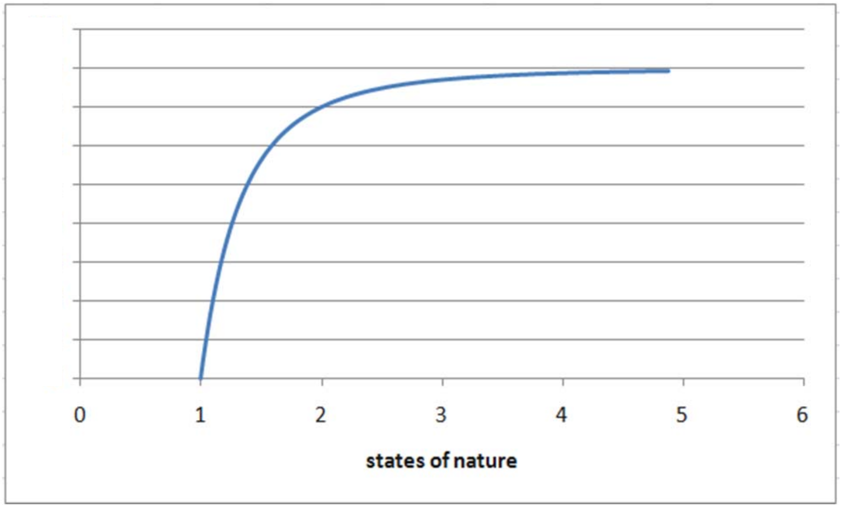

Intertemporal substitution

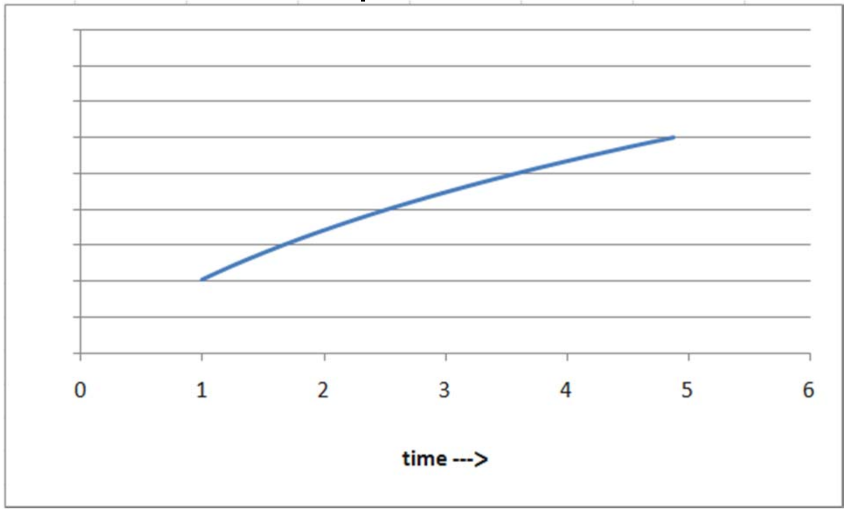

High curvature across states of nature can fit the very significant equity risk premiums that we observe in the market

While less intertemporal curvature can fit the relatively low risk free rates that we observe in the market

We pull these ideas together in a simple model

Following Weitzman, we add significant tails to the standard distributions of climate damage

We calibrate emissions reductions costs to previous studies

We specify an Epstein-Zin utility function

- we calibrate the intertemporal subsitution to market interest rates
- we investigate the impact of different risk aversion assumptions
- we find those results with risk aversion at least modestly consistent with equity market premia more interesting

In each case we solve for emissions reduction policies that maximize welfare

Economic impacts depend on future temperatures which are very uncertain

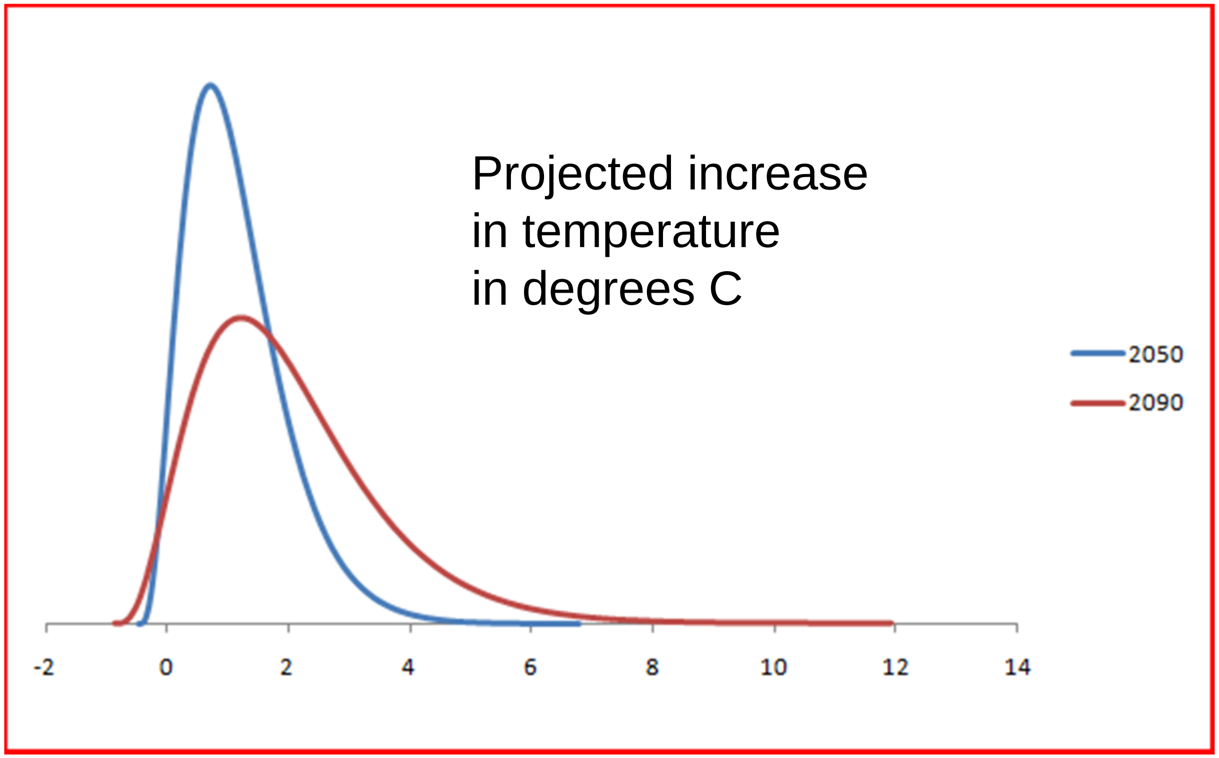

The distributions above are expected under “business as usual” scenarios

Many climate models take into account only known, or highly likely, damages

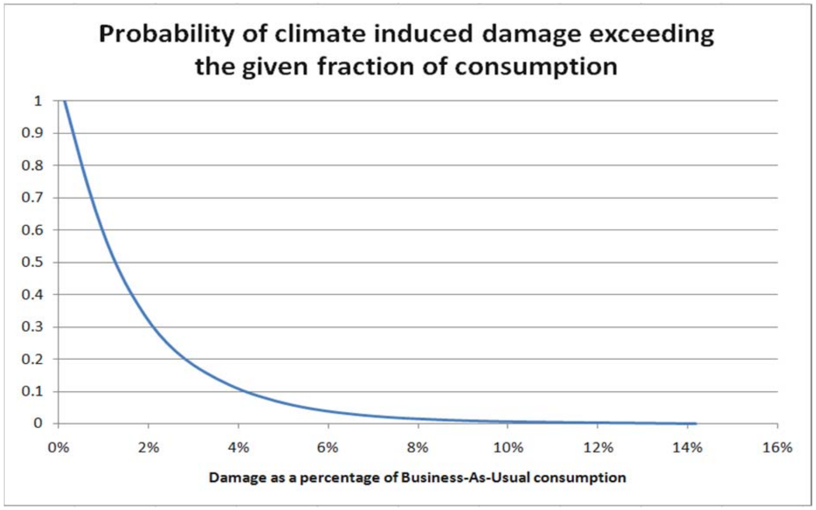

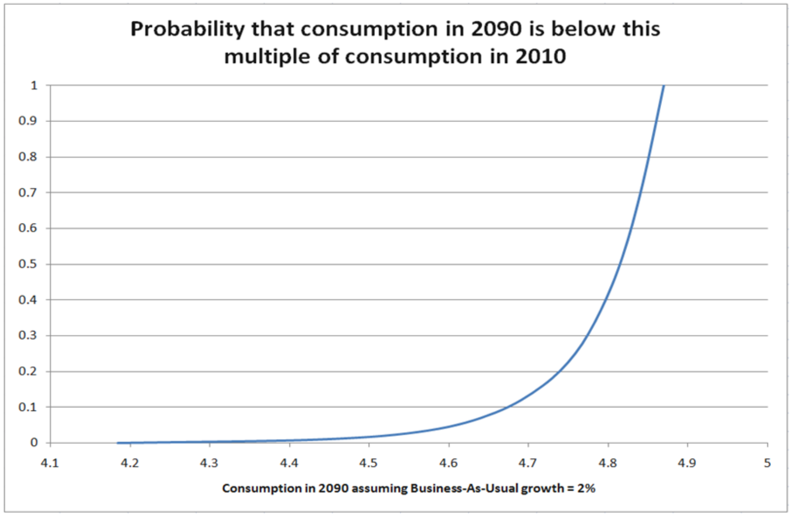

These calibrations, for example, which are designed to reflect the concensus damage estimates from integrated assessment models, come from Robert Pindyck's 2009 paper, "Uncertain outcomes and climate change policy"

Here the probability is less than .001 that real per capita consumption is impacted as much as 15% in 2090. Assuming 2% annual growth this implies we are 99.9% sure consumption will be > 4.1 times today's consumption

Many scientists, e.g. Sherwood and Huber (2010), feel "recent estimates of the costs of unmitigated climate change are too low"

Our analysis starts with the Pindyck distribution but recognizes additional uncertainty

- we add disaster scenarios
- in which case we expect significant damages

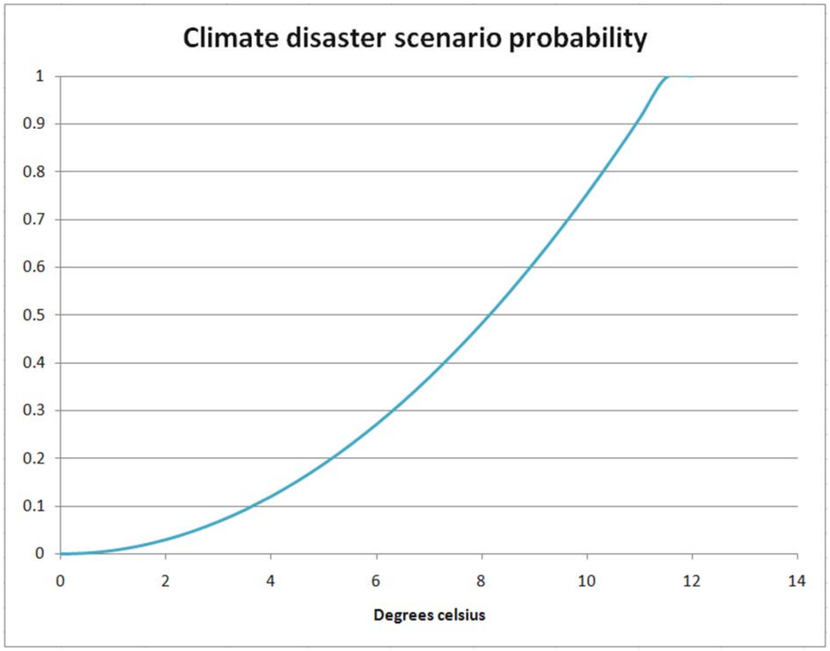

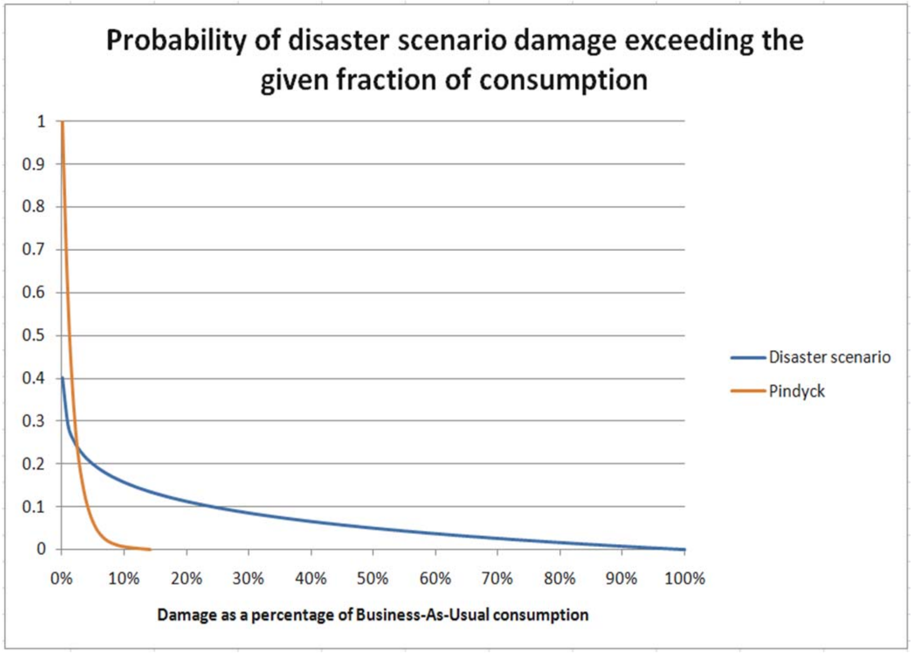

We specify priors over these scenarios and damage curves in order to generate more realistic uncertainty about future consumption

The probability of a significant reduction in consumption at future dates is much higher

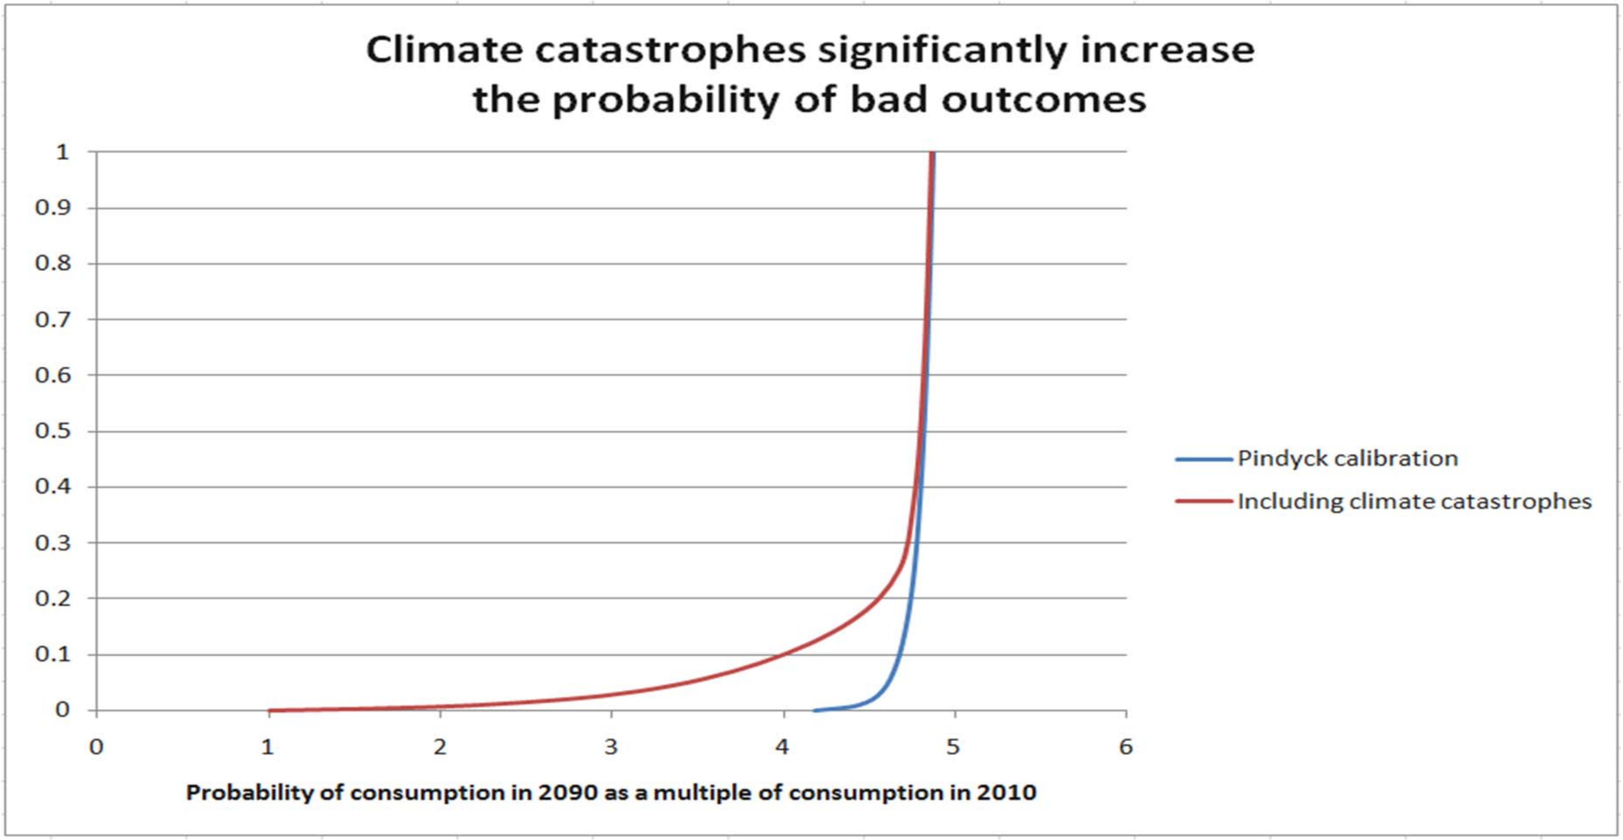

We calibrate our cost curve to the well known McKinsey analysis

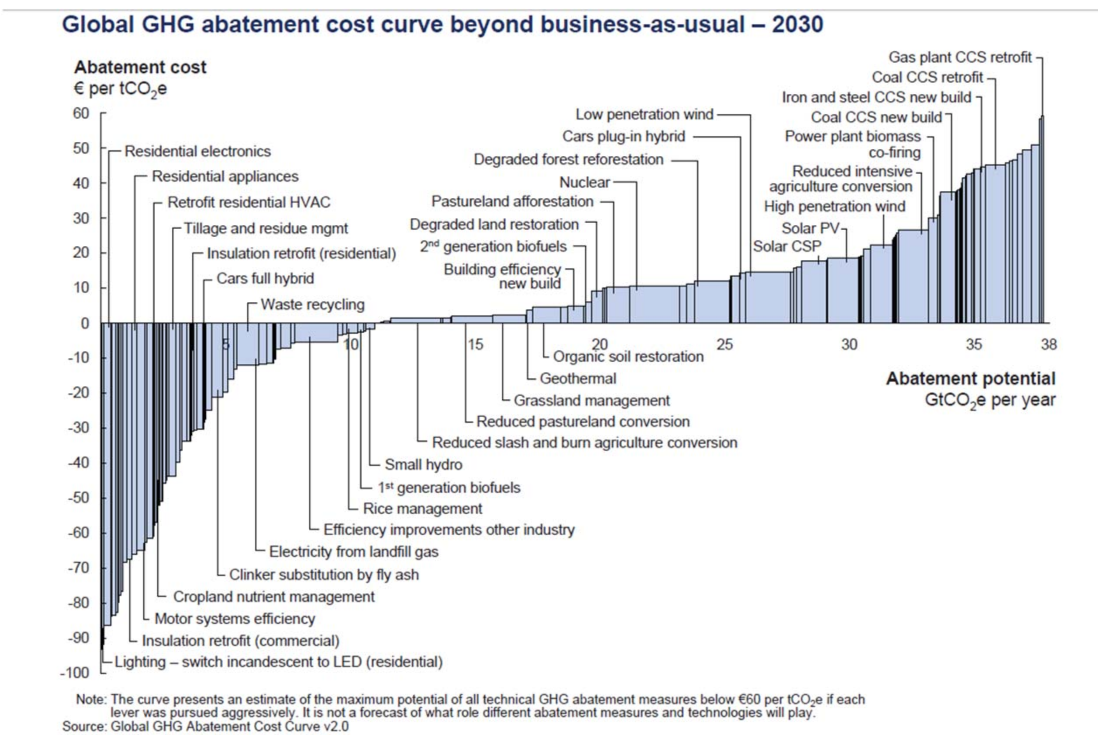

We use a simple two-period model adapted from Summers & Zeckhauser

Choose emissions reductions percentages e_{2010} and $e_{2050}(k)$

Observe the state, k, and implement emissions reduction $e_{2050}(k)$

k states index the uncertain fragility of the environment

2010

2050

2090

Consume $\tilde{C}_{2010} = [C_{2010} - cost_{2010}]$

Consume $\tilde{C}_{2050} = [C^{BAU}_{2050} - cost_{2050} - damage_{2050}]$

Consume $\tilde{C}_{2090} = [C^{BAU}_{2090} - damage_{2090}]$

$$\underset{[\text{over } e_{2010} \text{ and } e_{2050}(k)]}{\text{Max}} \quad \text{Welfare} = U_{EZ}(\tilde{C}_{2010}, \tilde{C}_{2050}, \tilde{C}_{2090})$$

Consumption

Business-as-usual consumption grows 2% per annum

$C_{2010} = 1$ $\quad C^{BAU}_{2050} = 2.2$ $\quad C^{BAU}_{2090} = 4.9$

Costs in 2010 of emissions reduction: e_{2010} are given by the cost curve:

$$cost_{2010} = cc \times e_{2010}^{\alpha}$$

Costs in 2050 of emissions reduction: $e_{2050}(k)$ are given by the cost curve:

$$cost_{2050} = cc_{2050} \times e_{2050}(k)^{\alpha}$$

Technology improves 1% per year

$$cc_{2050} = .67 \times cc_{2010}$$

Damages in 2050 from emissions are given by the damage function:

$$damage_{2050} = b \times d(k) \times (1 - e_{2010})^{q}$$

Damages in 2090 from emissions are given by the damage function:

$$damage_{2090} = d(k) \times (1 - .5 \times (e_{2010} \times e_{2050}(k)))^{q}$$

Optimization

We use Epstein-Zin utility and choose the elasticity of intertemporal substitution = 1.3, which together with a time discount rate, $\delta = .5\%$, implies a real interest rate of 2%

We investigate different levels of risk aversion

Risk Aversion	0.95	1.9	3.8	7.5	15	30	60
Equity Risk Premium	0.1%	0.2%	0.4%	0.8%	1.5%	2.5%	3.8%

Historical equity premia are greater than 6 percent.

Equity premia consistent with different values of risk aversion in the Epstein-Zin model (shown above) are calibrated from Weil (1989)

For each value we solve for a utility maximizing emissions reduction plan

We then look at various sensitivities

Model results as a function of risk aversion

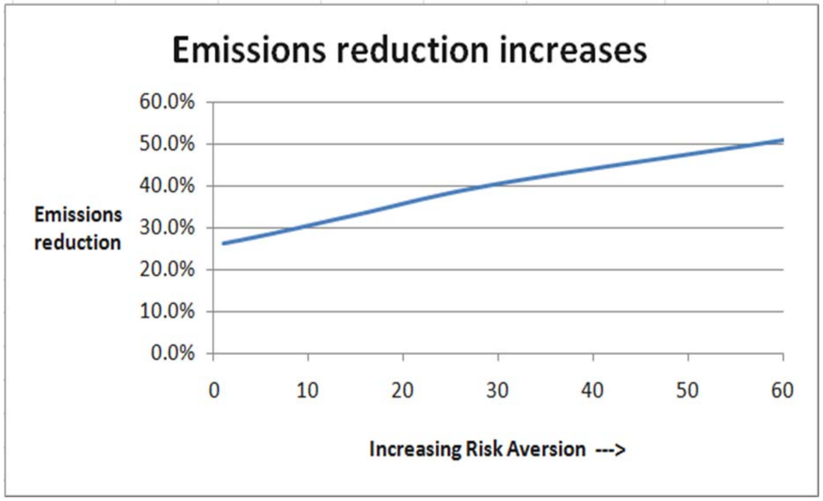

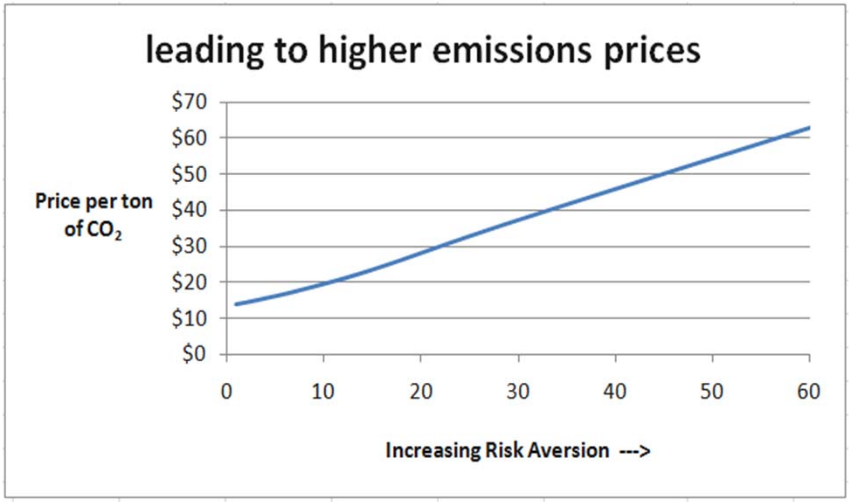

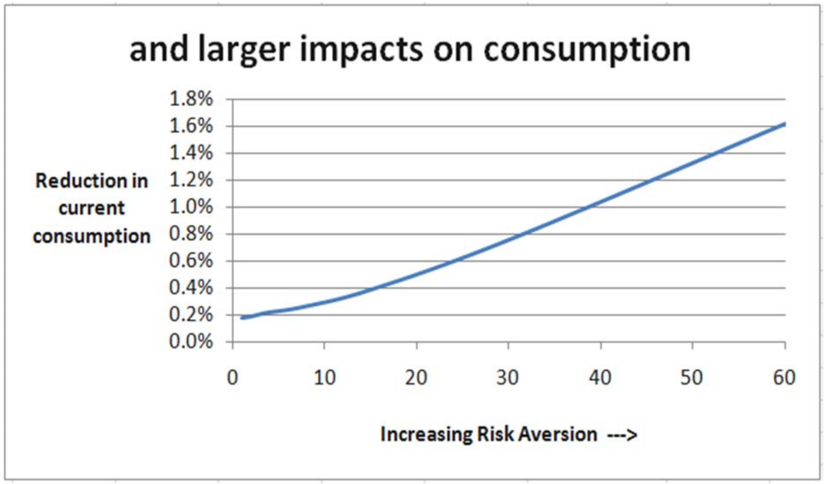

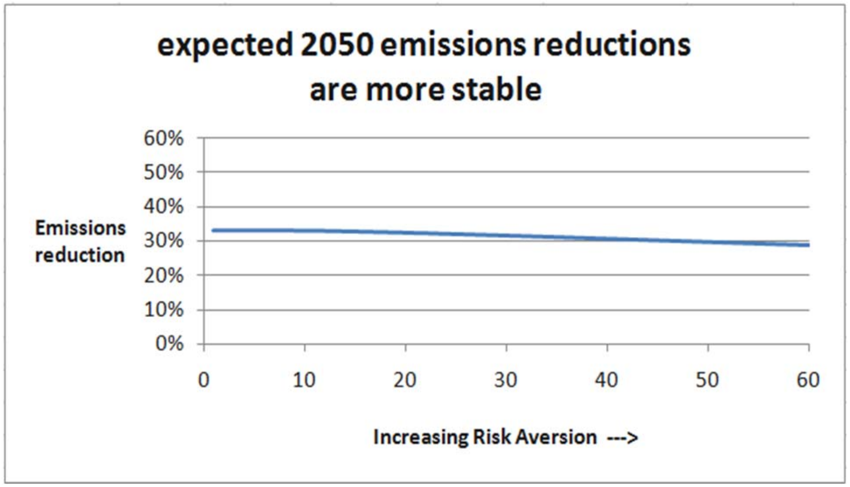

Results in 2050 as a function of the state

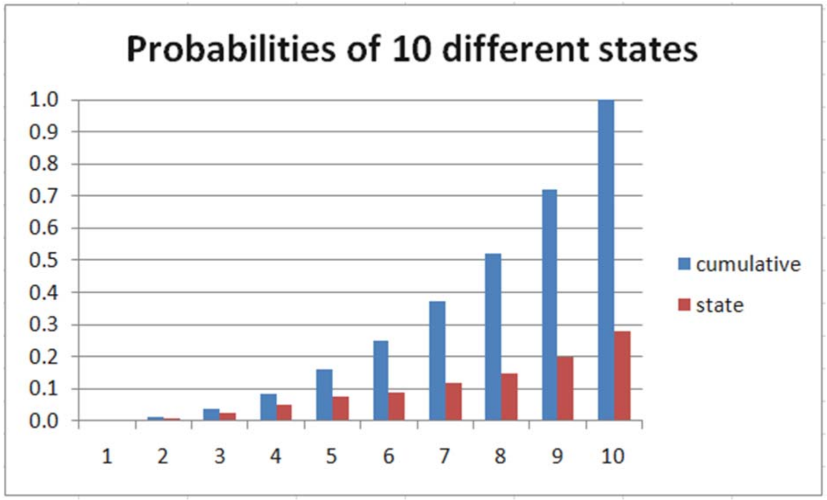

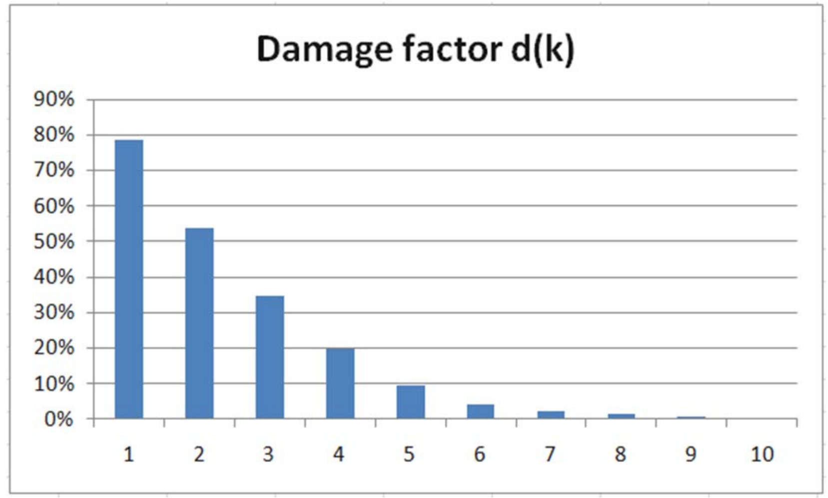

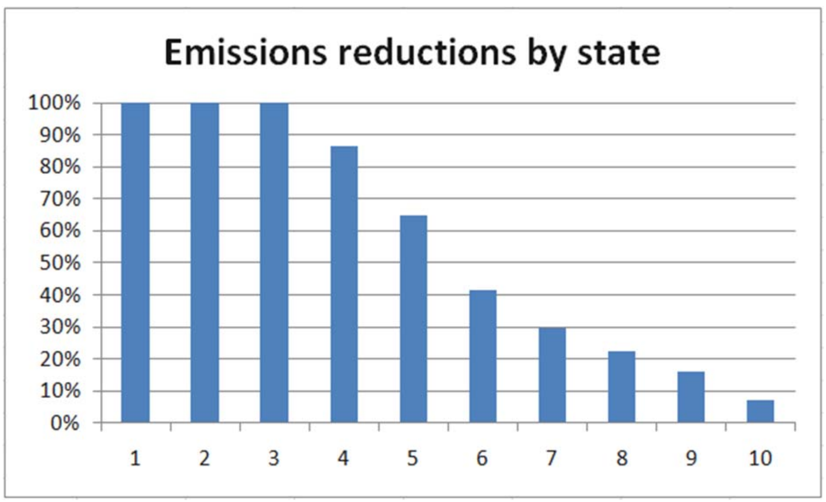

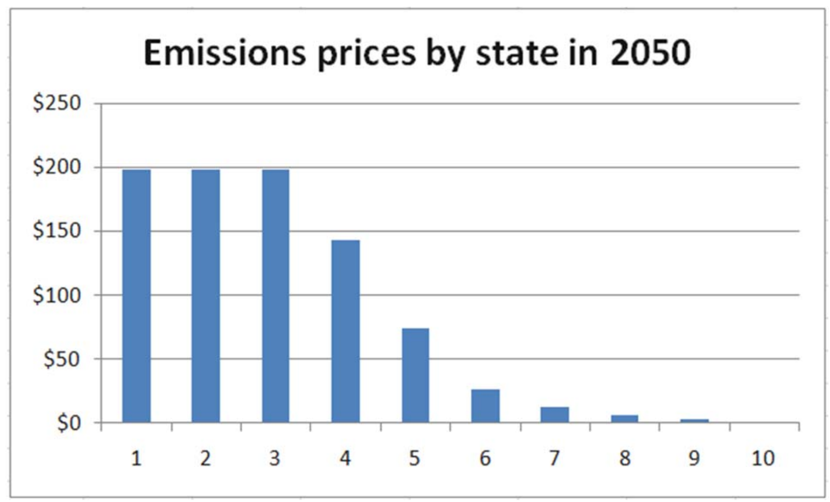

Current versus expected future emissions prices

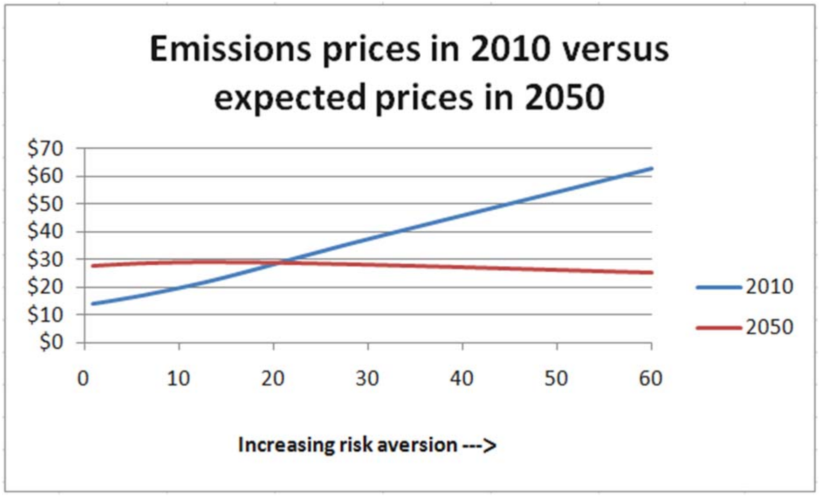

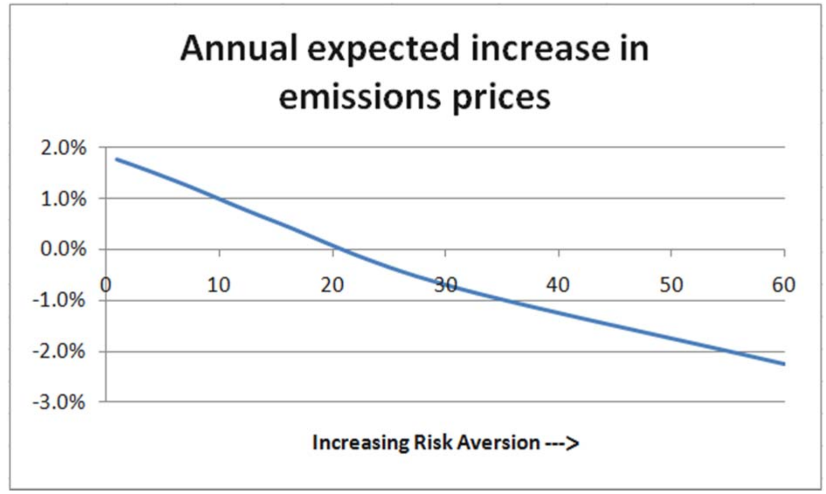

In a risk neutral world emissions prices start at around $13 and increase over time at the risk free rate

But in a risky and risk averse world where:

1. consumption in 2090 is expected to be reduced 4.2% by “business-as-usual” emissions,

2. there is a 10% probability that consumption in 2090 would be impacted > 20% by “business-as-usual” emissions and 1% chance of a > 50% impact, and

3. risk is priced to the same degree as seen in the equity market

- the risk premium on climate damages is -3% to -4% per year
- and the appropriate current price is in the range $40 to $60 per ton of CO_2

A simple cost-benefit view of optimal climate policy

	Business-as-usual	Log utility	Equity premium risk aversion
the impact on consumption in 2010	0	-0.20%	-1.59%
the expected impact on consumption in 2090	-4.23%	-2.17%	-1.47%
the 'worst case' scenario (probability = 1%)	-54%	-17%	-9.36%

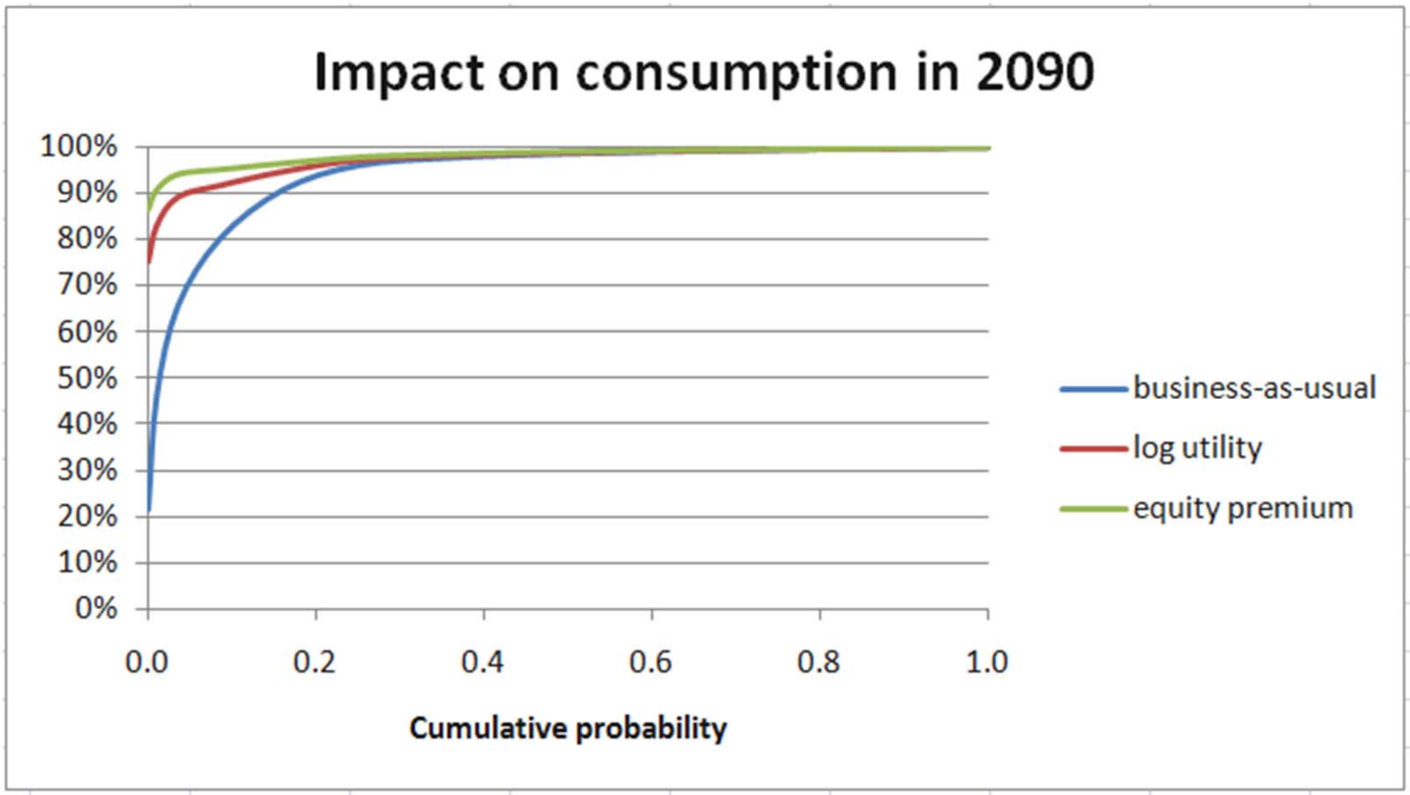

Sensitivity to price induced technological change

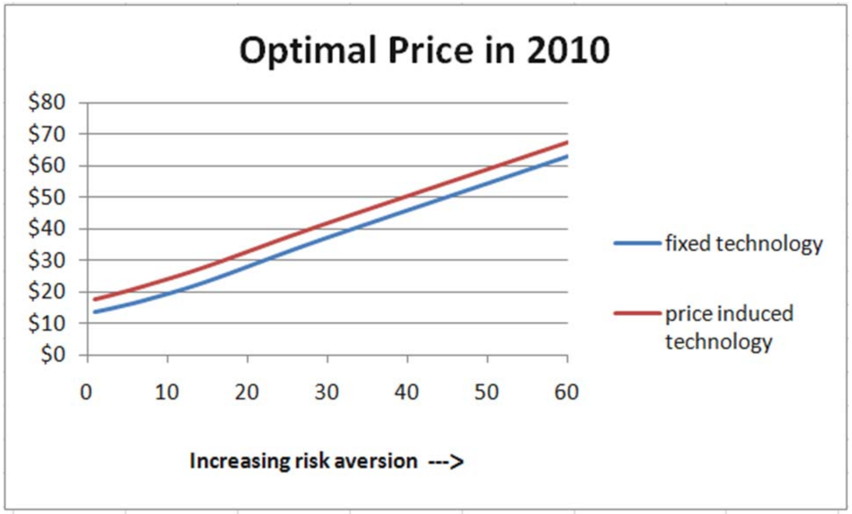

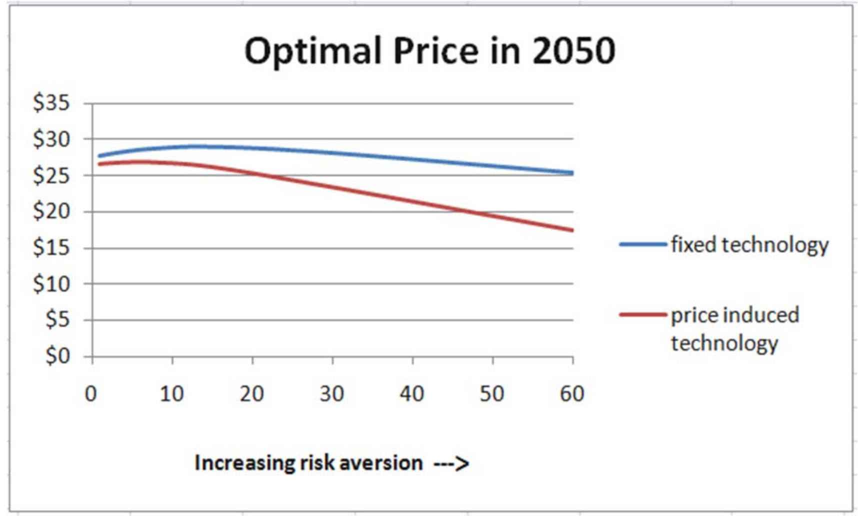

Our base case assumes technological innovation reduces the costs of emissions reduction decline by a fixed rate of 1percent per year

Here we investigate what happens if, instead, costs decline faster the higher emissions prices are. In particular we assume costs decline at a rate of $.8 + .03 \times \text{price}_{2010}$ percent per year from 2010 to 2050.

As expected, the optimal price is higher in 2010 and lower in 2050.

Sensitivity to the time discount rate

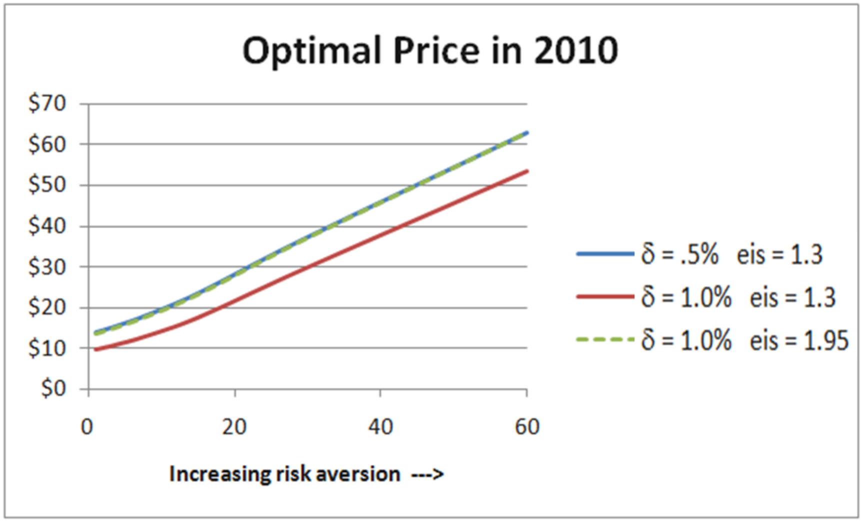

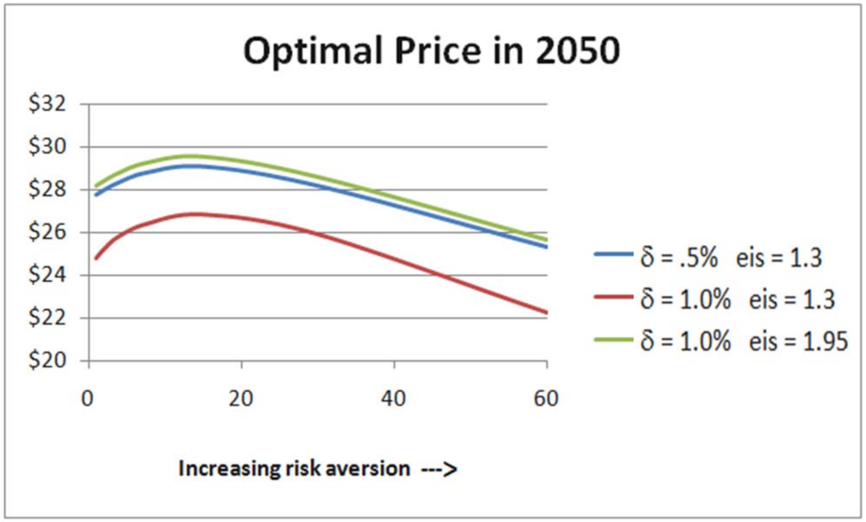

Our base case assumes the time discount rate δ = .5% per year. We then set the elasticity of intertemporal substitution, eis, to 1.3, which implies a real interest rate of 2%, matching historical data.

Here we investigate what happens if, instead, the time discount rate is raised to 1% per year, and if we then set the eis at 1.95, which also implies a real interest rate of 2%.

As long as eis is calibrated to set the real interest rate at 2% the rate of pure time discount has very little impact.

Impact of assuming no climate catastrophes

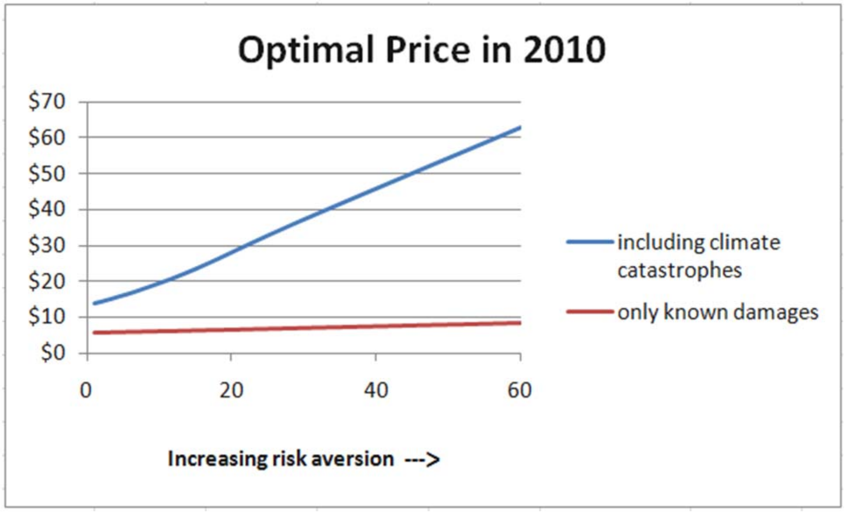

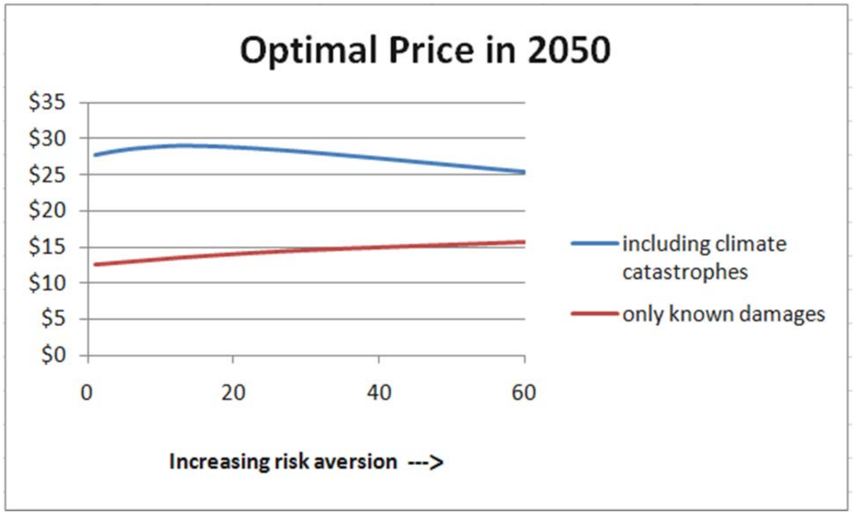

Our base case includes potential climate catastrophes with significant potential damages and with probabilities that increase with temperature.

Here we investigate what happens if, instead, damages are based only on the distribution reported by Pindyck. The Pindyck distribution, calibrated to those in many integrated assessment models, measures the impacts of damages that are known, or considered to be highly likely.

Viewed from today's consumption levels those damages represent very little real risk. Not surprisingly optimal emissions prices are much lower, particularly in the first period. We conclude that the main reason to price climate emissions is the risk that such optimistic assessments may be wrong.

CRRA versus Epstein-Zin Utility

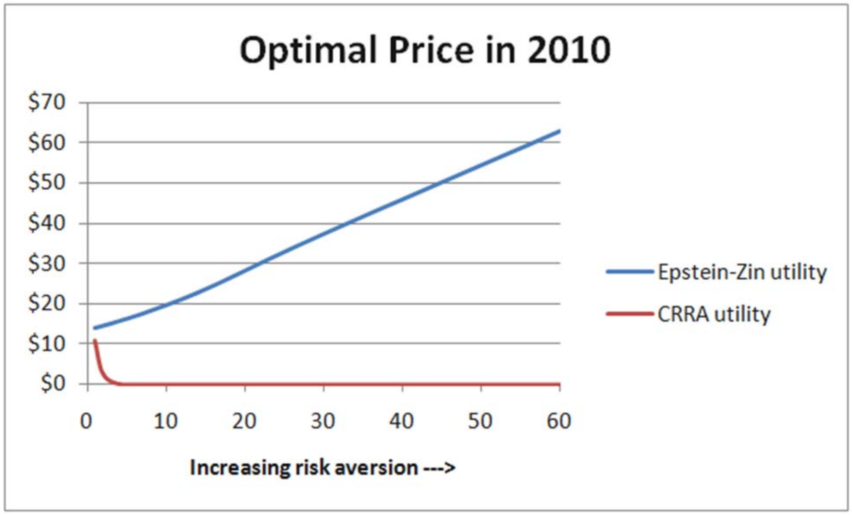

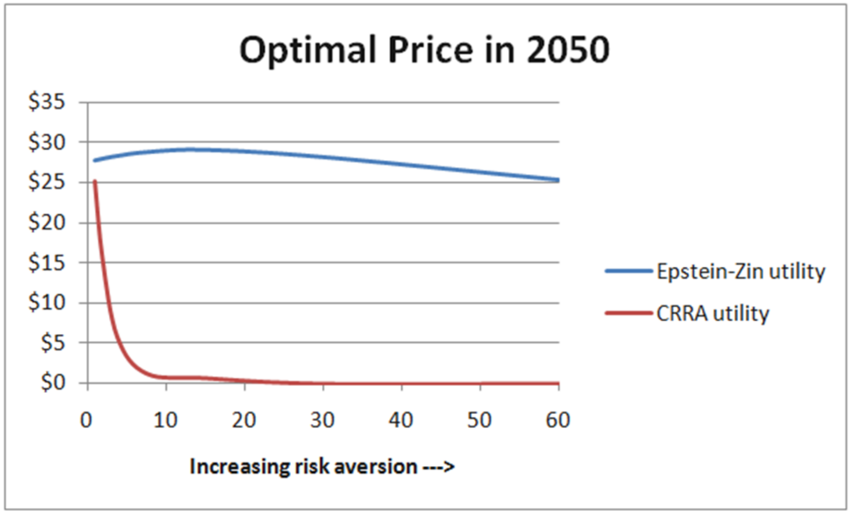

As seen earlier in the context of traditional integrated assessment models, increasing risk aversion in the context of CRRA utility leads to a fall in the optimal price of emissions.

The desire to smooth intertemporal consumption clearly dominates the aversion to future climate risk in the CRRA utility functional form.

Delay increases cost and reduces efficacy

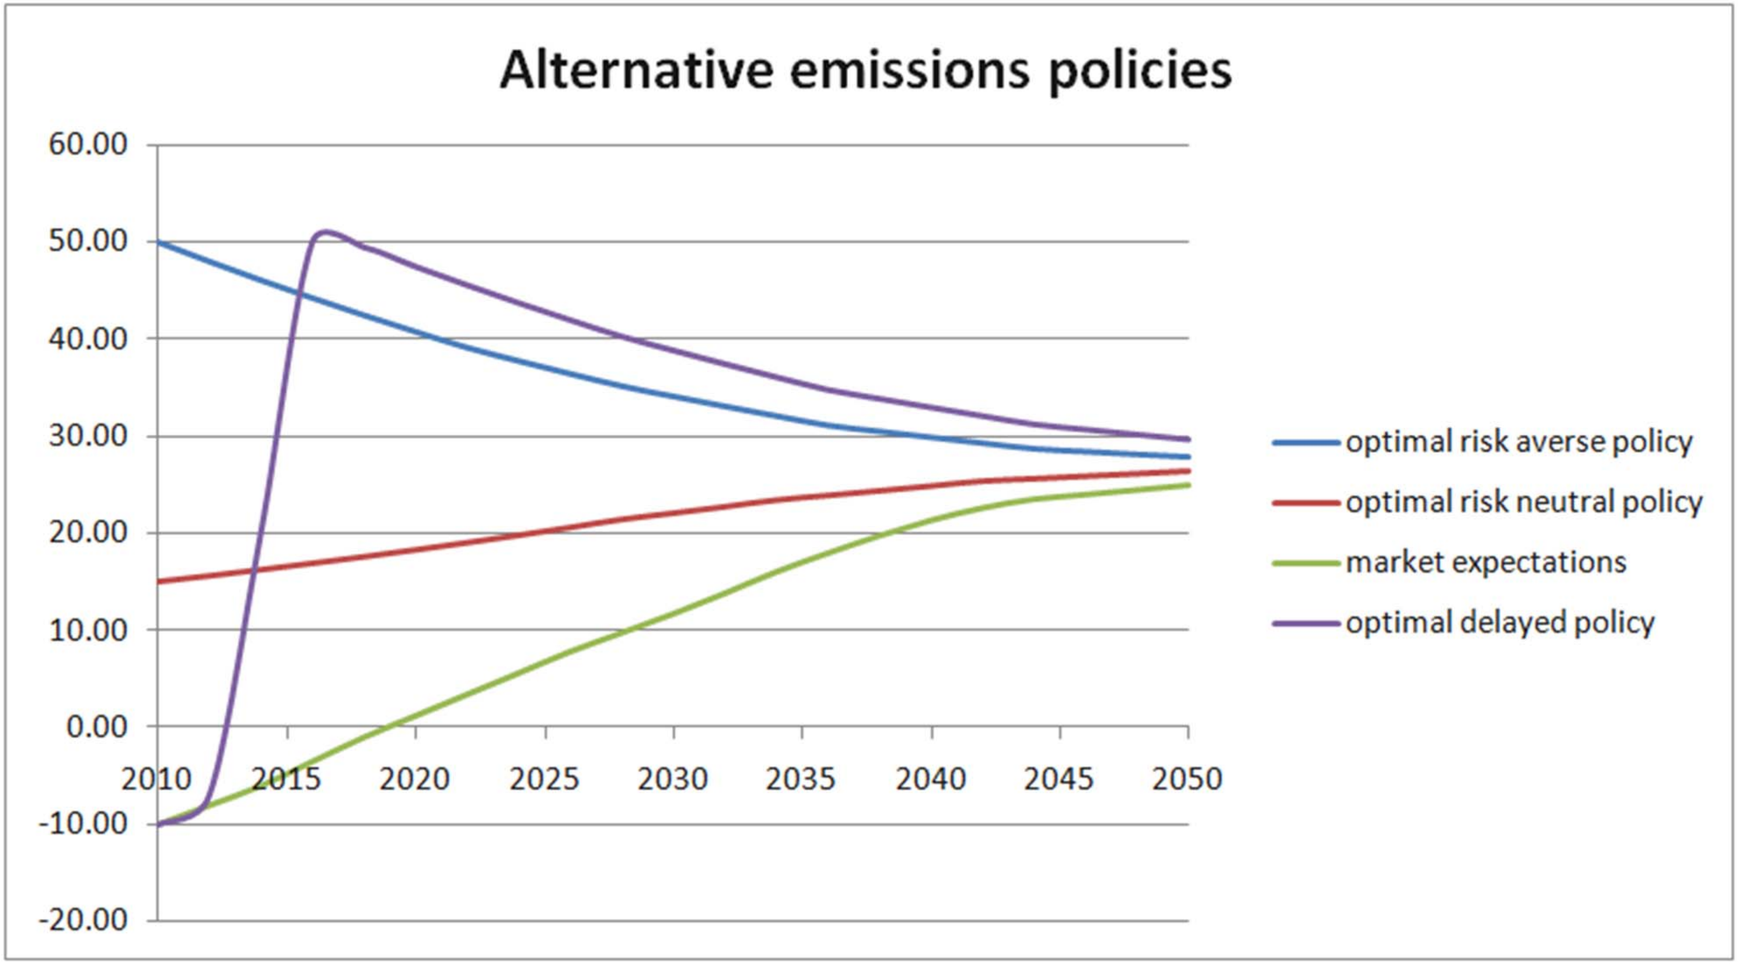

In the mean time:

benchmarks should reflect the current disequilibrium in the price of emissions

- Market capitalization weights currently seem to reflect expectations of a very slow increase in emissions prices over time

- At some point the market will recognize that emissions will need to be priced at much higher levels

- Long term investors should position themselves ahead of time

- Relative to current market capitalization weights:
  - Benchmarks should tilt
    - toward assets that will benefit, and
    - away from assets that will suffer

    from higher emissions prices

Summary Points

1. Consistent with the scientific consensus, we incorporate the possibility of tipping points, nonlinearities, and global catastrophes
   - A low probability of very bad outcomes
2. We calculate the damages associated with an additional ton of carbon emissions by discounting these damages using a stochastic discount rate that reflects the state of the world in which the damages are incurred
   - In the bad states of nature our progeny would be willing to give up a lot of ordinary consumption goods to have a more habitable world
3. We use market prices (i.e., the risk-free rate and the historical equity premium) to calibrate the stochastic discount factor parameters
4. The risk premium associated with the systematic risk of climate change significantly increases the appropriate current price of carbon emissions