

Five Essays in the Economics of Climate Engineering,  
Research, and Regulation under Uncertainty

Dissertation

zur Erlangung des akademischen Grades  
“doctor rerum politicarum” (Dr. rer. pol.)

eingereicht beim Promotionsausschuss  
der Fakultät für Wirtschafts- und Sozialwissenschaften  
der Ruprecht-Karls-Universität Heidelberg

von  
Daniel Heyen  
geboren 7. Juni 1981 in Köln

im Februar 2015



## Acknowledgements

First and foremost, thanks goes to Eva-Maria for her love and support, and to Ilias for joy and meaning. I also want to thank my parents, sisters, grand-parents and parents-in-law for constant support. It is a shame that many of them are not here anymore.

I owe much to my supervisor Timo Goeschl. For stimulating discussions, invaluable lessons, and for steering me thoughtfully through the first stages of my academic career.

I wish to thank Anastasios Xepapadeas for being the second supervisor of this thesis.

The members of the interdisciplinary Marsilius project “The Global Governance of Climate Engineering” provided an inspiring environment for the exchange of ideas, concepts and questions from which my research benefited in various ways.

Finally, I want to thank my former and current colleagues at the Chair of Environmental Economics, in particular Johannes Diederich, Dietrich Earnhart, Kathrine von Graevenitz, Ole Grogro, Johannes Jarke, Ole Jürgens, Sara Kettner, Johannes Lohse, Tobias Pfrommer, Daniel Römer, and Israel Waichman.

## Table of Contents

Introduction .....	1
<i>Paper 1</i>	
The Intergenerational Transfer of Solar Radiation Management Capabilities and Atmospheric Carbon Stocks .....	17
<i>Paper 2</i>	
Strategic Conflicts on the Horizon: R&D Incentives for Environmental Technologies .....	41
<i>Paper 3</i>	
Informativeness of Experiments for MEU – A Recursive Definition .....	67
<i>Paper 4</i>	
Learning under Ambiguity: A Note on the Belief Dynamics of Epstein and Schneider (2007) .....	72
<i>Paper 5</i>	
Information Acquisition under Ambiguity – Why the Precautionary Principle may Keep us Uninformed .....	91
Conclusion .....	122

# Introduction

At first unwittingly, but for the last couple of decades in full conscience, mankind has been substantially altering the climate by emitting greenhouse gases (GHG), mostly carbon dioxide. Due to the high inertia of the climate system, the warming effect of past emissions has only partly unfolded. A drastic cut of GHG emissions, yet currently at an all-time high, would thus be needed to at least curb global warming (IPCC 2013). Further aggravating the matter, the climate response from GHG concentrations to temperatures – though being only one link in the causal chain – is characterized by substantive uncertainty (Meinshausen et al. 2009), implying that the projections for global warming vary greatly. For instance, the likely range for the global mean surface temperature increase at the end of the 21st century under the RCP 6.0 concentration scenario extends from 1.4°C to 3.1°C (IPCC 2013). In addition, there are arguments that low probability scenarios of 'catastrophic' climate change with temperature increases well beyond the upper bound of the range deserve more attention (Weitzman 2011; Pindyck 2011; Nordhaus 2011).

Policy inaction and the outlook of climate change entailing substantial damages (IPCC 2014), potentially 'catastrophic', have spurred discussions about solar radiation management (SRM), a set of technologies potentially capable of quickly reducing global temperatures by "decreasing the amount of absorbed solar radiation through an increase in albedo" (Keith 2000). The most popular SRM proposal at present is 'stratospheric aerosol injection' with sulfur (Crutzen 2006; Keith 2013). Justifications for seriously considering SRM follow either a 'buying time' or a 'last resort' narrative (Gardiner 2011). In the 'buying time' narrative (Keith 2013), SRM stabilizes temperatures for a transformation to a low-emission economy by rebuilding the energy sector. Optimally managing the system may also involve to complement SRM deployment with the reduction of atmospheric GHG concentrations by means of carbon dioxide removal (CDR) techniques, besides SRM the second branch of climate engineering (CE) technologies. In the other justification, the 'last resort' narrative, the merit of SRM lies in the option to quickly respond to a 'climate emergency' that may result from an unlikely, yet devastating sharp temperature rise (Crutzen 2006; Victor et al. 2009; Weitzman 2011).

Despite possible unknowns unknowns from their deployment and already foreseeable, yet currently hard to quantifiable side-effects on precipitation patterns and stratospheric ozone (Shepherd 2009), it seems uncontroversial that "SRM methods, if realizable, have the potential to substantially offset a global temperature rise" (IPCC 2013).

Whether SRM is realizable is unclear at present; but even if so, costly and multi-year research and development (R&D) would be needed to make SRM available (Klepper and Rickels 2012; Keith 2013). It is against this background of potential benefits of SRM remaining inaccessible without a long and costly development process that voices urge small-scale outdoor experiments as the next research steps (Parson and Keith 2013; Parker 2014; Keith et al. 2014).

For coming to a fair assessment on whether R&D into SRM should proceed, one should bear in mind that both narratives favoring SRM rely on strong consistency assumptions. The 'buying time' narrative is based on a social planner who can commit to a specific SRM deployment profile stretching across decades, furthermore complemented by ambitious efforts in abatement and energy innovations (Barrett et al. 2014). The 'last resort' narrative relies on similar preconditions. Essential in this idea, current generation 'arms the future' with the 'lesser evil' option SRM (Gardiner 2010), only to be deployed in case of a 'climate emergency'. In so doing, the 'last resort' narrative silently rests on the assumption that a precise definition of a climate emergency is possible and persistent over time (Gardiner 2011; Markusson et al. 2014).

In light of these stark preconditions, various concerns have been raised with regard to SRM deployment profiles and its repercussions on other policy options. One concern is whether the sustained use of SRM can be guaranteed (Barrett et al. 2014). The failure to do so would cause temperature to rise rapidly, a harmful scenario as many climate damages are associated with the rate of temperature change rather than absolute temperature values (IPCC 2013). The probably best-known concern surrounding SRM runs under the slightly misleading label 'moral hazard'. In this concern, the prospect of a 'quick-fix' technology might create a perception of insurance and hence reduce the incentives for costly GHG mitigation (Shepherd 2009; Hale 2012; Morrow 2014; Gardiner 2010; Reynolds 2014; Betz 2012). Other concerns revolve around the deployment profile of SRM. Highlighting the possibility of irrational technology use, the American Meteorological Society (2009) is worried of "short-sighted and unwise" deployment. Similar, and closely related to the 'last resort' narrative, the definition of a 'climate emergency' is not unambiguous; rather, it can be re-defined and used rhetorically (Markusson et al. 2014). Other authors challenge the social planner perspective and emphasize that SRM would be used in a world with international actors and heterogeneous motives, thus stressing the potential of deployment conflicts (Robock 2008; Weitzman 2012).

In summary, society is confronted with pressing questions surrounding SRM. Does the prospect of a technology offering to temporarily stabilize temperatures and to stand by as a quick emergency measure justify considerable R&D expenditures, and, if yes, how should the R&D process proceed? And to what extent does this assessment of such an ideal technology change in light of the above mentioned concerns?

To contribute to this fascinating and important research field, environmental economics can rely on a rich toolkit including a vast literature on climate economics

(Nordhaus 1994; Stern 2006; Weitzman 2009), environmental innovation (Popp 2010; Goeschl and Perino 2007; Barrett 2006; Hall and Helmers 2013), and the rigorous game-theoretical analysis of strategic interaction in the presence of environmental externalities (Barrett 1994, 2013; Missfeldt 1999; Wagner 2001). Equipped with these general insights and tools, authors have made specific contributions to climate engineering and SRM in particular. Goes et al. (2011) analyze the danger of intermittent SRM and provide a framework for conducting a cost-benefit test of SRM. Bickel and Agrawal (2013) identify factors that crucially shape whether SRM passes such a cost-benefit test. Moreno-Cruz and Keith (2012) focus on the value of SRM as a last resort option under different assumptions on learning. The interplay of SRM and GHG mitigation has been analyzed in a portfolio set-up (Kousky et al. 2009), in intertemporal settings (Moreno-Cruz and Keith 2012), and along the international dimension (Moreno-Cruz 2010; Manoussi and Xepapadeas 2014). Strategic deployment conflicts have been the focus of contributions by Moreno-Cruz (2010) and Weitzman (2012).

The present literature on environmental technologies in general, and climate engineering in particular, however misses a relevant angle. What has received hardly any attention is a positive analysis of innovation incentives when actors anticipate that a technology they are about to develop would be prone to strategically induced deployment profiles. This constitutes an important research gap since, as was demonstrated above, current discussions about SRM are characterized by the joint presence of considerable concerns regarding technology deployment and the need for making consequential R&D decisions in the near future. The first two articles of this dissertation aim at filling this gap. They focus on two strategic problems about SRM – one of them intergenerational, the other intragenerational – and analyze the repercussion of these strategic conflicts for the incentives to undertake costly R&D. Both articles establish novel findings and thus contribute to an improved understanding of environmental innovation. A short summary of the articles in this dissertation can be found at the end of this introduction.

In order to put the strategic interplay of deployment and innovation of climate engineering center stage, the first two articles of this dissertation deliberately abstracted from a couple of issues, most notably uncertainty. Climate uncertainty, one of the main justifications for considering SRM in the first place, has been incorporated – if at all – only in the crudest form. This may have been justified in order to focus on the strategic dimension, but it has obviously left aside important normative questions surrounding SRM. Being a novel technology and thus, almost by definition, not well-understood, SRM deployment itself would add substantial uncertainty. Thus to prevent harm, a 'precautionary' case against deployment (Long and Winickoff 2010) or large-scale testing (Robock 2012) can be made. On the other hand, and in stark contrast, the idea of precaution can also be invoked for the development of SRM in order to be prepared for a 'climate emergency' (Hartzell-Nichols 2012) and, in light of possible climate tipping

points that might be crossed, even for precautionary SRM deployment (Barrett et al. 2014).

The ambivalent nature of SRM reveals, in a nutshell, some of the challenges modern society is confronted with. Potential risks are pervasive both in decisions related to complex systems like climate change, food safety, pandemics and biodiversity (Maslin and Austin 2012; May 2001) and in choices surrounding novel substances or technologies like asbestos, drugs, pesticides or climate engineering (Jasanoff 2007; Sunstein 2002; Wiener et al. 2013). The complexity of the system or the novelty of the interaction, or both, give rise to a fundamentally different uncertainty than standard risk (Randall 2011). It is this context of fundamental uncertainty and consequential societal decision in which the Precautionary Principle (PP) emerged. Lacking a clear and unique definition (Cooney 2004), but being intrinsically tied to the presence of fundamental uncertainty, the PP has emerged over the last decades and by now plays an important, yet controversial role in today's regulation, mostly in Europe (Wiener et al. 2013; Zander 2010).

The PP appears to comprise two notions with slightly different focus. The first notion of the PP, often referred to as 'better safe than sorry', urges the decision-maker to build up a 'margin of safety' or to impose 'safe minimum standards' against potential and uncertain harm (Ciriacy-Wantrup 1952; Pielke 2002). The second notion, referred to in the literature as 'look before you leap' or 'precautionary learning', emphasizes the importance of information and learning in decisions characterized by fundamental uncertainty (Doremus 2007; Bourg and Whiteside 2009). Closely connected to each other, the 'better safe than sorry' notion focuses more on how to make a decision under some static level of uncertainty in which learning already occurred or is not possible. In contrast, the 'look before you leap' notion of the PP puts potential improvements in the state of knowledge, maybe by active information acquisition, center stage.

Many commentators are critical of the PP and challenge its importance for risk regulation. Widely raised arguments are that the PP stifles innovation, privileges irrational fears and is inherently inconsistent (Graham 2004; Sunstein 2005; Peterson 2006). Opponents of the PP argue that there already exists a well-established practice of risk management. This 'ordinary risk management' (ORM, Randall 2011) rests on cost-benefit analysis (CBA) and incorporates uncertainty by probabilities based on objective risk assessments. Through these features, the ORM claims to avoid biases in decision-making, to uphold incentives for innovation, and to maximize overall welfare.

But there are a couple of issues with ORM. The first problem is that standard cost-benefit approaches often fail to give a full account of the options available to a decision-maker; as a consequence, ORM tends to ignore irreversibilities, leading to a biased assessment of investment and conservation decisions (Dixit and Pindyck 1994; Arrow and Fisher 1974; Henry 1974; Krutilla 1967; Bishop 1982). The other concerns about ORM specifically refer to the type of uncertainty that is present in the regulation of complex and novel risks. The first line of criticism holds that ORM is, due to its foundation

on expected utility theory, incapable of meaningfully compare unlikely events with high impacts, for instance climate related damages under very high temperature increases (Tol 2003; Chichilnisky 2010). A second and related line of criticism is that complex and novel risks, by definition, preclude the kind of precise description of uncertainty that is required in ORM (May et al. 2008; Betz 2007; Meinshausen et al. 2009).

Environmental economics has contributed to the discussion about the appropriate risk regulation and the PP in various ways. The first line of research keeps the basic premises of the ORM, in particular the expected utility assumption, and extends it into different directions. Gollier et al. (2000) and Gollier and Treich (2003) analyze sequential decision-making characterized by irreversibility and uncertainty in the tradition of Arrow and Fisher (1974) and Henry (1974), but with partial resolution of uncertainty and thus Bayesian learning. The 'precautionary effects' identified by Gollier et al. (2000) and Gollier and Treich (2003) consist of a reduction in an irreversible action, i.e. less consumption of a toxic product, when there is the prospect of learning by 'waiting for better information'. Their settings thus give rise to behavior consistent with the 'precautionary learning' notion of the PP. In a different approach, Naevdal and Oppenheimer (2007) and Margolis and Naevdal (2007) allow for the possibility of thresholds. They demonstrate that the presence of thresholds of unknown location induces a precautionary behavior of less economic activity, consistent with the 'better safe than sorry' notion of the PP. Similar, Weitzman (2013) shows that the existence of heavy tails in the probability distribution for adverse events may induce cautious behavior. These contributions are important extensions of the standard ORM framework and explore its relation with the PP. It is important to note that the precautionary behavior that emerges in these set-ups, as Randall (2011) puts it, is not 'principled' but 'circumstantial'. It is a specific behavior of an ORM decision-maker that results because of the underlying structure of the problem, not because the decision-maker adopted the PP.

The second line of research in environmental economics takes a different approach and directly associates the PP with the underlying principles of decision-making. In so doing, this strand of the literature goes beyond the expected utility framework. It describes scientific uncertainty with 'ambiguity', also known as 'Knightian uncertainty', and operationalizes the PP, according to the 'margin of safety' interpretation, as an ambiguity averse decision rule (Asano 2010; Athanassoglou and Xepapadeas 2012; Millner et al. 2013; Lemoine and Traeger 2012; Vardas and Xepapadeas 2010; Treich et al. 2013). The best-known decision rule in this context is maxmin expected utility (MEU), axiomatized by Gilboa and Schmeidler (1989). Scientific uncertainty is described by a set of probability distributions, and the decision-maker selects the decision that maximizes expected utility under the worst probability distribution. It is this focus on the worst probability scenario that makes MEU, in close connection to the idea safe minimum standards, "a conceptual framework for designing management rules which adhere to the PP" (Vardas and Xepapadeas 2010). The MEU approach has been used in bio-



diversity settings (Vardas and Xepapadeas 2010) and stock pollutant problems (Asano 2010; Athanassoglou and Xepapadeas 2012). The findings support the precautionary interpretation of MEU. The MEU decision-maker is more conservative, sooner adopts environmental policies, and invests more in damage control.

What is underrepresented in the present literature is the possibility of *active* information acquisition. If the papers cited above incorporate learning, it is of the passive form; neither the timing nor the extent of learning can be influenced by the decision-maker. The possibility to actively shape the learning process, however, is a standard regulatory task: Prior to consequential decisions surrounding novel substances and technologies, regulators typically undertake (costly) research to inform their decision-making. This active information acquisition is much broader than the 'wait to learn' specification in Gollier et al. (2000) and Gollier and Treich (2003). The underrepresentation of active research choices and a joint analysis of the two notions of the PP thus constitute an important research gap in the theory of regulation. This dissertation makes the first steps into addressing these issues. It follows the literature by modeling the PP as MEU. Then, based on a definition of the value of information under MEU (the third article) and making sure that learning dynamics under MEU are still compatible with ambiguity aversion (the fourth article), the fifth article focuses the implications of MEU for active learning and thus analyzes whether the 'better safe than sorry' notion of the PP translates into the 'precautionary learning' notion. As we will see, the answer is in general negative.

## Synopsis

The first article "**The Intergenerational Transfer of Solar Radiation Management Capabilities and Atmospheric Carbon Stocks**" (with Timo Goeschl and Juan Moreno-Cruz), in content almost identical with Goeschl et al. (2013), published in *Environmental and Resource Economics*, focuses on the intergenerational conflict that arises because the donor of the technological capability cannot stipulate its use by the recipient. An important contribution of the paper is to test existing narratives about the impact of SRM on abatement for their internal consistency.

The model features two non-overlapping generations. Current generation decides about abatement (quadratic abatement costs), and whether to develop SRM (a binary decision with fixed development costs) and thus to 'arm' future generation against severe climate change (Gardiner 2010; Betz 2012). Altruism of current generation towards the future is a prerequisite because both investments, abatement and technology provision, only pay-off to the future generation. If current generation developed the technology, future generation can decide on the SRM level in order to partially compensate excessive radiative forcing by the GHG stock, however incurring SRM side-effect damages (linear in the level of deployment). How much SRM future generation undertakes de-

depends on three things. First, the level of abatement by current generation (in other words the size of the atmospheric carbon stock), secondly the climate sensitivity that resolves in an intermediate stage (with two possible values and initially known probability distribution), and finally – a central component of the model – a bias parameter that captures to what extent future generation underestimates the damages of SRM deployment. There are various justifications for assuming such a biased view on SRM: in a behavioral dimension, future generation may be irrationally over-willing to stop a climate emergency; in a political economy dimension, capacity building for SRM may cause interest groups eager to promote technology deployment. Whatever the reason for the biased view on SRM damages, it gives rise to an intergenerational conflict about the use of the technology. Current generation, unable to stipulate a certain deployment profile, exerts influence on future generation’s behavior by adjusting the R&D decision and/or the level of abatement. The solution of this intergenerational game is derived by means of subgame perfectness.

We find a couple of possible equilibria. As a useful reference point, we first analyze the case of an unbiased damage assessment without the intergenerational conflict. The benchmark is calibrated such that current generation develops the technology as an emergency measure for the high climate sensitivity outcome; due to the absence of a bias, this specific deployment profile is consistent with the deployment incentives of future generation. Unsurprisingly, and as a simple consequence of the substitutability of costly abatement with technology deployment, the abatement level in the benchmark equilibrium is lower than in the absence of the SRM option. This equilibrium then serves as a benchmark for the equilibria with a non-vanishing bias. We find technology provision equilibria as well as equilibria in which current generation decides to deny future generation the technological option; this technology denial gets more plausible the higher the bias in the damage assessment and the higher the costs of R&D. Intriguingly, and demonstrating the inconsistency of some narratives in the SRM literature, we do not find any equilibrium involving a *decrease* in abatement relative to the benchmark. Instead, an *increase* in abatement together with technology provision is a possible outcome of this intergenerational game. The deeper reason is that current generation is altruistic towards its successor. In this line of thinking, increasing abatement is a strategic tool for current generation to partially correct for the anticipated SRM over-deployment in the future.

The second article, entitled **”Strategic Conflicts on the Horizon: R&D Incentives for Environmental Technologies”**, puts intragenerational conflicts about climate engineering center stage. The focus is on the question how anticipated deployment conflicts impact on innovation incentives. The contribution of my article is to provide a framework for analyzing this interplay and to demonstrate that the anticipated strategic conflicts surrounding deployment essentially carries forward to the R&D decision.

The model is a simple two-country setting with two decision stages. In the first stage, the countries non-cooperatively choose their R&D contributions. If the sum of contributions fails to exceed a commonly known threshold, no country can use the technology; otherwise the technology is available to both countries and they choose their technology deployment levels non-cooperatively in period 2. The quadratic deployment costs are fully private while the benefits, as usual in the public good literature, depend on the sum of technology contributions. This structure embraces CDR as well as SRM technologies: No country can be excluded from the reduced temperatures that result from removing greenhouse gases or adding reflective particles, but the costs of undertaking climate engineering fully fall on the deploying country. The benefit function is also quadratic (inverse U-shaped), reflecting the fact that excessive cooling is detrimental. A key feature of my model is to allow for heterogeneity in the preferred level of climate engineering deployment.

The first insight is that the technology deployment game in the second stage is not only able to produce standard free-riding equilibria, but also *free-driving* behavior (Weitzman 2012). In this only recently appreciated possibility, the country with the smaller preference for SRM cooling, not being able to exclude itself from the high deployment of the other country, does not contribute in equilibrium. Whether the technology deployment game is a free-rider or a free-driver game is determined by the the cost parameter. This comprehensive description of different climate engineering technologies within one framework is one of the deliverables of my paper.

The second insight is that the model offers a parsimonious, yet highly functional framework for analyzing the repercussion of deployment conflicts on R&D incentives: Because a threshold public good game in general, and the first period R&D game in particular, suffers in equilibrium neither from over- nor underprovision, any deviation from a social optimal R&D outcome can be fully attributed to the outlook of suboptimal technology deployment patterns. Indeed, I find that the prospect of free-riding for high cost technologies like CDR weakens R&D incentives, implying the existence of parameter constellations for which a CDR technology ought to be developed but remains undeveloped due to weak incentives. Less obvious is the outcome for low cost technologies like SRM characterized by free-driving. Even though the country with the smaller preference for global cooling anticipates to be dominated by the other country's high deployment, it may still be better off than without the technology. In other words, the willingness to pay for R&D can still be positive and high; for small costs and high heterogeneity across countries, however, the picture inevitably changes and the country's willingness for R&D turns negative, demonstrating that the anticipated conflict about the right SRM level may carry forward to the innovation stage. The flipside of this is that, with the lure of low deployment costs, the incentives of the other free-driver country to develop the technology may be so strong that technologies are developed even though it is against the global best.

In the following, the focus of the dissertation shifts away from the strategic aspects of climate engineering and turns to environmental regulation under fundamental uncertainty. The fifth and main article revolves around the value of research under the Precautionary Principle (PP). Bringing these two topics together not only establishes a novel angle in the literature on environmental regulation; also also on the fundamental decision-theoretic level, the interplay of ambiguity aversion, learning, and value of learning has not received adequate attention. Accordingly, the third and fourth article of this dissertation contribute to the decision-theoretic foundation of learning under ambiguity.

To appreciate the relevance of the third article "**Informativeness of Experiments for MEU – A Recursive Definition**" (with Boris Wiesenfarth), in content almost identical with Heyen and Wiesenfarth (2015), forthcoming in the *Journal of Mathematical Economics*, some background on the theory of decision-making under uncertainty may prove helpful. The basic formulation of the decision-problem under uncertainty involves a decision-maker without knowledge which value in a set  $\Theta$  is the true state of the world  $\theta$ . This uncertainty is economically relevant if the payoffs that result from the decision-maker's action  $a$  depend on the true state of world  $\theta$ .

The first important extension of this basic framework is when the decision-maker can learn. Learning means that, at some point in time prior to the decision, the decision-maker gets information about the true state  $\theta$ . This information is usually noisy, often not precluding any state of the world, but rendering some more and others less likely. This information, often called a *signal*, clearly has economic value as it enables the decision-maker to improve her decisions. The literature on this *value of information*, in more or less general settings, dates back to Blackwell (1953) with many applications in all economic fields, including environmental economics (e.g. Fisher and Hanemann 1987; Kolstad 1996; Gollier and Treich 2003; Karp and Zhang 2006).

The other extension of the simple decision-making under uncertainty framework revolves around the description of uncertainty. The traditional approach (Savage 1972), also pervasive in the value of information literature cited above, captures uncertainty in a unique probability distribution on  $\Theta$  (for instance in the form  $(p, 1-p)$  when there are two possible states of the world). This standard approach has been questioned for descriptive and normative reasons: on descriptive grounds because Ellsberg (1961) has shown that individual behavior is often inconsistent with such a unique probability distribution; on normative grounds because many decision situations are too poorly understood ('Knightian uncertainty' or 'ambiguity') as to justify a unique probability distribution ('risk') (Knight 1921). Thus, and as an alternative to decision-making based on subjective expected utility theory (Savage 1972), ambiguity averse decision rules have emerged (for a survey see Etner et al. 2012), often associated with the notion of precaution. As mentioned before, a prominent example, and also the choice in this dissertation, is to generalize the assumption of a unique probability distribution and model initial knowl-

edge by a *set* of priors, and to choose as the decision-rule the ambiguity averse maxmin expected utility MEU.

Having now introduced the two components *value of information* and *ambiguity aversion*, a natural question is about their interplay, namely the value of information under ambiguity aversion. As mentioned earlier, the novel regulatory dimension of this question is the interplay of two notions of the PP. The decision-theoretic significance of this question also seems obvious in light of ambiguity as 'uncertainty aversion' and learning as a possibility to reduce uncertainty. Hence even more surprising is that this interplay has received surprisingly little attention so far. The first attempt to give a general definition of the value of information under MEU preferences in the fundamental Blackwell setting was Çelen (2012). His definition, however, violates *dynamic consistency*, a central rationality criterion (Machina 1989). Dynamic consistency requires the decision-maker to take into account all possible contingencies, form plans, and stick to these plans if nothing unexpected happens. In that sense dynamic consistency is closely related to the standard solution concept *backward induction* in dynamic decision-making (Riedel 2009). In order to define a version of the value of information under MEU that respects this important rationality criterion, the third article in this dissertation takes the dynamic consistent axiomatization of MEU preferences by Epstein and Schneider (2003), applies it to a standard Blackwell setting of decision-making under uncertainty, and thus delivers a sound definition of the value of information und MEU preferences.

What was not central in the definition of the value of information under MEU was a specification of how the set of initial priors, after observing the signal conveying information about the true state of the world, is updated to the set of posteriors. Any economic application, in with particular multi-stage models, however makes it necessary to clearly specify the learning dynamics. The important work in this context is Epstein and Schneider (2007), embedding the dynamic consistent axiomatization of MEU preferences (Epstein and Schneider 2003) into a useful and simple to utilize intertemporal decision-framework. Not surprisingly, the cornerstone of the learning dynamics in Epstein and Schneider (2007) is to update every initial prior individually in the standard way (i.e. Bayesian Updating, see Savage 1972). An additional and less obvious feature of their model is to identify priors that have become implausible given the signal history and remove them, thus slimming down the set of posteriors. It is this rejection of priors, and the problematic consequences it entails, that is the focus of the fourth article, entitled "**Learning under Ambiguity: A Note on the Belief Dynamics of Epstein and Schneider (2007)**".

The motivation of Epstein and Schneider (2007) for incorporating the rejection of beliefs into their framework, even making rejection mandatory, may have been that the learning dynamics otherwise can be trivial. If a decision-maker initially holds the full set of possible priors, including extreme cases<sup>1</sup>, then full Bayesian updating without

---

<sup>1</sup>For instance (1,0) in the two states setting, reflecting subjective certainty that the first state is the

the rejection of priors will preclude any form of learning. The set of posteriors will, irrespective of the quality of the signal observed, remain the full set of beliefs forever. Rejection of implausible priors indeed addresses this problem, but does so at a high price. As my article demonstrates, the learning dynamics in Epstein and Schneider (2007) are in conflict with ambiguity aversion. Irrespective of how 'modest' the rejection of priors is designed, there always exist conditions under which a decision-maker characterized by the Epstein and Schneider (2007) learning dynamics, initially ambiguity averse, ex-post prefers to bet on an urn with unknown composition rather than a known composition urn. This is a clear contradiction to ambiguity aversion. This contradiction is particularly noteworthy as it is the main purpose of Epstein and Schneider (2007) to provide a tractable framework of intertemporal ambiguity averse preferences.

My article offers two modifications of the Epstein and Schneider (2007) framework to ensure the compatibility of learning dynamics with ambiguity aversion. The first modification still involves the rejection of priors, but makes sure that those priors 'essential' for ambiguity averse preferences are immune to the rejection procedure. The second modification, which appears preferable on grounds of simplicity, is to stick with full Bayesian updating and to abstain from the rejection of priors altogether. What is needed to avoid the collapse of the learning dynamics mentioned above is just to preclude the full set of priors as a possible prior set. I argue that this formal restriction is hardly limiting the modeling purposes in applications because the full prior set together with the focus on the worst prior inherent in MEU is an extreme case of ambiguity aversion hardly meaningful in applications anyway.

Finally, and building on these decision-theoretic foundations, the fifth article "**Information Acquisition under Ambiguity – Why the Precautionary Principle may Keep us Uninformed**" (with Timo Goeschl and Boris Wiesenfarth) compares research incentives under different regulatory mandates. Particular focus is on the question whether MEU, often associated with the precautionary notion 'better safe than sorry', is compatible with the notion of precautionary learning that urges better information in order to reduce erroneous decisions.

At the center of the model is uncertainty about a payoff relevant state of the world and the need for the decision-maker to choose among two actions, and none of the actions dominating the other. Concrete regulatory examples we discuss in the paper are (1) the European Food Safety Authority's need to single out and ban from the market the vegetable that caused a food-borne disease and (2) the Environmental Protection Agency's approval or non-approval decision on a novel pesticide. In both cases – and this is a general feature of the regulation of complex and novel matters – the decision-maker is confronted with the sort of fundamental scientific uncertainty that makes an empirically substantiated risk assessment infeasible. Possible regulatory mandates are, as explained above, to make the problem accessible to standard CBA by assuming a

---

true one

'best guess' probability distribution and, in sharp contrast, to adopt the PP. Decision theoretic formalizations of these regulatory mandates are subjective expected utility (Savage 1972) and maxmin expected utility (Gilboa and Schmeidler 1989).

Another common feature of the regulatory process is that the regulator can commission costly research to gain (incomplete) information about the true state of the world to improve the final regulatory decision. This research is modeled as a noisy one-shot signal with the learning dynamics, informed by the fourth article, assumed as full Bayesian updating. The precision of this signal is a choice variable to the decision-maker, reflecting the regulator's possibility to shape the research process. Our model thus involves active learning as a key component. Crucial in this context is the definition of the value of information under MEU. This definition is provided by the third article.

Our findings suggest that MEU is in general not compatible with the notion of precautionary learning. While we were actually able to isolate a mechanism how the PP, formalized as MEU, increases the demand for research precision relative to a standard CBA – an effect that we accordingly coin the 'Precautionary Learning Effect' – we also find a countervailing and thus research precision decreasing effect of MEU. This 'Research Pessimism Effect' however often outweighs the Precautionary Learning Effect. Hence, in total, PP regulation can in many settings be expected to *decrease* the regulator's information level. The article, informed by insights about the origin of the two countervailing effects, sketches a possible regulatory set-up that may avoid the inconsistency of MEU with the notion of precautionary learning. Splitting up the regulatory mandate – standard CBA for the research decision and PP for the final regulatory decision – would avoid the Research Pessimism Effect and unambiguously improve the regulator's state of knowledge.

## References

- American Meteorological Society (2009). Geoengineering the climate system: A policy statement of the American meteorological society, *Bulletin of the American Meteorological Society* **90**(9): 1369–1370.
- Arrow, K. J. and Fisher, A. C. (1974). Environmental preservation, uncertainty, and irreversibility, *Quarterly Journal of Economics* **88**(2): 312–319.
- Asano, T. (2010). Precautionary principle and the optimal timing of environmental policy under ambiguity, *Environmental and Resource Economics* **47**(2): 173–196.
- Athanassoglou, S. and Xepapadeas, A. (2012). Pollution control with uncertain stock dynamics: When, and how, to be precautionary, *Journal of Environmental Economics and Management* **63**(3): 304–320.
- Barrett, S. (1994). Self-enforcing international environmental agreements, *Oxford Economic Papers* **46**: 878–894.
- Barrett, S. (2006). Climate treaties and "breakthrough" technologies, *The American Economic Review* **96**(2): 22–25.
- Barrett, S. (2013). Climate treaties and approaching catastrophes, *Journal of Environmental Economics and Management* **66**(2): 235–250.

- Barrett, S., Lenton, T. M., Millner, A., Tavoni, A., Carpenter, S., Anderies, J. M., Chapin III, F. S., Crpin, A.-S., Daily, G., Ehrlich, P., Folke, C., Galaz, V., Hughes, T., Kautsky, N., Lambin, E. F., Naylor, R., Nyborg, K., Polasky, S., Scheffer, M., Wilen, J., Xepapadeas, A. and de Zeeuw, A. (2014). Climate engineering reconsidered, *Nature Climate Change* **4**(7): 527–529.
- Betz, G. (2007). Probabilities in climate policy advice: a critical comment, *Climatic Change* **85**(1-2): 1–9.
- Betz, G. (2012). The case for climate engineering research: an analysis of the 'arm the future' argument, *Climatic Change* **111**(2): 473–485.
- Bickel, J. and Agrawal, S. (2013). Reexamining the economics of aerosol geoengineering, **119**(3-4): 993–1006.
- Bishop, R. C. (1982). Option value: an exposition and extension, *Land Economics* pp. 1–15.
- Blackwell, D. (1953). Equivalent comparison of experiments, *The Annals of Mathematical Statistics* **24**: 265–272.
- Bourg, D. and Whiteside, K. H. (2009). Precaution and science-based environmental risk management: Complementary not contradictory, *Building Safer Communities: Risk Governance, Spatial Planning and Responses to Natural Hazards*, Landsdale, pp. 88–104.
- Çelen, B. (2012). Informativeness of experiments for MEU, *Journal of Mathematical Economics* **48**: 404–406.
- Chichilnisky, G. (2010). The foundations of probability with black swans, *Journal of Probability and Statistics* **59**(2): 184–192.
- Ciriacy-Wantrup, S. V. (1952). *Resource conservation: economics and policies*, University of California Press.
- Cooney, R. (2004). *The Precautionary Principle in Biodiversity Conservation and Natural Resource Management: An issues paper for policy-makers, researchers and practitioners*, IUCN.
- Crutzen, P. J. (2006). Albedo enhancement by stratospheric sulfur injections: A contribution to resolve a policy dilemma?, *Climatic Change* **77**(3-4): 211–220.
- Dixit, A. K. and Pindyck, R. S. (1994). *Investment under uncertainty*.
- Doremus, H. (2007). Precaution, science, and learning while doing in natural resource management, *Washington Law Review* **82**: 547–580.
- Ellsberg, D. (1961). Risk, ambiguity, and the savage axioms, *Quarterly Journal of Economics* **75**: 643–669.
- Epstein, L. G. and Schneider, M. (2003). Recursive multiple-priors, *Journal of Economic Theory* **113**(1): 1–31.
- Epstein, L. G. and Schneider, M. (2007). Learning under ambiguity, *The Review of Economic Studies* **74**(4): 1275–1303.
- Etner, J., Jeleva, M. and Tallon, J.-M. (2012). Decision theory under ambiguity, *Journal of Economic Surveys* **26**(2): 234–270.
- Fisher, A. C. and Hanemann, W. M. (1987). Quasi-option value: Some misconceptions dispelled, *Journal of Environmental Economics and Management* **14**(2): 183–190.
- Gardiner, S. M. (2010). Is 'arming the future' with geoengineering really the lesser evil? some doubts about the ethics of intentionally manipulating the climate system, *Climate Ethics: Essential Readings* pp. 284–314.
- Gardiner, S. M. (2011). Some early ethics of geoengineering the climate: a commentary on the values of the Royal Society report, *Environmental Values* **20**(2): 163–188.
- Gilboa, I. and Schmeidler, D. (1989). Maxmin expected utility with non-unique prior, *Journal of Mathematical Economics* **18**(2): 141–153.
- Goes, M., Tuana, N. and Keller, K. (2011). The economics (or lack thereof) of aerosol geoengineering, *Climatic Change* **109**(3-4): 719–744.



- Goeschl, T. and Perino, G. (2007). Innovation without magic bullets: Stock pollution and R&D sequences, *Journal of Environmental Economics and Management* **54**(2): 146–161.
- Goeschl, T., Heyen, D. and Moreno-Cruz, J. (2013). The intergenerational transfer of solar radiation management capabilities and atmospheric carbon stocks, *Environmental and Resource Economics* **56**(1): 85–104.
- Gollier, C. and Treich, N. (2003). Decision-making under scientific uncertainty: The economics of the precautionary principle, *Journal of Risk and Uncertainty* **27**(1): 77–103.
- Gollier, C., Jullien, B. and Treich, N. (2000). Scientific progress and irreversibility: an economic interpretation of the precautionary principle’, *Journal of Public Economics* **75**(2): 229–253.
- Graham, J. D. (2004). *The perils of the precautionary principle: lessons from the American and European experience*, Vol. 818, Heritage Foundation.
- Hale, B. (2012). The world that would have been: Moral hazard arguments against geoengineering, *Engineering the Climate: The Ethics of Solar Radiation Management*, Rowman and Littlefield, Lanham.
- Hall, B. H. and Helmers, C. (2013). Innovation and diffusion of clean/green technology: Can patent commons help?, *Journal of Environmental Economics and Management* **66**(1): 33–51.
- Hartzell-Nichols, L. (2012). Precaution and solar radiation management, *Ethics, Policy & Environment* **15**(2): 158–171.
- Henry, C. (1974). Investment decisions under uncertainty: The ‘irreversibility effect.’, *American Economic Review* **64**(6): 1006–1012.
- Heyen, D. and Wiesenfarth, B. R. (2015). Informativeness of experiments for MEU – a recursive definition, *Journal of Mathematical Economics*, forthcoming.
- IPCC (2013). Summary for policymakers, *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, Cambridge University Press.
- IPCC (2014). Summary for policymakers, *Climate Change 2014: Impacts, Adaptation, and Vulnerability. Part A: Global and Sectoral Aspects. Contribution of Working Group II to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, Cambridge University Press.
- Jasanoff, S. (2007). Technologies of humility, *Nature* **450**(7166): 33–33.
- Karp, L. and Zhang, J. (2006). Regulation with anticipated learning about environmental damages, *Journal of Environmental Economics and Management* **51**(3): 259–279.
- Keith, D. (2000). Geoengineering the climate: History and prospect, *Annual Review of Energy and the Environment* **25**(1): 245–284.
- Keith, D. (2013). *A Case for Climate Engineering*, MIT Press.
- Keith, D. W., Duren, R. and MacMartin, D. G. (2014). Field experiments on solar geoengineering: report of a workshop exploring a representative research portfolio, *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* **372**(2031): 20140175.
- Klepper, G. and Rickels, W. (2012). The real economics of climate engineering, *Economics Research International* **2012**: 1–20.
- Knight, F. H. (1921). *Risk, uncertainty and profit*, Hart, Schaffner and Marx.
- Kolstad, C. D. (1996). Learning and stock effects in environmental regulation: the case of greenhouse gas emissions, *Journal of environmental economics and management* **31**(1): 1–18.
- Kousky, C., Rostapshova, O., Toman, M. and Zeckhauser, R. (2009). Responding to threats of climate change mega-catastrophes, *RFF Discussion Paper*.
- Krutilla, J. V. (1967). Conservation reconsidered, *The American Economic Review* **57**(4): 777–786.

- Lemoine, D. M. and Traeger, C. P. (2012). Tipping points and ambiguity in the economics of climate change, *NBER Working Paper 18230*, National Bureau of Economic Research.
- Long, J. and Winickoff, D. (2010). Governing geoengineering research: Principles and process, *Governing* **1**(5): 60–62.
- Machina, M. J. (1989). Dynamic consistency and non-expected utility models of choice under uncertainty, *Journal of Economic Literature* **27**(4): 1622–1668.
- Manoussi, V. and Xepapadeas, A. (2014). Cooperation and competition in climate change policies: Mitigation and climate engineering when countries are asymmetric, *SSRN Scholarly Paper ID 2535720*, Social Science Research Network, Rochester, NY.
- Margolis, M. and Naevdal, E. (2007). Safe minimum standards in dynamic resource problems: Conditions for living on the edge of risk, *Environmental and Resource Economics* **40**(3): 401–423.
- Markusson, N., Ginn, F., Singh Ghaleigh, N. and Scott, V. (2014). 'In case of emergency press here': framing geoengineering as a response to dangerous climate change, *Wiley Interdisciplinary Reviews: Climate Change* **5**(2): 281–290.
- Maslin, M. and Austin, P. (2012). Uncertainty: Climate models at their limit?, *Nature* **486**(7402): 183–184.
- May, R. (2001). Risk and uncertainty, *Nature* **411**(6840): 891–891.
- May, R. M., Levin, S. A. and Sugihara, G. (2008). Complex systems: Ecology for bankers, *Nature* **451**(7181): 893–895.
- Meinshausen, M., Meinshausen, N., Hare, W., Raper, S. C. B., Frieler, K., Knutti, R., Frame, D. J. and Allen, M. R. (2009). Greenhouse-gas emission targets for limiting global warming to 2°C, *Nature* **458**(7242): 1158–1162.
- Millner, A., Dietz, S. and Heal, G. (2013). Scientific ambiguity and climate policy, *Environmental and Resource Economics* **55**(1): 21–46.
- Missfeldt, F. (1999). Game-theoretic modelling of transboundary pollution, *Journal of Economic Surveys* **13**(3): 287.
- Moreno-Cruz, J. (2010). Mitigation and the geoengineering threat, *Unpublished Manuscript*. Available at: <http://works.bepress.com/morenocruz/3>.
- Moreno-Cruz, J. B. and Keith, D. W. (2012). Climate policy under uncertainty: a case for solar geoengineering, *Climatic Change* **121**(3): 431–444.
- Morrow, D. R. (2014). Ethical aspects of the mitigation obstruction argument against climate engineering research, *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **372**(2031): 20140062.
- Naevdal, E. and Oppenheimer, M. (2007). The economics of the thermohaline circulation—a problem with multiple thresholds of unknown locations, *Resource and Energy Economics* **29**(4): 262–283.
- Nordhaus, W. D. (1994). *Managing the global commons: the economics of climate change*, Vol. 31, MIT Press.
- Nordhaus, W. D. (2011). The economics of tail events with an application to climate change, *Review of Environmental Economics and Policy* **5**(2): 240–257.
- Parker, A. (2014). Governing solar geoengineering research as it leaves the laboratory, *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **372**(2031): 20140173.
- Parson, E. A. and Keith, D. W. (2013). End the deadlock on governance of geoengineering research, *Science* **339**(6125): 1278–1279.
- Peterson, M. (2006). The precautionary principle is incoherent, *Risk Analysis* **26**(3): 595–601.
- Pielke, R. (2002). Better safe than sorry, *Nature* **419**(6906): 433–434.

- Pindyck, R. S. (2011). Fat tails, thin tails, and climate change policy, *Review of Environmental Economics and Policy* **5**(2): 258–274.
- Popp, D. (2010). Innovation and climate policy, *NBER Working Paper 15673*, National Bureau of Economic Research.
- Randall, A. (2011). *Risk and precaution*, Cambridge University Press.
- Reynolds, J. (2014). A critical examination of the climate engineering moral hazard and risk compensation concern, *Available at SSRN 2492708*.
- Riedel, F. (2009). Optimal stopping with multiple priors, *Econometrica* **77**(3): 857–908.
- Robock, A. (2008). 20 reasons why geoengineering may be a bad idea, *Bulletin of the Atomic Scientists* **64**(2): 14–18.
- Robock, A. (2012). Is geoengineering research ethical?, *Atmospheric Research* **17**(18): 19.
- Savage, L. (1972). *The foundations of statistics*, Dover Publications.
- Shepherd, J. (2009). *Geoengineering the climate: science, governance and uncertainty*, Royal Society.
- Stern, N. (2006). *The Stern review report on the economics of climate change*, Cambridge University Press.
- Sunstein, C. R. (2002). *Risk and reason: Safety, law, and the environment*, Cambridge University Press.
- Sunstein, C. R. (2005). *Laws of fear: Beyond the precautionary principle*, Cambridge University Press.
- Tol, R. S. J. (2003). Is the uncertainty about climate change too large for expected cost-benefit analysis?, *Climatic Change* **56**(3): 265–289.
- Treich, N., Rey, B. and Courbage, C. (2013). Prevention and precaution, *TSE Working Paper 445*.
- Vardas, G. and Xepapadeas, A. (2010). Model uncertainty, ambiguity and the precautionary principle: Implications for biodiversity management, *Environmental and Resource Economics* **45**(3): 379–404.
- Victor, D. G., Morgan, M. G., Apt, J., Steinbruner, J. and Ricke, K. (2009). The geoengineering option: A last resort against global warming?, *Foreign Affairs* **88**(2): 64–76.
- Wagner, U. J. (2001). The design of stable international environmental agreements: Economic theory and political economy, *Journal of Economic Surveys* **15**(3): 377–411.
- Weitzman, M. L. (2009). On modeling and interpreting the economics of catastrophic climate change, *Review of Economics and Statistics* **91**(1): 1–19.
- Weitzman, M. L. (2011). Fat-tailed uncertainty in the economics of catastrophic climate change, *Review of Environmental Economics and Policy* **5**(2): 275–292.
- Weitzman, M. L. (2012). A voting architecture for the governance of free-driver externalities, with application to geoengineering, *NBER Working Paper 18622*, National Bureau of Economic Research.
- Weitzman, M. L. (2013). A precautionary tale of uncertain tail fattening, *Environmental and Resource Economics* **55**(2): 159–173.
- Wiener, J. B., Hammit, J., Rogers, M. and Sand, P. (2013). *The reality of precaution: Comparing risk regulation in the United States and Europe*, RFF Press.
- Zander, J. (2010). *The Application of the Precautionary Principle in Practice: Comparative Dimensions*, Cambridge University Press.

# The Intergenerational Transfer of Solar Radiation Management Capabilities and Atmospheric Carbon Stocks <sup>\*†</sup>

Timo Goeschl      Daniel Heyen      Juan Moreno-Cruz

## Abstract

Solar radiation management (SRM) technologies are considered one of the likeliest forms of geoengineering. If developed, a future generation could deploy them to limit the damages caused by the atmospheric carbon stock inherited from the current generation, despite their negative side effects. Should the current generation develop these geoengineering capabilities for a future generation? And how would a decision to develop SRM impact on the current generation's abatement efforts? Natural scientists, ethicists, and other scholars argue that future generations could be more sanguine about the side effects of SRM deployment than the current generation. In this paper, we add economic rigor to this important debate on the intergenerational transfer of technological capabilities and pollution stocks. We identify three conjectures that constitute potentially rational courses of action for current society, including a ban on the development of SRM. However, the same premises that underpin these conjectures also allow for a novel possibility: If the development of SRM capabilities is sufficiently cheap, the current generation may for reasons of intergenerational strategy decide not just to develop SRM technologies, but also to abate more than in the absence of SRM.

**Keywords:** Geoengineering; Climate Change; Intergenerational Issues; Strategic Behavior.

**JEL Codes:** D9; O33; Q54; Q55.

## 1 Introduction

The economics of climate change have been emphasizing for a long time that concerns for intergenerational equity imply that the current generation needs to reduce its greenhouse gas (GHG) emissions. Such a reduction will benefit future generations by decreasing the damages they are expected to suffer as a result of climate change (e.g. Stern 2006). Consequently, a central concern of the literature on optimal climate policy is to determine the optimal scale of abatement efforts vis-a-vis the business-as-usual scenario and thus

---

\*We are grateful to the associate editor and three anonymous referees whose advice and comments have greatly improved the paper. We also owe thanks to audiences at the University of Kiel, the University of Alberta, Heidelberg University, the AUROE 2011 workshop, the EAERE 2011 Annual Conference, and the AERE 2012 Summer Conference. Funding for this research was provided through the project "The Global Governance of Climate Engineering" at the Marsilius-Kolleg, Heidelberg University.

<sup>†</sup>An almost identical form of this article has been published 2013 in *Environmental and Resource Economics* 56(1), doi: 10.1007/s10640-013-9647-x.

to determine the optimal intertemporal trajectory of mitigation (Stern 2006; Nordhaus 2007; Tol 2001).

Recent developments have started to challenge the almost exclusive role of emissions reductions in climate policy. Apart from progress on the economics of adaptation to climate change (Agrawala and Fankhauser 2008), there is growing awareness among economists of the increasing plausibility of so-called “climate engineering” (CE). This term is a shorthand for deliberate large-scale interventions in the Earth’s climate system with the aim of limiting the damages of excessive atmospheric carbon stocks (Keith 2000). A variety of approaches run under the heading of “climate engineering” or “geoengineering”. Most of the attention focuses—due to feasibility, effectiveness, and cost—on solar radiation management (SRM) technologies.<sup>1</sup> These technologies enable an increase in the Earth’s albedo with respect to solar radiation, e.g. by the dispersal of reflective aerosols in the stratosphere. This allows the Earth to tolerate higher atmospheric carbon stocks while keeping global mean surface temperatures within acceptable bounds. Merely a theoretical possibility decades ago (Fleming 2010), it is now considered feasible that an ambitious R&D program could deliver effective SRM technologies within a few decades (Ridgwell et al. 2012).

A setting in which GHG emissions reductions are not the only option for addressing climate change damages, but compete with or are complemented by investment in geoengineering capabilities raises an entirely new set of questions. Some early assessments of climate engineering conclude that a future generation with access to such technologies would be able to handle the damages associated with stock pollutants at a surprisingly low direct variable cost (Klepper and Rickels 2012; Barrett 2008). However, these technologies are not understood to be “magic bullets”. The current consensus is that SRM interventions will involve significant indirect costs through side effects, such as changes in precipitation patterns, and that these side effects will themselves raise issues of how benefits and costs are distributed (Klepper and Rickels 2012; Moreno-Cruz and Keith 2012; Ricke et al. 2010). As a result, a decision-maker will have to carefully weigh the social benefits of avoided carbon damages and the social costs of SRM-induced side effects when making a decision on the deployment of geoengineering measures. Some papers have started to establish the conditions under which the deployment of SRM may or may not be meaningful (Bickel and Agrawal 2013; Goes et al. 2011; Moreno-Cruz and Keith 2012), to consider the impacts of a geoengineering capability on international negotiations (Barrett 2008; Moreno-Cruz 2010) and how to regulate deployment (Barrett 2008; Victor 2008).

The starting point of this paper is the observation that at the present time, the technological capabilities for a deployable SRM system do not exist. These capabilities

---

<sup>1</sup>The alternative to SRM technologies is geoengineering in the form of carbon dioxide removal (CDR), for example by means of direct air capture (Keith 2009), ocean fertilization (Lampitt et al. 2008), or enhanced weathering (Schuiling and Krijgsman 2006). See Klepper and Rickels (2012) for a comparative economic assessment.

would have to be created at a cost, and the costs of developing these capabilities are considered substantial (Klepper and Rickels 2012). It would fall to the current generation to devote some of its resources to investment into a R&D program that would make these capabilities available, perhaps 30 years from now, to a future generation (Schelling 1996). But is it a rational course of action for the current generation to develop these capabilities for a future generation that can then decide on its use? And how would a decision to develop such SRM capabilities impact on abatement efforts?

Despite agreement on the basic premises of the decision problem, the current literature offers at least five different conjectures regarding the rational course of action. One conjecture, termed “arming the future” argues in favor of SRM R&D since the technology would act as some type of insurance in the event that the sensitivity of the climate system with respect to the carbon stock turns out to be high (Moreno-Cruz and Keith 2012). This rationale of SRM R&D has been advanced by a number of commentators (Gardiner 2010; Schneider 1996) and enjoys considerable support. Some have even gone so far as to claim that, since SRM is an imperfect substitute for emissions abatement, a positive SRM R&D decision will not affect GHG emissions abatement (Bunzl 2009). This conjecture, which we dub “abatement invariance”, contrasts with a third conjecture, namely that investment in SRM R&D detracts from mitigation and will lead to less abatement effort (American Meteorological Society 2009; Shepherd 2009). This possibility of a reduction in mitigation effort by the current generation has been characterized as a manifestation of “moral hazard” that a rational course of action would avoid (Shepherd 2009; Hale 2012).

In other conjectures, the concerns lie with the behavior of the future generation. Future decision-makers may engage in climate engineering under circumstances and on a scale that the current generation would not hold to be optimal. At least three mechanisms are thought to explain why future generations may deviate from the judicious course of action. One is behavioral: Given the impression of an impending or immediate “climate emergency”, there is a belief that decision-makers will succumb to emotional factors that favor SRM deployment as a ‘quick fix’ (e.g. Bodansky 1996).<sup>2</sup> Similarly, while sunk cost of its development should on rational grounds be immaterial to the decision on whether to deploy SRM technology, there is ample historical experience that use of a capability is often regarded as necessary to justify large sunk cost (Gardiner 2011). The second mechanism lies in the political economy of science and technology: Researchers and industry funded to carry out research on certain technologies become interest groups in favor of technology deployment, creating a vocal and influential lobby for SRM use (Jamieson 1996).<sup>3</sup> The final mechanism is the potential for a genuine

---

<sup>2</sup>The American Meteorological Society, for instance, believes that “[...]geoengineering technologies, once developed, may enable short-sighted and unwise deployment decisions, with potentially serious unforeseen consequences.” (American Meteorological Society 2009).

<sup>3</sup>This concern is also present among the general public. Mercer et al. (2011) report very strong agreement on the statement that “Research will lead to a technology that will be used no matter what

change of preferences regarding SRM once its availability has become a fact of life for a future generation.<sup>4</sup> Irrespective of the mechanism, a bias in favor of SRM interventions, once those technological capabilities have been created, is the argument underpinning two additional conjectures advanced in the literature. One is that the rational course of action for today’s generation is to rule out research on geoengineering measures in order to prevent the future generation from acting against its own best interests (Keith et al. 2010).<sup>5</sup> The fifth conjecture is a different interpretation of the “moral hazard” argument by Bunzl (2009) and turns the logic on its head. It postulates that in the face of a pro-CE bias among the future generation, the rational course of action is to offer SRM as a backstop for a “climate emergency” and to abate less. The reason is that in a world in which the future generation uses SRM indiscriminately, the carbon stock damages imposed on the future generation are consistently lower than in the absence of SRM capabilities. As a result, the current generation would be treated more equitably if it allowed itself to emit more (Hale 2012).

Against this background of competing conjectures, this paper makes two contributions. The first is to extend the economic literature on the intergenerational aspects of climate policy in the direction of technology transfers. To our knowledge, the transfer of a mitigating technology, such as SRM, to a future generation that bears the damage costs of a stock pollutant has not been explicitly considered before. The second contribution is to demonstrate how simple economic analysis can provide a useful starting point for discriminating between the different conjectures on the ‘right’ course of action regarding a portfolio of both mitigation and SRM R&D activities. We provide these contributions by developing the simplest two-generations model that captures the four joint premises that appear to underpin all conjectures. These common elements are (i) that the current generation cares about the future generation sufficiently to be concerned about the stock damages of atmospheric carbon, (ii) that there may be a pro-SRM bias in the future, (iii) that both abatement and R&D involve a cost today and (iv) that both determine, in an environment characterized by uncertainty about the damages associated with atmospheric carbon, the future generation’s carbon stock and technological capabilities. Using this model, we study the subgame perfect behavior of the current generation in order to determine which of the conjectures above can arise, and if so, under which conditions.

Our results demonstrate that the possibility of transferring mitigating technologies

---

the public thinks.” among participants in a large-scale international survey of public perceptions of climate engineering.

<sup>4</sup>Already in one of the early reports by the US National Academy of Sciences on climate change, the authors raise the possibility that interest in CO<sub>2</sub> may generate or reinforce a lasting interest in national or international means of climate and weather modification; once generated, that interest may flourish independent of whatever is done about CO<sub>2</sub> (p. 470).

<sup>5</sup>Ruling out CE research could take the shape of the recent explicit ban on climate engineering R&D that was declared at the Conference of the Parties of the U.N. Convention on Biological Diversity in Nagoya in 2010.

into the future both re-emphasizes important economic insights on climate policy and raises important new issues. For example, even in the absence of a pro-SRM bias, the presence of an SRM option offsets current abatement. Far from constituting an instance of “moral hazard” (Bunzl 2009), this is simply a result of the partial substitutability between abatement and SRM that a current generation will rationally want to exploit. At the same time, the presence of a pro-SRM bias will constitute an important source of potentially powerful strategic distortions between generations that support some of the conjectures, but not all. A failure to carefully define those comparisons that are meaningful may be to blame for some of the confusion. As the analysis makes clear, rising R&D costs lead *ceteris paribus* always to ruling out SRM research for rational reasons. The same holds for an increasing pro-SRM bias. However, providing no SRM R&D will always be combined with higher abatement levels than if SRM was made available. At the same time, we find no support for the conjecture that abatement weakens due to a distortion between generations if R&D is undertaken. Quite the opposite: If SRM is made available, abatement will never fall below the non-distortion benchmark but increase the higher the distortion between generations. That abatement will increase even if SRM is provided is a new finding that has not been discussed before. The intuition is that an altruistic current generation can and will want to partially offset a pro-SRM bias among the future generation by providing more abatement today, thus reducing the incentives to deploy SRM in the future.

We proceed as follows. In the following section, we develop the simple two-generations model that captures the salient components of the SRM R&D debate in the most parsimonious and tractable fashion. In section 3, we provide a major step in the debate by defining the meaningful comparison point by the way of a benchmark. Section 4 derives the equilibria of the intergenerational game and establishes the main propositions. In section 5, we conclude.

## 2 The Model

### 2.1 The setting

Here, we develop a setting that pares an intergenerational decision problem involving both a stock pollutant and the development of an imperfect backstop technology down to its very essentials. Figure 1 provides a graphical representation of the setting, which features four periods. While the building blocks on stock pollution, abatement, and damages are taken from the mainstream literature, the novel aspects are the inclusion of the climate engineering option and the associated intergenerational issues so as to capture the common narrative elements described in the introduction.

The basic set-up consists of two non-overlapping generations, termed “current” and “future”. The current generation decides in period 1 on whether to invest in R&D for



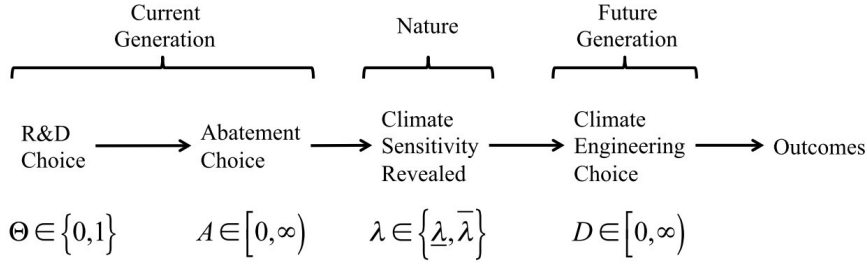


Figure 1: Timing of the game.

future SRM capabilities ( $\Theta = 1$ ) or not ( $\Theta = 0$ ) and chooses pollution abatement level  $A$  in period 2.<sup>6</sup> R&D of SRM involves technological as well as regulatory and institutional costs that will allow the future generation to use the climate engineering technology. These costs could be significant (Klepper and Rickels 2012). In period 3, nature resolves the key scientific uncertainty about climate change (Roe and Baker 2007), namely the climate sensitivity parameter  $\lambda$  that in turn determines the marginal damages associated with the future global carbon stock. In period 4, the future generation decides whether to deploy the SRM technology in order to counteract the damages  $T$  from the unabated pollution stock. If it deploys SRM ( $D > 0$ ), the future generation reduces damages from the carbon stock, but suffers environmental damages  $G$  associated with the use of SRM, e.g. in the form of disruptions of the hydrological cycle. If no SRM is used ( $D = 0$ ), future society faces the full temperature damages  $T$ .

As other tractable economic models of climate change, we specify the damages associated with the global carbon stock as caused by increased temperatures above historical averages. We employ the temperature damage function  $T$  by Moreno-Cruz and Keith (2012) of the form

$$T = \lambda^2 (R_0 - A)^2 . \quad (1)$$

The damage function consists of two arguments: The first argument is the squared carbon sensitivity of the climate to a doubling of  $\text{CO}_2$ ,  $\lambda$ . The second argument is the square of the net deviation in the carbon stock from historical levels ( $R_0 - A$ ), which consists of the business-as-usual increase in the carbon stock  $R_0$  minus the abatement effort  $A$ .<sup>7</sup> From the vantage point of the current generation, i.e. in periods 1 and 2, the carbon sensitivity is a random variable with two possible realizations: a carbon-sensitive value of  $\bar{\lambda} > 0$  with probability  $p$  and a carbon-insensitive value  $\underline{\lambda}$  ( $0 < \underline{\lambda} < \bar{\lambda}$ ) with probability  $(1 - p)$ . As expression (1) makes clear, abatement is productive as it reduces the expected value of pollution damages associated with climate change. Along with all moves  $\{\Theta, A\}$ , the pollution damage function (1) is common knowledge.

<sup>6</sup>The sequentiality of the decisions on  $\Theta$  and  $A$  is purely for ease of presentation. A simultaneous choice of  $\Theta$  and  $A$  leads to identical results.

<sup>7</sup>Temperature damages are assumed to be quadratic in temperature increases. The latter are assumed to be linear in the pollutant stock change

Abatement costs are usually assumed to be convex in abatement efforts. As a simple approximation, we model the abatement cost function of the quadratic type

$$X = \alpha A^2, \quad (2)$$

with increasing marginal abatement cost  $2\alpha A$ . The R&D process is modeled in a deterministic fashion, following Goeschl and Perino (2007): R&D for a functioning SRM technology requires payment of a fixed amount  $K$  in period 2 and delivers the climate engineering technology by period 4. The cost of not providing the SRM technology ( $\Theta = 0$ ) is zero.

The second option to counteract temperature damages, which can be combined with abatement, is the deployment of SRM. Like the abatement level  $A$ , we measure the amount  $D$  of deployed SRM in terms of compensated carbon stock (cf. Moreno-Cruz and Keith 2012) such that temperature damages read  $T = \lambda^2 (R_0 - A - D)^2$ . In that sense, abatement and SRM are equivalent regarding temperature damage compensation. They differ, however, in terms of timing and costs.

The use of SRM involves negligible direct costs (Barrett 2008), but causes collateral damages because of unintended negative impacts. These negative impacts comprise changes in the hydrological cycle and increase in air pollution, according to current assessments (Ricke et al. 2010; Kravitz et al. 2009). The general form of the damages  $G$  from SRM is a higher-order polynomial involving the volume of aerosol injected into the stratosphere ( $D$ ), the net carbon stock increase ( $R_0 - A$ ), and various particle characteristics (Bala et al. 2010; Caldeira and Wood 2008; Ricke et al. 2010). Recent simulations with general circulation models have shown, however, that the changes in temperature and precipitation are disproportionately driven by the linear term attached to the volume of SRM (Ban-Weiss and Caldeira 2010; Moreno-Cruz et al. 2012). A first approximation of the current generation's assessment of SRM damages would therefore specify an essentially linear relationship between damages and the amount  $D$  of SRM of the form

$$G = \rho D \quad (3)$$

with  $\rho$  denoting the constant marginal value of SRM damages. An immediately obvious economic viability condition for SRM is that  $\rho < 2\alpha R_0$ . Otherwise, the marginal cost of even the last unit of abatement is lower than the marginal damage of SRM, implying that SRM is never a competitive substitute for abatement.

The final component to make the setting reflect the current literature is a device that captures the possible presence of a bias of the future generation in favor of SRM deployment. The bias implies that unit for unit, the future generation values SRM damages generally less than the current generation for reasons of behavioral mechanisms, political economy, or genuine preference shifts. As a shorthand for this divergence between current and future generations, we introduce a bias parameter  $\beta \in [0, 1]$  such that from

the future generation's vantage point, damages are  $\tilde{G} = (1 - \beta)\rho D$ . If there is no bias ( $\beta = 0$ ), then the current and the future generation have an identical relative valuation between temperature damages and SRM damages.

## 2.2 Objectives and equilibrium concept

Against this background, we now define the objective functions, payoffs, and strategies of both generations. The future generation's objective is to minimize the sum of pollution damages and SRM damages, taking the choices of the current generation on abatement and R&D and the realization of climate sensitivity  $\lambda$  as given. Given R&D on SRM, the future generation can determine its optimal volume of SRM,  $D^*$ , that is how much of the inherited carbon stock should be compensated for by means of SRM. The optimal volume of SRM is a function of the abatement level  $A$ , the realized climate sensitivity  $\lambda$  and the bias parameter  $\beta$ . The future generation's problem is therefore

$$\min_{D \in [0, R_0 - A]} \left\{ \lambda^2 (R_0 - A - D)^2 + (1 - \beta)\rho D \right\}. \quad (4)$$

It is both intuitively obvious and clear from the first-order conditions that use of SRM requires that  $(1 - \beta)\rho < 2\lambda^2(R_0 - A)$ , i.e. the marginal damage of the first unit of SRM must be below the temperature damages thus avoided. Otherwise, the corner solution  $D^* = 0$  is optimal. For  $D^* > 0$ , the optimal amount is given by  $D^*(A, \lambda, \beta) = R_0 - A - \frac{(1-\beta)\rho}{2\lambda^2}$ .

From the vantage point of the current generation, the combination of the two possible realizations of climate sensitivity  $\lambda$  and the conditions on corner and interior solutions give rise to three possible deployment profiles. One profile,  $D_{\text{always}}$ , implies use of SRM by the future generation irrespective of whether the climate turns out to be sensitive or insensitive to carbon. Another,  $D_{\text{never}}$ , features no use of SRM, even if the climate turns out to be sensitive to carbon. The third,  $D_{\text{cond}}$ , conditions the use of SRM on the climate sensitivity, using SRM for a sensitive outcome and not using SRM for an insensitive one. The profiles are defined by

$$D_{\text{always}} = \begin{cases} R_0 - A - \frac{(1-\beta)\rho}{2\lambda^2} & \text{if } \lambda = \bar{\lambda} \\ R_0 - A - \frac{(1-\beta)\rho}{2\lambda^2} & \text{if } \lambda = \underline{\lambda} \end{cases}, \quad D_{\text{cond}} = \begin{cases} R_0 - A - \frac{(1-\beta)\rho}{2\lambda^2} & \text{if } \lambda = \bar{\lambda} \\ 0 & \text{if } \lambda = \underline{\lambda} \end{cases}$$

$$D_{\text{never}} = \begin{cases} 0 & \text{if } \lambda = \bar{\lambda} \\ 0 & \text{if } \lambda = \underline{\lambda} \end{cases}. \quad (5)$$

Due to a decrease of perceived damages caused by SRM, an increase in the bias  $\beta$  increases the SRM amount deployed. For  $\beta = 1$ , which refers to future generation attributing no damages to the use of SRM, temperature damages will be fully compensated,  $T = 0$ .

A quick inspection of the deployment profiles shows that the current generation, given its belief about the future generation's  $\beta$ , determines which profile will be chosen in the future through its choice of abatement  $A$ . More specifically, there is a lower and an upper threshold on abatement, with  $\underline{A} = R_0 - \frac{(1-\beta)\rho}{2\lambda^2}$  and  $\bar{A} = R_0 - \frac{(1-\beta)\rho}{2\lambda^2}$  respectively, such that

$$D^* = \begin{cases} D_{\text{always}} & \text{if } A \leq \underline{A} \leq \bar{A} \\ D_{\text{cond}} & \text{if } \underline{A} \leq A \leq \bar{A} \\ D_{\text{never}} & \text{if } \underline{A} \leq \bar{A} \leq A. \end{cases} \quad (6)$$

Through sufficiently high abatement, therefore, the current generation can ensure that SRM will never be used. Conversely, by abating little, available SRM would always be used at some positive level, even if the climate turns out to be relatively insensitive to carbon stocks. Finally, abatement levels between the two critical thresholds give rise to SRM deployment that is conditional on the realization of carbon sensitivity of the climate. Here, SRM would only be used in the eventuality of a sensitive climate.<sup>8</sup>

The optimal choice of costly SRM R&D  $\Theta$  in period 1 and abatement  $A$  in period 2 constitute the essence of the current generation's problem. In its choice, the current generation takes into account both the costs to itself (in the form of R&D and abatement costs) and the costs borne by the future generation (in the form of damages). The cost minimization objective requires that current and future costs need to be made comparable through an appropriate discount factor  $\delta$ . The objective function is then given by

$$\begin{aligned} \min_{\Theta, A} C(\Theta, A | \beta) = \min_{\Theta, A} \left\{ \Theta K + \alpha A^2 + \right. \\ \left. \delta(1-p) (\lambda^2(R_0 - A - D(A, \lambda, \beta))^2 + \rho D(A, \lambda, \beta)) + \right. \\ \left. \delta p (\bar{\lambda}^2(R_0 - A - D(A, \bar{\lambda}, \beta))^2 + \rho D(A, \bar{\lambda}, \beta)) \right\}. \quad (7) \end{aligned}$$

For ease of exposition, we set  $\delta = 1$  (no discounting). Also note that the objective function of the current generation makes explicit reference to  $\beta$  since the bias determines profile and amount of SRM deployment. We thus analyze functions of the form  $C(\Theta, A | \beta)$  or, equivalently, the form  $C(\Theta, A | D)$  where  $D$  denotes a certain deployment profile.

The optimal abatement levels differ depending on whether there is R&D on SRM or not. If no SRM R&D is carried out,  $\Theta = 0$ , the future generation cannot deploy the technology. This simplifies the objective function since  $D = 0$  for all  $A$ ,  $\lambda$  and  $\beta$ . The

---

<sup>8</sup>Note that the deployment profiles  $D_{\text{always}}$  and  $D_{\text{cond}}$  become indistinguishable at the abatement level  $\underline{A}$ . The same holds for  $D_{\text{cond}}$  and  $D_{\text{never}}$  at  $\bar{A}$ .

associated optimal abatement level is

$$A_{\text{NoR\&D}} = \frac{p\bar{\lambda}^2 + (1-p)\underline{\lambda}^2}{\alpha + p\bar{\lambda}^2 + (1-p)\underline{\lambda}^2} R_0 > 0, \quad (8)$$

which is a fraction of the business-as-usual increase in the carbon stock  $R_0$ . In line with intuition, abatement in the absence of SRM increases with a higher probability of a carbon-sensitive climate ( $\frac{d}{dp}A_{\text{NoR\&D}} > 0$ ) and for higher levels of sensitivity ( $\frac{d}{d\lambda}A_{\text{NoR\&D}} > 0$ ,  $\frac{d}{d\lambda}A_{\text{NoR\&D}} > 0$ ). The higher marginal abatement costs (higher  $\alpha$ ), the lower the abatement level: While costless abatement ( $\alpha = 0$ ) leads to full compensation,  $A_{\text{NoR\&D}} = R_0$ , abatement goes down to zero if marginal abatement costs tend to infinity.

If SRM R&D is carried out,  $\Theta = 1$ , the choice of the optimal  $A$  is more intricate because a change in  $A$  influences the future generation's discrete decision which of the three deployment profiles to choose. The objective function is therefore only piecewise differentiable and the optimal abatement level in each segment is not necessarily an interior solution. At the same time, note that the bias parameter  $\beta$  influences the choice of  $A$  only through its impact on the deployment profile; within one profile  $\beta$  has no impact on the abatement choice as  $\frac{d}{d\beta} \left( \frac{d}{dA} D(A, \lambda, \beta) \right) = 0$ . This is an interesting property that results from the dominance of the linear component in SRM damages and that will be exploited in section 4.

From an abstract point of view, the model set-up and objectives that capture the narratives on whether the current generation should develop SRM capabilities for the future generation define a sequential game with incomplete information. Its basic structure, in particular the technology transfer decision, is a variant of the trust game by Kreps (1990). However, the intergenerational decision problem here features two important differences: One is the availability of the second instrument in the form of abatement, the other the presence of exogenous uncertainty in the form of the random variable  $\lambda$ . The proper solution concept for determining the equilibrium played by the current and future generation is that of sub-game perfection (SP). The current generation, looking forward, employs backward induction to solve problem (7): By determining the optimal play of future generation ( $D^*|A, \lambda, \beta$ ) in period 4 contingent on current generation's choices in periods 1 and 2 and nature's move in period 3, it identifies its own optimal play  $\{\Theta^*, A^*\}$  in periods 1 and 2. However, it has to do so not knowing  $\lambda$  (see Figure 1). The equilibrium concept of SP will admit some combinations of abatement and R&D choices but not others, and thus impose elementary consistency checks on current society's rational course of action.

### 3 The Benchmark

Using the framework presented above, the purpose of the following two sections is to explore the consistency of different conjectures regarding the current generation's rational

course of action.

We proceed by constructing and establishing, as the first building block, a suitable benchmark case. Such a case is one in which (i) the bias of the future generation with respect to damages from SRM is set to zero by assumption and (ii) where equilibrium play features the current generation engaging in SRM R&D and the future generation adopting a conditional deployment profile. This choice of a benchmark has three benefits: The first is that in the comparisons in section 4 with situations involving a bias, the strategic distortions introduced by the bias will be clearly identifiable. The second benefit is that this benchmark case demonstrates the consistency of the “arming the future” conjecture for a given parameter set and thus establishes one of the five conjectures as a candidate for the rational course of action. Thirdly, the benchmark case also demonstrates that even in the absence of the bias, the rational course of action involves certain subtleties that the discussion on developing SRM technologies has so far failed to identify.

In the spirit of keeping notation clutter to a minimum, we assign values to some parameters. We restrict the level of abatement to the unit interval by setting  $R_0 = 1$ . Furthermore, marginal temperature and SRM damages will be regarded in terms of abatement costs by assuming  $\alpha = 1$  such that abatement costs are simply  $A^2$ . Finally, we set  $\underline{\lambda} = 1$  such that  $\bar{\lambda} > 1$  captures temperature damages for high carbon sensitivity, relative to a fixed baseline. These restrictions preserve the relevant degrees of freedom of the model and simplify the analysis substantially.

The benchmark case assumes the absence of a bias of the future generation with respect to SRM damages,  $\beta = 0$ .<sup>9</sup> To fix notation, equilibrium play in the benchmark case is characterized by current generations SRM R&D decision  $\hat{\Theta}$  and abatement level  $\hat{A}$  as well as the future generation’s deployment profile  $\hat{D}$ . For the conjecture of “arming the future” to survive the consistency check then requires that equilibrium play can give rise to a decision to conduct R&D on SRM,  $\hat{\Theta} = 1$ , and an abatement level  $\hat{A}$  that induces the future generation to choose the conditional deployment profile of the SRM technology:  $\hat{D}(\hat{A}, \beta = 0) = D_{\text{cond}}$ . Under these conditions, the current generation will make SRM capabilities available to the future generation alongside an optimal abatement effort to be determined, and the future generation will use SRM technologies only in case that the climate system turns out to be carbon sensitive. The future generation, in other words, will be “armed” against inadvertent outcomes in the climate system.

We begin the construction of the benchmark case with the convenient assumption that the abatement level associated with equilibrium play in the “arming the future” conjecture ought to be an interior solution. The interior solution to the current generation’s cost-minimization problem (7) assuming a conditional SRM deployment profile is

---

<sup>9</sup>In terms of interpretation,  $\beta = 0$  could mean that the future generation has the same preferences as current generation or that current generation is just ignorant of an existing asymmetry.

given by

$$A_{\text{cond}} = \frac{2(1-p) + p\rho}{2(2-p)} . \quad (9)$$

What parameter restrictions follow from designating  $A_{\text{cond}}$  as our solution  $\hat{A}$ ? These parameter restrictions can be summarized as restrictions on the marginal damage from SRM,  $\rho$ , and on the R&D cost of SRM,  $K$ . Note first that designating  $A_{\text{cond}}$  as  $\hat{A}$  requires that the abatement level  $A_{\text{cond}}$  chosen by the current generation actually gives rise to a conditional deployment profile by the future generation. Formally,

$$\underline{A}(\beta = 0) < A_{\text{cond}} < \bar{A}(\beta = 0) . \quad (10)$$

The level of abatement therefore has to lie in the middle segment defined in equation (6). Designating  $A_{\text{cond}}$  as  $\hat{A}$  therefore translates into a condition on the parameter  $\rho$ : The marginal damages of SRM  $\rho$  have to fulfill that  $1 < \rho < \frac{2\bar{\lambda}^2}{2+p(\bar{\lambda}^2-1)}$ . Together with the economic viability condition for SRM  $\rho < 2\alpha R_0 = 2$  (cf. section 2.1), this requires that the marginal damages of SRM fulfill

$$\rho \in \left( 1, \min \left\{ 2, \frac{2\bar{\lambda}^2}{2+p(\bar{\lambda}^2-1)} \right\} \right) \quad (11)$$

in order for  $A_{\text{cond}}$  to be feasible as  $\hat{A}$ . The intuition is straightforward: If  $\rho < 1$ , SRM damages are too small to give rise to conditional deployment of SRM and an unconditional deployment profile would be optimal instead. On the other hand, if  $\rho > \frac{2\bar{\lambda}^2}{2+p(\bar{\lambda}^2-1)}$  damages caused by SRM are too large and the optimal deployment profile would be to never use the technology. In this last case the amount of abatement would be larger than  $\bar{A}$ .<sup>10</sup>

The restrictions on  $\rho$  ensure that the deployment profile chosen will indeed be conditional on the realization of  $\lambda$ . Since these restriction cannot by themselves guarantee that there is not some other abatement level that involves lower cost than  $A_{\text{cond}}$ , restrictions on  $K$  are also required. To do so, we compare the total costs of the conditional profile, and associated abatement level  $A_{\text{cond}}$ , with the costs of the profile without SRM deployment and the unconditional deployment profile (see section 2.2), with their respective optimal abatement levels. First, the total costs of the conditional profile have to be lower than those of the unconditional deployment profile. That is,

$$C(\Theta = 1, A_{\text{cond}} | D_{\text{cond}}) < C(\Theta = 1, \underline{A} | D_{\text{always}}) , \quad (12)$$

---

<sup>10</sup>It should not come as a surprise that the upper bound increases with  $\bar{\lambda}$ ; that is, if the climate system is more climate sensitive, SRM damages can be larger and still allow for a conditional use profile. Less obvious is the fact that the upper bound decreases with  $p$ . An increase in  $p$  makes the upper limit more stringent because  $A_{\text{cond}}$  increases with  $p$ , but  $\bar{A}$  does not (whether the future generation deploys SRM does not depend on ex-ante probabilities). As  $p$  increases, the current generation has an incentive to increase  $A_{\text{cond}}$  and minimize now higher expected future costs which in turn reduces the need for SRM deployment.

where  $C(\Theta = 1, \underline{A} | D_{\text{always}})$  are the minimum costs under the indiscriminate use of SRM.<sup>11</sup> This is equivalent<sup>12</sup> to  $\frac{(\rho-1)^2}{2-p} > 0$ , which is always satisfied. Restrictions on  $K$  therefore need to come from the second comparison between the total costs of the conditional profile and those of the no SRM deployment profile, taking into account that conditional deployment requires investing in SRM R&D. For conditional deployment to involve lower cost requires that

$$C(\Theta = 1, A_{\text{cond}} | D_{\text{cond}}) < C(\Theta = 0, A_{\text{NoR\&D}} | D_{\text{never}}) . \quad (13)$$

This requirement translates into a limit on how costly SRM R&D can be:

$$K < \bar{K} := \frac{p(\rho(2 + p(\bar{\lambda}^2 - 1)) - 2\bar{\lambda}^2)^2}{4(2-p)\bar{\lambda}^2(2 + p(\bar{\lambda}^2 - 1))} . \quad (14)$$

Taking the above restrictions on  $\rho$  and  $K$  together, the parameter requirements for the benchmark case are thus fully characterized. More importantly still, there is no evidence of logical contradictions that would rule out equilibrium play according to the “arming the future” conjecture if there is no bias among the future generation: There is nothing inherently contradictory about the current generation providing both abatement and SRM R&D to the future generation such that this technology can be used as a backstop in the event of a carbon-sensitive climate.

Before exploring the impact of a bias regarding SRM use among the future generation, it is useful to point out two features of the benchmark equilibrium. One is the abatement level: The abatement level  $\hat{A} = A_{\text{cond}}$  is smaller than the abatement level which would be optimal without SRM,  $A_{\text{cond}} < A_{\text{NoR\&D}}$ . From an economic perspective, this is not surprising: Since abatement is costly, the ability to counteract the adverse effects of a pollutant implies that a higher pollution stock can be tolerated. Contrary to the characterization as “moral hazard” (see introduction), this reduction in abatement is therefore an optimal response. Whether the “moral hazard” argument has traction in a setting in which there is a bias among generations is a key question of the following section: it is conceivable that strategic considerations may lead to a suboptimal reduction on abatement, that is a level of abatement that is lower than  $A_{\text{cond}}$ . The second instructive feature of the benchmark are its comparative statics. The benchmark abatement level  $A_{\text{cond}}$  responds in a predictable fashion to an increasing likelihood of a “climate emergency” as well as higher marginal damages of SRM. In both cases, we

<sup>11</sup>The function  $C(\Theta, A | D_{\text{always}})$  is a convex quadratic function in  $A$  with first derivative  $-2(\rho - 1)$  at  $A = \underline{A}$ . This is negative due to the benchmark condition (4). This shows that  $\underline{A}$  is the minimizer of  $C(\Theta = 1, A | D_{\text{always}})$  in  $[0, \underline{A}]$ .

<sup>12</sup>In this section, the cost functions simplify due to  $\beta = 0$ . For some  $\beta \neq 0$ , expression (7) evaluated at the respective abatement levels yields  $C(\Theta = 1, A_{\text{cond}} | D_{\text{cond}}) = K + \frac{4(1-p) + p\rho(4-p\rho) - p\rho^2}{4(2-p)}(1-\beta^2)$ ,  $C(\Theta = 1, \underline{A} | D_{\text{always}}) = K + \frac{4-\rho(4-(2+p)\rho) - p\rho^2}{4} - (\rho-1)\rho\beta + \frac{(p+(2-p)\bar{\lambda}^2)\rho^2}{4\bar{\lambda}^2}\beta^2$  and  $C(\Theta = 0, A_{\text{NoR\&D}} | D_{\text{never}}) = 1 - \frac{1}{2+p(\bar{\lambda}^2-1)}$ .



see higher levels of abatement since  $\frac{d}{dp}A_{\text{cond}} > 0$  and  $\frac{d}{d\lambda}A_{\text{cond}} > 0$ . The threshold level of R&D costs also responds predictably to stronger climate damages (the willingness to equip future generation with the technology increases) and higher marginal damages of SRM (opposite effect) since  $\bar{K}$  reveal  $\frac{d}{d\lambda}\bar{K} > 0$  and  $\frac{d}{dp}\bar{K} < 0$  where the inequalities follow from (14). Its response to an increase in the probability of a “climate emergency” is less intuitive, however: the sign of  $\frac{d}{dp}\bar{K}$  is ambiguous. Increases in  $p$  at small levels of  $p$  tend to increase the threshold  $\bar{K}$  while the opposite is true at high levels of  $p$ . The intuition behind this result is that if the “climate emergency” is a low-probability event, the backstop characteristic of SRM dominates. An increase in the probability in this region strengthens the incentives to provide SRM irrespective of R&D costs. It would not make sense here to sacrifice too much in abatement cost but rather provide the option to counteract climate damages if these turn out to be high. If, however, a certain threshold for  $p$  is reached the “climate emergency” is more prevalent and, in order to minimize expected damages, the current generation draws larger levels of abatement. But with higher abatement levels, the role for SRM is reduced and with it, the incentives to pay for R&D. This intuition is reflected in  $A_{\text{NoR\&D}}$  being a concave function in  $p$  with large increases for small  $p$  while  $A_{\text{cond}}$  is convex in  $p$ , thus featuring higher increases for higher levels of the  $p$ . The surprising result here is that there is no monotonic relationship between the severity of the responsive climate and the willingness to develop the technology. While the incentives to provide the future generation with SRM increase in the marginal damages  $\bar{\lambda}$ , that is not necessarily the case in terms of increases in the likelihood  $p$  of the “climate emergency”.

## 4 Strategic distortions

Having established the benchmark equilibrium, we can now study how the rational course of action is impacted when the current generation anticipates a bias of the future generation when assessing SRM damages from geoengineering technologies relative to temperature damages from the carbon stock. In particular, we are interested in the effects on the current generation’s R&D decision and abatement choice  $\{\hat{\Theta}, \hat{A}\}$ . In the interest of space, we immediately proceed to a full characterization of the intergenerational carbon stock and technology transfer game in  $\beta$ - $K$ -space before providing the analysis that underpins this characterization. We then relate this characterization to the question of which of the five conjectures on current society’s rational course of action survive the consistency test. We find that only some conjectures pass this test. However, a sixth conjecture that has so far not been discussed in the literature emerges as a candidate for how current society could rationally approach this intergenerational decision.

#### 4.1 The full characterization

Figure 2 depicts the subgame perfect equilibria of the intergenerational carbon stock and technology transfer game in  $\beta$ - $K$ -space. The  $x$ -axis is defined by the bias parameter  $\beta$ , which denotes the degree to which the future generation discounts SRM damages vis-a-vis temperature damages relative to the current generation. The  $y$ -axis is defined by the cost parameter  $K$ , which denotes how much the current generation has to sacrifice to provide the future generation with SRM technologies. The upper bound of the  $x$ -axis,  $\beta = 1$ , refers to a level of bias at which the future generation attaches no damages to SRM deployment and therefore compensates temperature damages completely through geoengineering. The upper bound of the  $y$ -axis,  $\bar{K}$ , derives from expression (14) of the benchmark case and is the level of R&D costs at which the current generation is indifferent between providing the technology and not.

The benchmark case can be seen in Figure 2 as the line for which  $\beta = 0$ , the  $y$ -axis. The area to the right of the benchmark consists of four distinct zones. The north-east of the area is taken up by zone I in which no R&D is provided for the future generation, i.e.  $\Theta = 0$ . This zone is associated with a fixed abatement level  $A_{\text{NoR\&D}}$  irrespective of the bias  $\beta$  because the future generation's relative valuation of SRM damages is immaterial when no SRM technology is provided. On the  $y$ -axis, the boundary of this zone,  $K^{\text{Ban}}$ , naturally starts at  $\bar{K}$  for  $\beta = 0$  (see above) and decreases as  $\beta$  increases. As a result, zone I takes up an increasing amount of parameter space for higher levels of bias  $\beta$ . The other three zones, namely zones II, III, and IV, are all associated with the provision of R&D by the current generation, i.e.  $\Theta = 1$ , but differ with respect to the abatement levels chosen by the current and the deployment profiles chosen by the future generation. The boundaries between zones II and III and zones III and IV are each defined by a critical threshold condition on  $\beta$  and do not depend on  $K$ .

#### 4.2 Equilibria under SRM R&D ( $\Theta = 1$ )

With these basic features of Figure 2 established, we now proceed to explain how its geometry depends on the bias parameter  $\beta$  and the cost of R&D,  $K$ . To do so, assume for the moment that R&D is always carried out, i.e.  $\Theta = 1$ .

**Small bias (zone II)** Then, starting from the benchmark case ( $\beta = 0$ ), we can analyze what happens to abatement and deployment as  $\beta$  increases. Recall that in the benchmark case, equilibrium play consists of abatement of amount  $\hat{A} = A_{\text{cond}}$  by the current generation and of a conditional deployment profile  $D_{\text{cond}}$  by the future generation. Also recall from section 2.2 that as long as the conditional deployment profile is the future generation's best response, the current generation will find it optimal to stick to the benchmark level of abatement  $\hat{A}$  even at higher levels of  $\beta$ .

Expression (6) in section 2 defined the conditions under which sticking to the condi-




Figure 2: Equilibria in the intergenerational decision problem. The axes are the bias of the future generation  $\beta$  and the costs of SRM R&D  $K$ .

tional deployment profile is no longer a best response for the future generation. These conditions were set by two thresholds regarding the current generation's abatement,  $\underline{A}$  and  $\bar{A}$ . Observe now that both thresholds increase in  $\beta$ . As a result, the abatement level  $\hat{A}$  that gave rise to conditional deployment at low levels of  $\beta$  can now give rise to unconditional deployment as a best response. This is the case if  $\beta$  exceeds some critical value  $\underline{\beta}$ . The critical value is the level of  $\beta$  at which  $\underline{A}(\beta)$  exceeds  $\hat{A} = A_{\text{cond}}$  and is given by

$$\underline{\beta} := \frac{2(\rho - 1)}{(2 - p)\rho}. \quad (15)$$

Formally, the best-response deployment profile follows

$$D^*(\hat{A}, \beta) = \begin{cases} D_{\text{cond}} & 0 \leq \beta \leq \underline{\beta} \\ D_{\text{always}} & \underline{\beta} < \beta \leq 1. \end{cases} \quad (16)$$

A straightforward implication of (16) and the assumption of an interior equilibrium for the benchmark case, which is associated with  $\rho > 1$  in equation (11), is the existence of a non-empty interval of  $\beta \in [0, \underline{\beta}]$  for which the benchmark solution will be the equilibrium play despite the presence of a bias. For the geometry of Figure 2,  $\underline{\beta}$  defines the boundary between zones II and III and implies that zone II always exists and that equilibrium play within zone II follows the benchmark case.

**Large bias (zone IV)** Starting from the opposite end of the interval of bias, as announced in the beginning of this section, we now examine optimal play for the case of  $\beta = 1$  when  $\Theta = 1$ . Again, from expression (6) and (5) in section 2, it is clear that

the future generation will always deploy SRM in a way to fully offset the temperature damages. Deployment will therefore be  $D = 1 - A$  irrespective of  $\lambda$ . The current generation's best response to this is to choose abatement level  $A_{\text{always}} = \frac{\rho}{2}$  because  $A_{\text{always}}$  minimizes the total cost to the current generation in the face of unconditional deployment in the future, given that R&D is carried out. This abatement level is strictly greater than that which is optimal under conditional deployment, which reflects that it is cheaper for the current generation to counteract the future generation's "abuse" of SRM technologies through increased abatement, and hence less SRM deployment. Just as with the benchmark case, the question arises over what interval of  $\beta$  this equilibrium play persists, now going in the opposite direction (decreasing  $\beta$ ). In the benchmark case, we observed that as long as conditional deployment remained the best response of the future generation, the current generation did not adjust its abatement. The same logic applies here: As long as unconditional deployment remains the best response of the future generation, the current generation does not deviate from  $A_{\text{always}}$ . How low can  $\beta$  be for unconditional deployment to remain the best response to  $A_{\text{always}}$ ? Broadly analogously to the previous case, the critical value now is the level of  $\beta$  at which  $\underline{A}(\beta)$  falls below  $A_{\text{always}}$  and is given by

$$\bar{\beta} := \frac{2(\rho - 1)}{\rho} . \quad (17)$$

Casual inspection makes clear that there is always a non-empty interval of  $\beta \in [\bar{\beta}, 1]$  for which a combination of abatement at  $A_{\text{always}}$  by the current generation and unconditional deployment by the future generation will be the equilibrium play when  $\Theta = 1$ . For the geometry of Figure 2,  $\bar{\beta}$  defines the boundary between zones IV and III and implies that zone IV always exists.

**Medium bias - abatement as a strategic instrument (zone III)** The strategically subtlest case arises for degrees of bias between the two boundaries  $\underline{\beta}$  and  $\bar{\beta}$ . This is always a non-empty interval (except for the degenerate case that  $p = 1$ ) as is obvious from inspecting the boundary expressions. Zone III therefore exists. In order to understand equilibrium play in this interval, recall that  $\underline{\beta}$  denoted the threshold for values of  $\beta$  above which  $A_{\text{cond}}$  no longer induced conditional deployment as a best response and that  $\bar{\beta}$  denoted the threshold for values of  $\beta$  below which  $A_{\text{always}}$  no longer induced unconditional deployment as a best response. To find the amount for abatement that is both cost-minimizing and subgame-perfect in this interval, it is helpful to start from the threshold  $\underline{A}$  which separates those abatement levels that induce unconditional (to the left of  $\underline{A}$ ) from those that induce conditional deployment (to the right of  $\underline{A}$ ). It should be clear that for unconditional deployment, if subgame perfection was not a requirement, the cost-minimizing level of abatement  $A_{\text{always}}$  would lie to the right of  $\underline{A}$ . Subgame perfection in this  $\beta$ -zone, however, requires an abatement level not above  $\underline{A}$  and thus

smaller than  $A_{\text{always}}$ . At  $\underline{A}$ , marginal total costs are still decreasing. Therefore,  $\underline{A}$  must be the cost-minimizing subgame perfect choice of abatement if the current generation wants to induce unconditional deployment.

The case of the current generation wanting to induce conditional deployment is a mirror image of the case above: If subgame perfection was not a requirement, the cost-minimizing level of abatement  $A_{\text{cond}}$  would lie to the left of  $\underline{A}$ . However, subgame perfection in this zone requires an abatement level not below  $\underline{A}$  and thus greater than  $A_{\text{cond}}$ . At  $\underline{A}$ , marginal total costs are already increasing. Therefore,  $\underline{A}$  must be the cost-minimizing subgame perfect choice of abatement if the current generation wants to induce conditional deployment. Taken together,  $\underline{A}$  is the cost-minimizing subgame perfect choice of abatement in zone III. At  $\underline{A}$ , the conditional and unconditional deployment profiles become indistinguishable because the level of unconditional deployment in the case of carbon-insensitive climate becomes zero. Note that in contrast to the cost-minimizing subgame perfect abatement levels in zones II and IV which are independent of  $\beta$ , equilibrium play in zone III features abatement that increases in  $\beta$ .

We now have a complete picture of equilibrium abatement and SRM deployment in zones II, III, and IV, i.e. under the assumption that R&D is carried out by the current generation.<sup>13</sup> What remains to be shown is under what conditions the provision of R&D indeed constitutes a rational course of action by the current generation. Given the presence of strategic distortions in equilibrium play in zones II, III, and IV depending on the degree of the bias  $\beta$ , not providing SRM technologies in the first place may be the cost-minimizing choice. The key determinant of this decision is the cost of R&D. To this we turn now.

### 4.3 Equilibrium R&D decision ( $\Theta = 0$ vs. $\Theta = 1$ )

Recall from above that if no R&D is carried out ( $\Theta = 0$ ), then there is a unique abatement level  $A_{\text{NoR\&D}}$  chosen by the current generation. The future generation is forced into never deploying SRM. Therefore,  $D = 0$ . The expected total costs from this course of action are  $C(\Theta = 0, A_{\text{NoR\&D}} | D_{\text{never}})$ . For equilibrium play under  $\Theta = 1$  to be selected as the rational course of action, it needs to feature lower cost than  $C(\Theta = 0, A_{\text{NoR\&D}} | D_{\text{never}})$  despite involving R&D cost  $K$ . What is required therefore is to compare total expected costs in each of the zones II, III, and IV with the total expected costs from not carrying out R&D and to determine the condition on  $K$  for R&D still to be provided. The

---

<sup>13</sup>An interested reader might ask how the thresholds that separate zones II, III and IV,  $\underline{\beta}$  and  $\bar{\beta}$ , depend on model parameters. We have  $\frac{d}{d\rho}\underline{\beta} > 0$  and  $\frac{d}{d\rho}\bar{\beta} > 0$ . Increasing technology damages and likelihood of bad outcomes thus both expand the first region. The intuition is that both changes lead to a higher benchmark level  $\hat{A}$  which is then more robust to deviations in SRM damage assessment  $\beta$ . The second threshold,  $\bar{\beta}$ , depends only on  $\rho$ . The derivative is  $\frac{d}{d\rho}\bar{\beta} > 0$  and larger than  $\frac{d}{d\rho}\underline{\beta}$  because higher SRM damages reduce its desirability stronger if it is to be always deployed compared to conditional deployment. Thus, increasing damages of the technology make region IV smaller and enlarge both region II and region III.

condition that results for each zone each constitutes a segment of the boundary  $K^{\text{Ban}}$  in Figure 2 that separates the no-R&D zone I from the three R&D zones II, III, and IV.

We first compare the boundary between zone II and zone I. The costs we have to compare are  $C(\Theta = 1, A_{\text{cond}} | D_{\text{cond}})$  and  $C(\Theta = 0, A_{\text{NoR\&D}} | D_{\text{never}})$ . We are looking for a condition under which no R&D involves lower cost than providing R&D together with abatement of amount  $A_{\text{cond}}$ . Simple algebraic operations translate this comparison into a condition on  $K$  of the form

$$K > K_{II}^{\text{Ban}}(\beta) = \bar{K} - \frac{p\rho^2}{4\bar{\lambda}^2}\beta^2. \quad (18)$$

The boundary between zone III and zone I involves a comparison of costs  $C(\Theta = 1, \underline{A} | D_{\text{cond}})$ , which are by the definition of  $\underline{A}$  equal to  $C(\Theta = 1, \underline{A} | D_{\text{always}})$ , and  $C(\Theta = 0, A_{\text{NoR\&D}} | D_{\text{never}})$ . For no R&D to be cost-minimizing requires that

$$K > K_{III}^{\text{Ban}}(\beta) = \bar{K}_{III} - \frac{(p + (2-p)\bar{\lambda}^2)\rho^2}{4\bar{\lambda}^2}(\beta - \beta_{III})^2 \quad (19)$$

with  $\beta_{III} = \frac{\rho-1}{\rho} \frac{2\bar{\lambda}^2}{2\bar{\lambda}^2 - p(\bar{\lambda}^2 - 1)}$  and  $\bar{K}_{III} = \bar{K} - \frac{p(\rho-1)}{(2-p)(2\bar{\lambda}^2 - p(\bar{\lambda}^2 - 1))}$ . It is easy to show that  $\beta_{III} \in [0, \underline{\beta}]$ .

The boundary between zone IV and zone I involves a comparison of costs  $C(\Theta = 1, A_{\text{always}} | D_{\text{always}})$  and  $C(\Theta = 0, A_{\text{NoR\&D}} | D_{\text{never}})$ . For no R&D to be cost-minimizing requires that

$$K > K_{IV}^{\text{Ban}}(\beta) = \bar{K}_{IV} - \frac{p + (1-p)\bar{\lambda}^2}{4\bar{\lambda}^2}\beta^2 \quad (20)$$

where  $\bar{K}_{IV} = \bar{K} + \frac{(1-p)(\rho-1)^2}{2-p}$ .

Note that the segments derived above form a continuous boundary  $K^{\text{Ban}}$  since  $K_{II}^{\text{Ban}} = K_{III}^{\text{Ban}}$  and  $K_{III}^{\text{Ban}} = K_{IV}^{\text{Ban}}$  at the relevant thresholds  $\underline{\beta}$  and  $\bar{\beta}$ , respectively. Also note that  $K^{\text{Ban}}$  varies monotonously in the bias  $\beta$  since  $\frac{d}{d\beta} K_i^{\text{Ban}} < 0$  for all segments  $i$ . This confirms our earlier intuition that for the same level of R&D costs  $K$ , an increase in the bias  $\beta$  renders not carrying out R&D a relative more attractive course of action because the strategic distortions between the generations increase in  $\beta$ .<sup>14</sup>

<sup>14</sup>An interested reader might again ask how the boundary  $K^{\text{Ban}}$  depends on key model parameter. While the comparative statics are algebraically messy, we can build on the continuity and monotonicity of the boundary observed above using a simple trick. This trick involves examining the comparative statics of point  $K_{IV}^{\text{Ban}}(\beta = 1)$ , i.e. the point at which the boundary coincides with the upper bound on the bias  $\beta$ . By continuity and monotonicity, the comparative statics of this point are qualitatively the same as the comparative statics of the entire boundary. The comparative statics are that  $\frac{d}{dp} K_{IV}^{\text{Ban}}(\beta = 1) > 0$ ,  $\frac{d}{d\lambda} K_{IV}^{\text{Ban}}(\beta = 1) > 0$  and  $\frac{d}{d\rho} K_{IV}^{\text{Ban}}(\beta = 1) < 0$ . A greater likelihood or severity of a “climate emergency” increase the relative size of those zones that involve R&D into SRM while higher SRM damages have the opposite effect.

#### 4.4 Surviving conjectures

In this final substantive section of the paper, we relate the full characterization of the intergenerational carbon stock and technology transfer game back to the conjectures about the rational course of action for current society that appear in the current literature and that we reviewed in the introduction. Based on the preceding analysis, it is now possible to make statements about the degree to which different conjectures can be replicated in a simple model that captures the four common elements of (i) intergenerational altruism, (ii) possible pro-SRM bias, (iii) costly abatement and R&D, and (iv) intergenerational transmission of carbon stocks and technology under uncertainty about the climate system.

We identified five conjectures that share these common elements, yet differ markedly in their conclusions. One conjecture termed “arming the future” postulates that current society will find it rational to provide SRM technologies as a backstop for inadvertent climate outcomes while providing significant abatement. We were able to replicate this conjecture without problems and used it as the benchmark case for understanding the strategic distortions introduced by anticipating a bias among the future generation that implicitly favors the deployment of SRM. By contrast, the “abatement invariance” conjecture, namely that abatement should be unaffected by the decision to conduct R&D into SRM (Bunzl 2009), does not pass the consistency check of our model. The common elements that underpin all conjectures make it almost inevitable that abatement will change as soon as a decision in favor of SRM is taken since both abatement and SRM address the same problem, but feature different cost structures. We also fail to find support for the hypothesis that the anticipation of a large bias among the future generation that favors SRM deployment leads to the current generation slashing abatement. The reason is that if the current generation cares sufficiently for the future generation to take abatement action today, then it will also care sufficiently to take into account the negative side effects of large-scale geoengineering that would result from slashing abatement.

On the same grounds that challenge the last two conjectures, the model generates a much more upbeat view about the link between SRM R&D and abatement in the presence of a bias. As zone II in Figure 2 illustrates, if the bias is sufficiently weak ( $\beta < \underline{\beta}$ ), the benchmark abatement level  $\hat{A}$  will be maintained in equilibrium. In zone III with  $\underline{\beta} < \beta \leq \bar{\beta}$ , the bias is strong enough to render the initial abatement level  $\hat{A}$  inconsistent with conditional use by future generation. Anticipating that the technology will be used indiscriminately, current generation reacts by *increasing* abatement to the level  $\underline{A}(\beta)$  at which  $D_{\text{cond}}$  and  $D_{\text{always}}$  merge. The larger the bias  $\beta$ , the larger the necessary increase in abatement in order to induce the future generation into this specific use of the SRM technology. Finally, in zone IV ( $\beta > \bar{\beta}$ ), a further increase of  $\beta$  has no additional impact on the abatement level which remains at  $A_{\text{always}}$ , irrespective of  $\beta$ .

Interestingly,  $A_{\text{always}}$  can be even higher than the technology denial abatement level  $A_{\text{NoR\&D}}$ . We have

$$A_{\text{always}} > A_{\text{NoR\&D}} \quad \text{if and only if} \quad \rho > \frac{2 + 2p(\bar{\lambda}^2 - 1)}{2 + p(\bar{\lambda}^2 - 1)}. \quad (21)$$

This is, of course, only meaningful if this condition does not preclude technology provision equilibria with  $\Theta = 1$ . We can actually find equilibria in zone IV that feature “overabatement” relative to  $A_{\text{NoR\&D}}$ .<sup>15</sup> Figure 3 summarizes the abatement level over all equilibria.




Figure 3: Abatement in case of  $\Theta = 1$  as a function of the bias  $\beta$ . Whether  $A_{\text{always}}$  is larger or smaller than the optimal abatement level under no technology provision,  $A_{\text{NoR\&D}}$ , depends on model parameter.

In addition to this new equilibrium with SRM R&D provision and abatement increases, Figure 2 captures another consistent conjecture supported by the premises of the model. This is that current society will rationally choose not to engage in SRM R&D activities, i.e. a ban on SRM R&D is another course of action that our model confirms in a robust fashion (zone I). This course of action becomes particularly relevant for a large anticipated bias and for high development costs associated with SRM technologies.

We reserve a final comment on the conjecture that investment in SRM R&D will reduce abatement and thus give rise to “moral hazard”. As we point out in section 3, comparing abatement levels under positive and no R&D in the benchmark case shows that abatement in the presence of SRM R&D is smaller. However, rather than constituting a situation where an economic party is imposing a risk burden on some other

<sup>15</sup>Even though  $\rho > \frac{2+2p(\bar{\lambda}^2-1)}{2+p(\bar{\lambda}^2-1)}$  implies that  $K_{IV}^{\text{Ban}}(\beta = 1) < 0$  and thus a higher abatement level  $A_{\text{always}} > A_{\text{NoR\&D}}$  will not realize for  $\beta = 1$ , there are equilibria with  $\beta < 1$ . Take, for instance,  $p = \frac{1}{10}$ ,  $\bar{\lambda} = 2$  and  $\rho = 1.15$ . This leads to  $A_{\text{always}} > A_{\text{NoR\&D}}$ ; at the same time  $K_{IV}^{\text{Ban}}(\bar{\beta}) > 0$ , implying that the proposed equilibrium with  $\Theta = 1$  and  $A = A_{\text{always}} > A_{\text{NoR\&D}}$  exists in zone IV.



party without proper compensation, this is an efficient decision that reflects a rational readjustment of its abatement efforts by the current generation.

In sum, our model is able to successfully replicate three out of the five conjectures reviewed in the introduction. “Arming the future”, a R&D ban, and abatement reductions relative to a situation with no SRM R&D all constitute courses of action for current society that a model capturing the common elements among the conjectures can generate as consistent conclusions. Two conjectures, one postulating “abatement invariance” with respect to SRM R&D decisions and one postulating a drastic abatement reduction by current society, cannot be replicated and appear inconsistent with the basic premises in the literature. At a minimum, this means that auxiliary hypotheses and assumptions are necessary in order to demonstrate the consistency of these conjectures. At a maximum, the conclusion is that these conjectures are erroneous. In addition to the replication test of five existing conjectures, our model also shows that the same four common elements that underpin these conjectures give rise to a novel conjecture: It can be a rational course of action of current society to provide more abatement the higher the degree of anticipated bias. In fact, it is possible for this optimal abatement level to even exceed the level that society would rationally provide in the absence of SRM R&D.

## 5 Concluding discussion

If feasible, human interventions into the Earth’s climate system would represent a novel method for future generations to limit the damages of the atmospheric carbon stock that they will inherit from current society. Such interventions, running under the term of “geoengineering” or “climate engineering”, would create undesirable side effects of their own, but could conceivably be considered a realistic option. Solar radiation management (SRM) technologies are considered one of the likeliest forms of geoengineering to be developed and deployed.

R&D into SRM raises the possibility of passing on to the next generation not just abatement efforts in the form of a carbon stock that is below business-as-usual. In addition, or instead, the current generation could pass on a technology that can partially remedy the damages of excessive atmospheric carbon stocks. Natural scientists, philosophers, ethicists, and other scholars have started to develop several conjectures on how current society should decide on the right combination of SRM R&D and abatement efforts. By contrast, economists have so far not examined the intergenerational issues implicit in this technological possibility. This is despite the fact that economics has the potential to contribute substantially to the discussion thanks to its powerful conceptual tools for the analysis of intergenerational transfers. It is also despite the fact that the issue of intergenerational technology transfers to address the intergenerational transfer of stock pollution has so far attracted little attention in economics, in contrast to other types of intergenerational transfers.

The results of the present paper are a first step towards addressing the intergenerational issues of SRM R&D in the context of climate change. By developing a simple analytical model that formalizes several common elements in the wider debate about the correct course of action for the current generation in this intergenerational game, we harness the powers of game theoretic analysis for the purpose of understanding more about the problem. The comparison of the diverse conjectures about the rational course of action to develop SRM capabilities and the attempts to replicate their logic in a tractable model adds rigor to the debate and allows distinguishing between consistent conjectures and those that require either auxiliary hypotheses or correction. The same rigor allows us to identify solutions that have so far escaped attention, such as our finding that abatement may actually be higher when SRM capabilities are developed, and challenges loose argumentation, such as the claim regarding the presence of “moral hazard”.

For the economic debate, this paper represents a starting point for considering more systematically than before how to integrate the development of technological capabilities into intergenerational models. This integration is far from complete. It also adds to an emerging literature that examines the potential role of behavioral factors such as hyperbolic discounting, paternalism, and bounded rationality in the context of interactions across generations. The context of geoengineering provides a conducive and potentially consequential setting for more research of this type.

## References

- Agrawala, S. and Fankhauser, S. (2008). *Economic Aspects of Adaptation to Climate Change: Costs, Benefits and Policy Instruments*, OECD Publishing.
- American Meteorological Society (2009). Geoengineering the climate system: A policy statement of the American Meteorological Society, *Bulletin of the American Meteorological Society* **90**(9): 1369–1370.
- Bala, G., Caldeira, K. and Nemani, R. (2010). Fast versus slow response in climate change: implications for the global hydrological cycle, *Climate dynamics* **35**(2): 423–434.
- Ban-Weiss, G. and Caldeira, K. (2010). Geoengineering as an optimization problem, *Environmental Research Letters* **5**(3): 034009.
- Barrett, S. (2008). The incredible economics of geoengineering, *Environmental and Resource Economics* **39**(1): 45–54.
- Bickel, J. and Agrawal, S. (2013). Reexamining the economics of aerosol geoengineering, **119**(3-4): 993–1006.
- Bodansky, D. (1996). May we engineer the climate?, *Climatic Change* **33**(3): 309–321.
- Bunzl, M. (2009). Researching geoengineering: should not or could not?, *Environmental Research Letters* **4**(4): 045104.
- Caldeira, K. and Wood, L. (2008). Global and arctic climate engineering: Numerical model studies, *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **366**(1882): 4039–4056.
- Fleming, J. (2010). *Fixing the sky: the checkered history of weather and climate control*, Columbia University Press.
- Gardiner, S. M. (2010). Is ‘arming the future’ with geoengineering really the lesser evil? some doubts about the ethics of intentionally manipulating the climate system, *Climate Ethics: Essential Readings* pp. 284–314.

- Gardiner, S. M. (2011). Some early ethics of geoengineering the climate: a commentary on the values of the Royal Society report, *Environmental Values* **20**(2): 163–188.
- Goes, M., Tuana, N. and Keller, K. (2011). The economics (or lack thereof) of aerosol geoengineering, *Climatic Change* **109**(3-4): 719–744.
- Goeschl, T. and Perino, G. (2007). Innovation without magic bullets: Stock pollution and R&D sequences, *Journal of Environmental Economics and Management* **54**(2): 146–161.
- Hale, B. (2012). Getting the bad out: Remediation technologies and respect for others, in W. Kabasenche, M. O'Rourke and M. Slater (eds), *The Environment. Philosophy, Science, and Ethics*, MIT Press (MA).
- Jamieson, D. (1996). Ethics and intentional climate change, *Climatic Change* **33**(3): 323–336.
- Keith, D. (2000). Geoengineering the climate: History and prospect, *Annual Review of Energy and the Environment* **25**(1): 245–284.
- Keith, D. (2009). Why capture co2 from the atmosphere?, *Science* **325**(5948): 1654–1655.
- Keith, D., Parson, E. and Morgan, M. (2010). Research on global sun block needed now, *Nature* **463**(7280): 426–427.
- Klepper, G. and Rickels, W. (2012). The real economics of climate engineering, *Economics Research International* **2012**: 1–20.
- Kravitz, B., Robock, A., Oman, L., Stenchikov, G. and Marquardt, A. (2009). Sulfuric acid deposition from stratospheric geoengineering with sulfate aerosols. doi: 0.1029/2009JD011918.
- Kreps, D. (1990). *Game theory and economic modelling*, Clarendon Press Oxford.
- Lampitt, R. S., Achterberg, E., Anderson, T., Hughes, J., Iglesias-Rodriguez, M., Kelly-Gerreyn, B., Lucas, M., Popova, E., Sanders, R., Shepherd, J. et al. (2008). Ocean fertilization: a potential means of geoengineering?, *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **366**(1882): 3919–3945.
- Mercer, A., Keith, D. and Sharp, J. (2011). Public understanding of solar radiation management, *Environmental Research Letters* **6**(4): 044006.
- Moreno-Cruz, J. (2010). Mitigation and the geoengineering threat, *Unpublished Manuscript*. Available at: <http://works.bepress.com/morenocruz/3>.
- Moreno-Cruz, J. B. and Keith, D. W. (2012). Climate policy under uncertainty: a case for solar geoengineering, *Climatic Change* **121**(3): 431–444.
- Moreno-Cruz, J., Ricke, K. and Keith, D. (2012). A simple model to account for regional inequalities in the effectiveness of solar radiation management, *Climatic change* **110**(3): 649–668.
- Nordhaus, W. (2007). A review of the “stern review on the economics of climate change”, *Journal of Economic Literature* **45**(3): 686–702.
- Ricke, K., Morgan, M. and Allen, M. (2010). Regional climate response to solar-radiation management, *Nature Geoscience* **3**(8): 537–541.
- Ridgwell, A., Freeman, C. and Lampitt, R. (2012). Geoengineering: taking control of our planet's climate?, *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **370**(1974): 4163–4165.
- Roe, G. and Baker, M. (2007). Why is climate sensitivity so unpredictable?, *Science* **318**(5850): 629–632.
- Schelling, T. (1996). The economic diplomacy of geoengineering, *Climatic Change* **33**(3): 303–307.
- Schneider, S. H. (1996). Geoengineering: Could or should we do it?, *Climatic Change* **33**(3): 291–302.
- Schuiling, R. and Krijgsman, P. (2006). Enhanced weathering: An effective and cheap tool to sequester co 2, *Climatic Change* **74**(1): 349–354.
- Shepherd, J. (2009). *Geoengineering the climate: science, governance and uncertainty*, Royal Society.
- Stern, N. (2006). *The Stern review report on the economics of climate change*, Cambridge University Press.
- Tol, R. (2001). Equitable cost-benefit analysis of climate change policies, *Ecological Economics* **36**(1): 71–85.
- Victor, D. (2008). On the regulation of geoengineering, *Oxford Review of Economic Policy* **24**(2): 322–336.

# Strategic Conflicts on the Horizon: R&D Incentives for Environmental Technologies \*

Daniel Heyen

## Abstract

The innovation of environmental technologies is a key strategy against environmental problems, requiring however substantial R&D expenditures. This paper focuses on a specific mechanism that may influence countries' innovation incentives, thus impeding global optimal R&D outcomes. In this mechanism, the outlook of future conflicts about the technology deployment directly impacts on the willingness to sacrifice means for innovation. The possible deployment conflicts are free-riding on other countries' technology deployment and, recently discussed, so-called 'free-driving' (Weitzman 2012) that is expected to occur for technologies with low deployment cost and heterogeneous preferences like stratospheric aerosol injection, a 'solar radiation management' technique (Keith 2013). This paper develops a simple two stage model for analyzing how innovation incentives of two asymmetric countries are shaped by taking into account future deployment equilibria. The framework gives rise to rich findings, demonstrating that future deployment conflicts pull forward to the R&D decision. It can be shown that the outlook of free-riding behavior in the technology deployment weakens the incentives for innovation, underpinning a narrative by Popp (2010). More surprisingly and novel, the repercussions of free-driving on R&D incentives are sharply different. The paper demonstrates the possibility of *superoptimal* innovation incentives as well as countries' willingness to undertake counter-R&D in order to preclude a 'free-driver' from the option to deploy the technology.

**Keywords:** Environmental Technologies; Innovation; R&D Incentives; Public Goods; Strategic Conflicts; Climate Engineering.

**JEL Codes:** Q54; Q55; Q52.

## 1 Introduction

Technological innovation plays a key role in combating environmental problems. Important examples include CO<sub>2</sub> abatement technologies (Bosetti et al. 2009; Perino and Requate 2012; Poyago-Theotoky 2007) and the development of 'breakthrough technologies', for instance no-emission energy sources (Barrett 2006; Hoffert et al. 2002). A defining characteristic of novel technologies is that they require costly R&D to become available (Popp 2006; Harstad 2012; Golombek and Hoel 2011). In the absence of a

---

\*I thank conference participants at the World Congress of Environmental and Resource Economists (WCERE) 2014 in Istanbul, the AUROE Young Academics Workshop 2014 in Kiel, seminar participants at the FZU-ZEW Monthly Brown Bag Seminar, the Brown Bag Seminar at the Chair of Environmental Economics, and participants of the climate engineering summer schools at Harvard University 2013 and Heidelberg 2014 for helpful comments.

supranational entity undertaking or enforcing the development of these technologies, the main burden for R&D expenditures falls on countries. Domestic interests and strategic considerations, however, often stand in the way of implementing actions that would improve global well-being. The prospects for successful development of potentially welfare enhancing environmental technologies thus crucially depend on countries' incentives to engage in R&D.

In this R&D game, incentives can deviate from the global optimal for different reasons. A possible cause for an overall insufficient willingness to develop technologies is the public good nature of knowledge (Stiglitz 1999), leading to free-riding on other countries' R&D efforts (Popp 2010; Hall and Helmers 2013). Another reason for inefficiencies in the technological innovation process is that future deployment of the technologies, usually involving significant externalities, itself is prone to strategic considerations. Such outlook of suboptimal deployment patterns, in turn, may also shape the incentives for technology R&D. This mechanism has been explicitly raised by Popp (2010)<sup>1</sup> and is present, yet less prominently, in a couple of other contributions (Hall and Helmers 2013; Perino and Requate 2012; Barrett 2006; Hoel and De Zeeuw 2010). It is this impact of anticipated technology deployment profiles on R&D incentives that is the focus of the present inquiry.

The mechanism described by Popp (2010) is that R&D incentives are weakened due to the anticipation of free-riding behavior in future use of the technology. But free-riding is not the only deviation from global optimal deployment that can occur with environmental technologies. To explore the possibility of other deployment conflicts, it is helpful to focus on a specific set of environmental technologies. 'Climate engineering' (CE), or 'geoengineering', is the "deliberate large-scale manipulation of the planetary environment to counteract anthropogenic climate change" (Shepherd 2009). The main categories are so-called *carbon dioxide removal* (CDR) techniques that aim at reducing the stock of greenhouse gases in the atmosphere, for instance by removing CO<sub>2</sub> from the ambient air by means of chemicals ('direct air capture', see Keith et al. 2006), and *solar radiation management* (SRM) techniques that would alter the earth's radiation balance, for instance by the release of sulfur particles ('stratospheric aerosol injection', see Caldeira et al. 2013).

CE technologies raise a set of new economical and political questions (Barrett 2008; Finus et al. 2013; Victor 2008). While CDR techniques like direct air capture may, in terms of the surrounding incentive structure, be very similar to the mitigation of greenhouse gases and thus prone to free-riding (Chen and Tavoni 2013), the low deployment costs together with the necessarily global effect of SRM are expected to induce a novel strategic conflict, 'free driving' (Weitzman 2012): If preferences for global cooling are different across countries, the country aiming at the strongest temperature reduction,

---

<sup>1</sup>"Thus, without appropriate policy interventions, the market for technologies that reduce emissions will be limited, reducing incentives to develop such technologies" (Popp 2010).

the 'free-driver', will use SRM as to maximize its own payoffs; due to the low deployment costs, the resulting amount of cooling may considerably exceed the level of cooling other countries prefer. The domination of the outcome by one single country jeopardizes the global benefits SRM could provide if deployed as to maximize global welfare.

The present paper asks about the repercussions of these strategic conflicts, free-riding and free-driving, on the incentives for technological innovation. For this purpose, it develops a simple game-theoretical framework in which two countries non-cooperatively play a threshold R&D game in the first period and, conditional on successful R&D, a deployment game for a transboundary environmental technology in the second period. The general set-up combines well-known building blocks from the literature: The two stage R&D game structure is borrowed from the industrial organization literature (D'Aspremont and Jacquemin 1988; Kamien et al. 1992), the threshold R&D structure stems from Barrett (2006), and the technology deployment game is a standard public good game (Barrett 1994; Finus and Rübbelke 2013; Diamantoudi and Sartzetakis 2006) in which the countries have heterogeneous preferences for the technology deployment (McGinty 2006).

CE technologies are a fitting illustration for this general framework. Different technologies are expected to bring about the different strategic conflicts free-riding and free-driving. Also, CE technologies are not developed yet and thus require sufficient R&D incentives to become available.

The paper makes three main contributions. The first contribution is to demonstrate that free-rider and free-driver equilibria can emerge in the same standard public good framework. In comparison to Weitzman (2012), the smooth (inverse U-shaped) benefit function is more tractable than a kinked benefit function. Also, the cost parameter in the present framework plays an important and informative role, while Weitzman (2012) abstracts from technology deployment costs. The second contribution is to provide a simple and tractable framework to analyze the impact of anticipated deployment equilibria on the willingness to undertake R&D. In particular, the framework is capable of underpinning the narrative of Popp (2010) of future free-riding weakening today's innovation incentives. The third contribution of this paper is to offer a toolkit for disentangling R&D incentives into a 'non-spillover technology' part and a pure 'technology interaction' part and thus enabling a deeper understanding of R&D.

The framework produces novel and partially surprising findings. The main message is that strategic conflicts looming on the horizon directly impact on RnD incentives. It is this mechanism at work that weakens R&D incentives due to the anticipation of free-riding. What is more to it, free-driving – itself being a non-standard strategic outcome – has striking implications for R&D: Not only are *superoptimal* innovation incentives possible for a low cost technology like stratospheric aerosol injection, it is also possible that a country is willing to undertake 'counter-R&D' in order to preclude a potential free-driver from the technological option. What makes the possible outcomes of the model

more rich and subtle is the finding that even in free-driver equilibria, if the divergence in preferences is not so extreme as to cause counter-R&D, the 'dominated' country benefits even more of the transboundary nature of the technology than the free-driver does.

The present study makes contributions to three strands of literature. The first strand of literature revolves around the implications of externalities on strategic interaction. This literature can be applied, as the present paper does, to technology deployment, but the most common field of application in environmental economics is mitigation of greenhouse gas emissions. This literature started from Barrett (1994) and mostly developed into the direction of International Environmental Agreements (for a survey see Wagner 2001). The present paper abstracts from any possible form of treaty, issue-linkage and governance. The topic of asymmetric countries, a crucial feature of the present study, has received attention in Barrett (2001) and McGinty (2006).

The second strand of literature this paper contributes to is environmental innovation. This environmental R&D can either be analyzed in a social planner framework (Goeschl and Perino 2007; Teubal 1978), with regard to an intergenerational dimension (Goeschl et al. 2013), or – most common – the international dimension. The latter literature is closely connected to the industrial organization literature on R&D games (Brander and Spencer 1983; D'Aspremont and Jacquemin 1988; Kamien et al. 1992; Cozzi 1999) and covers different forms of environmental innovation: R&D can be cost-reducing (Bosetti et al. 2009; Hall and Helmers 2013), pollution-reducing (Perino and Requate 2012; Poyago-Theotoky 2007) or, as in the present paper, technology enabling (Barrett 2006; Hoel and De Zeeuw 2010).

Finally, this paper adds to the literature on the economics of CE. Most of the existing literature is interested in the interplay of SRM and abatement efforts, focusing either on intergenerational heterogeneity (Goeschl et al. 2013) or heterogeneity across countries (Moreno-Cruz 2010; Manoussi and Xepapadeas 2014). For the sake of simplicity, the present paper abstracts from the interplay of SRM with abatement. It is closely connected to Weitzman (2012) who, for the first time, analyzed 'free-driving' behavior in SRM deployment. The new angle the present paper adds to Weitzman (2012) is to put the repercussions on R&D incentives center stage.

The organization of the paper is as follows. Section 2 develops the simple game-theoretical setting in which two asymmetric countries play a non-cooperative R&D and technology deployment game. Section 3 characterizes deployment equilibria for the case of symmetric countries and analyzes the resulting R&D incentives. These findings serve as a useful benchmark for the general case of asymmetric countries that is the focus of section 4, which is the main section of this paper. Section 5 demonstrates how R&D incentives can be disentangled into two effects that enhance the understanding of the complex findings of section 4. Finally, section 6 concludes.

## 2 The Model

Two countries play a non-cooperative R&D and technology deployment game. In stage one, countries simultaneously make their individual R&D investments for an environmental technology. The R&D game is a threshold public good game with perfect spillovers: If the sum of R&D contributions exceeds a (commonly known) threshold, the technology is in the second stage available to both countries; otherwise no country can use it. Conditional on successful R&D, countries in the second stage simultaneously choose their deployment levels of the technology. This technology deployment game is a standard public good game with perfect spillovers. Each country bears the full (quadratic) costs of its technology deployment; in contrast to that and as usual in public good games, the other country cannot be excluded from the (inverse U-shaped) benefits. But, due to heterogeneity in the benefit-function, a central feature of the model, this has a flip-side. Non-excludability also implies that a country cannot protect against undesired high technology deployment levels.

The set-up combines different well-established building blocks. The general two-period setting with R&D and successive deployment stage, here transboundary technology use instead of Cournot competition, is borrowed from seminal work in the industrial organization literature (D’Aspremont and Jacquemin 1988; Kamien et al. 1992). The design of the threshold R&D process for a new environmental technology, but without treaty formation, is taken from Barrett (2006).

This public good structure with asymmetric benefits is quite general and capable of embracing a variety of technologies with transboundary effects and heterogeneous preferences. A particular fitting set of technologies however are the above mentioned climate engineering (CE) technologies.<sup>2</sup> What makes CE a good working example is that, first, these technologies are not developed yet so that the relevant R&D question is whether they are made available – in contrast to cost-reducing R&D (for instance Hall and Helmers 2013). Second, CE deployment fits well into the public good structure: Deployment costs are borne by the deploying country alone while the effect of reduced GHG levels and temperatures is inevitably global. Finally, and central in this paper, countries differ in their assessment of an ‘optimal climate’ (Rosenzweig and Parry 1994; Porter et al. 2014; Manoussi and Xepapadeas 2014; Heyen et al. 2015); as a consequence, the assumption of heterogeneous preferences for the *public good*<sup>3</sup> is particular reasonable here.

The magnitude of deployment costs will crucially shape the outcome of the game. Direct air capture, a CE technology of the carbon dioxide removal (CDR) class, is a

---

<sup>2</sup>It should be noted that the simple framework deliberately abstracts from dissimilarities of CE technologies – how quickly they act and how big their unintended side-effects will be – that are crucial in other contexts.

<sup>3</sup>“A pure public good is a pure public good (more of it is better) for some people under some circumstances and a pure public bad (more of it is worse) for some other people under some other circumstances.” (Weitzman 2012)



suitable example for an environmental technology with rather high deployment costs; in contrast, stratospheric aerosol injection, the best-known proposal for a solar radiation management (SRM) technology, is a good example of a low deployment cost technology. The rest of this paper will constantly refer to these two CE technologies.

The solution concept of the model is standard subgame perfectness (SPNE). The natural tool to solve the game is thus by backward induction. In the following we will explain the two stages of the model in detail, starting with the technology deployment stage (2.1) and then turning to the R&D stage (2.2).

## 2.1 Technology deployment stage

In case of successful R&D, countries in the second period choose their technology level  $q_i \geq 0$  simultaneously and non-cooperatively. In terms of the working example CE, think of  $q_i$  as the reduction in global temperatures that is accomplished by either removing CO<sub>2</sub> from the atmosphere (direct air capture) or putting sulfur into the stratosphere (stratospheric aerosol injection). The cost and benefit structure is of the quadratic-quadratic type (Barrett 1994; Finus and Rübbelke 2013; Diamantoudi and Sartzetakis 2006). The cost function, the same for both countries, is

$$C(q_i) = \frac{c}{2} q_i^2 \quad , \quad i = 1, 2 \quad (1)$$

with  $c > 0$ . Costs only depend on the private contribution  $q_i$ . In contrast, benefits feature the usual public good structure with perfect spillovers, such that benefits are a function of the total technology level  $Q = q_1 + q_2$ ,

$$B_i(Q) = b \left( a_i Q - \frac{1}{2} Q^2 \right) \quad (2)$$

with  $b > 0$  and  $a_i > 0$ . The marginal benefits  $dB_i/dQ$  vanish at  $Q = a_i$  which justifies calling  $a_i$  *country  $i$ 's preferred technology level*. A central component of the model is to allow for  $a_1 \neq a_2$  and hence heterogeneous technology preferences. Without loss of generality suppose  $a_1 \leq a_2$ . Instead of  $a_1$  and  $a_2$ , it is often more meaningful to focus on the *mean technology optimum*  $\bar{a} = (a_1 + a_2)/2$  and the *preference asymmetry*  $a_\Delta = (a_2 - a_1)/2$ . The condition  $a_i > 0$  translates into the restriction  $a_\Delta < \bar{a}$ .

The countries choose their deployment levels non-cooperatively so that standard Nash equilibrium is the appropriate solution concept. Denote the Nash equilibrium of this game, to be determined in section 3 and 4, by  $(q_1^*, q_2^*)$ . The social optimal configuration is  $(q_1^{**}, q_2^{**})$ . The payoff of country  $i$  of this technology deployment game is denoted  $\pi_i(q_1, q_2) = B(q_1 + q_2) - C(q_i)$ .<sup>4</sup>

---

<sup>4</sup>The reason to consider only the second stage payoffs, and not total payoffs that would also take into account R&D expenditures, is the following: The main focus of this paper is on R&D incentives, and the second period payoffs crucially shape the willingness to undertake costly R&D.

## 2.2 R&D stage

In the first period, countries simultaneously choose their R&D levels  $r_i \geq 0$  (the possibility of non-negative R&D contributions will be discussed in section 4.3). As in Barrett (2006), R&D is a threshold public good game with perfect knowledge spillovers. If  $r_1 + r_2 \geq \bar{R}$ , the technology is available to both countries. If  $r_1 + r_2 < \bar{R}$ , neither country has access to the technology. The threshold  $\bar{R}$  is common knowledge. Modeling R&D as a threshold process, and not as cost-reducing R&D (Hoel and De Zeeuw 2010; Hall and Helmers 2013; Bosetti et al. 2009) or emission-reducing (Poyago-Theotoky 2007; Perino and Requate 2012), is a realistic assumption in the context of technologies like direct air capture and stratospheric aerosol injection because CE technologies are at present merely theoretical concepts. The countries choose their individual R&D levels non-cooperatively so that, again, standard Nash equilibrium is the appropriate solution concept.

The analysis of R&D incentives crucially depends on how much the countries are willing to sacrifice to have the environmental technology available.

**Definition 1.** *Country  $i$ 's willingness to pay for the technology (WTP) is  $R_i = \pi_i(q_1^*, q_2^*) - \pi_i(0, 0)$ , where  $(q_1^*, q_2^*)$  is the Nash equilibrium of the technology deployment game in the second period. The total willingness to pay is  $R = R_1 + R_2$ . Important for comparison is  $R^{**} = (\pi_1(q_1^{**}, q_2^{**}) - \pi_1(0, 0)) + (\pi_2(q_1^{**}, q_2^{**}) - \pi_2(0, 0))$ , the maximal amount society would be willing to pay for making the technology available.*

Well known (for instance Barrett 2013) is that the Nash equilibria of a threshold public good game do not suffer from underprovision. In our context:

**Lemma 1.** *The Nash equilibria of the R&D game are all combinations  $(r_1, r_2)$  with  $r_1 + r_2 = \bar{R}$  and  $r_i \leq \max(0, R_i)$ .*

The intuition is that every country is willing to fill up R&D investments up to the necessary threshold if the necessary contribution does not exceed its WTP  $R_i$ . The reason for the condition  $r_i \leq \max(0, R_i)$  in Lemma 1 is that  $R_1$  can be negative (see section 4.3). In this case, Nash equilibria involving zero R&D contributions by country 1,  $r_1 = 0$ , are still possible.

Lemma 1 demonstrates a further merit of the threshold R&D assumption. In general, due to the public good nature of knowledge, we would expect total R&D contributions in equilibrium to fall short of its optimal level (Stiglitz 1999; Popp 2010). At the same time, and being the central topic of this paper, we also expect strong implications of the anticipated strategic behavior in the deployment stage on R&D incentives. Disentangling both effects would be cumbersome. Due to the favorable equilibrium conditions of the threshold R&D game, it is clear that any effect on total R&D in this model can be fully ascribed to the anticipated strategic outcome of the technology deployment game.

In particular, this paper is not about equilibrium selection. Which of the, in general infinitely many, R&D equilibria is more or less likely to become reality is not the ambition of this inquiry. Rather, the focus is on (i) a comparison of R&D incentives across countries and (ii) the question whether total R&D incentives are strong enough for successful technology development. Both questions are fully determined by analyzing  $R_i$  and  $R_1 + R_2$ .

### 3 The symmetric benchmark

This section is dedicated to the case in which both countries have homogeneous preferences for the level of technology deployment,  $a_1 = a_2$ , that is  $a_\Delta = 0$ . In a first step, section 3.1 derives the deployment equilibrium, reestablishing well-known free-riding results from the literature. Building on this, section 3.2 characterizes the R&D equilibria, demonstrating that the anticipated free-riding deployment profile weakens R&D incentives. Thus, the first contribution of this section is to pin down the narrative of Popp (2010), giving us confidence that the model is adequately designed for making statements about R&D incentives due to anticipated strategic conflicts. The second contribution of this section is to provide benchmark equilibria for the asymmetric case  $a_\Delta > 0$  that will be covered in section 4.

#### 3.1 Technology deployment

This section derives the Nash equilibrium  $(q_1^*, q_2^*)$  and contrasts it with the social optimal configuration  $(q_1^{**}, q_2^{**})$ . The main step for deriving the Nash equilibrium is to determine the reaction function. Given the other country's contribution  $q_{-i}$ , country  $i$  optimally chooses

$$q_i(q_{-i}) = \max \left\{ 0, \frac{b}{b+c}(\bar{a} - q_{-i}) \right\}, \quad i = 1, 2. \quad (3)$$

Recall that in the symmetric case  $\bar{a}$  is the global technology level that maximizes the countries' benefit function. If the other country's contribution does not exceed  $\bar{a}$  (if it does,  $q_i = 0$  is the best reply), the optimal response is to deploy some fraction of the remaining amount  $\bar{a} - q_{-i}$  that would maximize country  $i$ 's benefits, and this fraction approaches unity as the deployment costs converge to zero. As usual, the condition for the Nash equilibrium is  $q_1(q_2^*) = q_1^*$  and  $q_2(q_1^*) = q_2^*$ .

**Proposition 1** (Deployment equilibrium in the symmetric benchmark,  $a_\Delta = 0$ ). *Let  $c > 0$ . The deployment equilibrium in the symmetric benchmark with  $a_1 = a_2 = \bar{a}$  is unique and has the following properties:*

- (i) *The contributions  $q_i^* = \frac{\bar{a}b}{c+2b}$  are positive and monotonically decreasing in  $c$  with  $\lim_{c \rightarrow \infty} q_i^* = 0$ .*

- (ii) The sum of contributions  $Q^* = q_1^* + q_2^*$  is smaller than the socially optimal amount  $Q^{**} = \frac{4ab}{c+4b}$ , and the fraction  $Q^*/Q^{**}$  decreases in  $c$ .
- (iii) The equilibrium  $(q_1^*, q_2^*)$  is not Pareto optimal. The social optimal configuration  $(q_1^{**}, q_2^{**})$  is Pareto optimal and a Pareto improvement to  $(q_1^*, q_2^*)$ .

*Proof.* See Appendix B. □

These findings are hardly surprising and, for abatement instead of technology deployment choices, widely found in the literature (e.g. Barrett 1994). The key properties of the strategic conflict surrounding the technology use is that free-riding on other countries' contribution exists, giving rise to suboptimal low deployment levels. The social optimal configuration would make both countries better off, but is not stable against unilateral deviations. The higher the deployment costs, the more severe is, in relative terms, the gap between social optimal and actually undertaken technology deployment. In that sense, lower cost technologies are not only beneficial because they boost total net benefits, but also because small deployment costs alleviate the free-riding problem.

Regarding our working example climate engineering technologies, the findings of Proposition 1 imply that, if countries are symmetric and hence regard the same global temperature as optimal, any strategic problem that we can expect, irrespective of the cost structure of the CE technology, is free-riding and thus underprovision of the technology. We would expect this conflict to be stronger for cost-intensive technologies like direct air capture and significantly attenuated for low cost technologies like stratospheric aerosol injection. We will see in section 4 that heterogeneity in technology preferences substantially changes this favorable picture of low deployment costs.

### 3.2 R&D

Based on the results from the previous section we can determine the countries' willingness to pay  $R_i$  and contrast the total WTP  $R$  with the amount  $R^{**}$  a global planner would be willing to sacrifice to make the technology available (cf. Definition 1).

**Proposition 2** (WTP for R&D in the symmetric benchmark). *The WTP in the symmetric benchmark has the following properties:*

- (i) The individual WTP  $R_1 = R_2$ , and hence also the total WTP  $R = R_1 + R_2$ , are positive and decreasing in  $c$  with  $\lim_{c \rightarrow \infty} R = 0$ .
- (ii) The total WTP  $R$  falls short of the social optimal amount  $R^{**}$ .

*Proof.* See Appendix B. □

Figure 1 gives a graphical illustration of Proposition 2. What is intuitive in light of the deployment equilibrium in Proposition 1 is the decrease of WTP in the costs




Figure 1: Comparison of the total WTP (solid line) and the social optimal amount (dashed line) as a function of the cost parameter  $c$  in the symmetric benchmark ( $a_{\Delta} = 0$ ). The total WTP  $R$  is the sum  $R_1 + R_2$  with  $R_1 = R_2$ . The parameter settings are  $\bar{a} = 2$  and  $b = 1$ .

parameter  $c$ : The higher the costs, the lower will be the levels deployed, and thus the less willing are countries to spend money to get the technology.

The interesting feature is that the total WTP  $R$  is lower than what a social planner would be ready to pay for the technology's availability,  $R^{**}$ . Whether this has implications for the success of technology R&D depends on the threshold  $\bar{R}$ . If  $\bar{R}$  is higher than  $R^{**}$ , then the discrepancy of  $R$  and  $R^{**}$  is inconsequential because R&D should not proceed anyway. Likewise, if  $\bar{R}$  is lower than  $R$ , the R&D incentives for the (social desirable) technology are strong enough. The divergence of  $R$  and  $R^{**}$  however implies the existence of threshold values  $\bar{R}$  for which development of the technology should proceed, but fails to do so due to weak incentives.

The reason why  $R$  falls short of  $R^{**}$ , to stress this again, is not the public good nature of R&D. The threshold assumption precludes underprovision due to R&D free-riding. Rather, the reason for  $R < R^{**}$  is the anticipated strategic conflict in the technology deployment stage. Foreseeing free-riding in technology deployment, countries' incentives to develop the technology are substantially weakened. The first deliverable of section 3 is thus to underpin the narrative of Popp (2010) in a rigorous framework. The second deliverable is that the findings of Proposition 1 and Proposition 2 serve as a benchmark for the asymmetric case.

## 4 Asymmetric countries

This central section extends the analysis of R&D incentives to asymmetric countries. In Section 4.1 it will become clear that the cost parameter  $c$  plays a crucial role in the analysis. High cost parameter values give rise to deployment and R&D equilibria very similar to the benchmark case. These 'free-riding' equilibria will be covered in Section 4.2. In contrast, and topic of section 4.3, low deployment costs substantially change the

strategic set-up, giving rise to 'free-driving' behavior (Weitzman 2012) with far-reaching implications for R&D incentives.

#### 4.1 Two types of deployment equilibria

The first step is again to look at the deployment stage and specifically the reaction functions. The counterpart of 3, now for non-vanishing preference asymmetry  $a_\Delta > 0$ , is

$$q_1(q_2) = \max \left\{ 0, \frac{b}{b+c}(\bar{a} - a_\Delta - q_2) \right\}, \quad q_2(q_1) = \max \left\{ 0, \frac{b}{b+c}(\bar{a} + a_\Delta - q_1) \right\}. \quad (4)$$

Recall that  $\bar{a} - a_\Delta = a_1$  and  $\bar{a} + a_\Delta = a_2$  are the preferred technology levels of country 1 and country 2, respectively.

**Definition 2.** *We call a Nash equilibrium  $(q_1^*, q_2^*)$  of the technology deployment game a free-driver equilibrium if the country with the lower preference for the technology does not deploy in equilibrium,  $q_1^* = 0$ . If both countries contribute positive amounts  $q_i^* > 0$ , we call  $(q_1^*, q_2^*)$  a free-rider equilibrium.*

In particular, the deployment equilibrium in the symmetric benchmark is of the free-rider type. Weitzman (2012) gives a similar definition for free-driving in an  $n$  country setting with a kinked benefit function and without deployment costs.<sup>5</sup>

As a direct consequence of the reaction functions in (4), the following Lemma presents the key role of the cost parameter  $c$  in determining the type of deployment equilibrium.

**Lemma 2.** *The technology deployment game has a unique Nash equilibrium if not  $c = 0$  and  $a_\Delta = 0$  at the same time. For  $a_\Delta > 0$ , the cost parameter  $\bar{c} := 2ba_\Delta/(\bar{a} - a_\Delta)$  separates the two different strategic outcomes. The equilibrium is of the free-rider type for  $c > \bar{c}$  and of the free-driver type for  $c < \bar{c}$ .*

*Proof.* See Appendix C. □

The separating cost parameter  $\bar{c}$  increases in the preference asymmetry  $a_\Delta$  and vanishes for  $a_\Delta = 0$ .

Figure 2 illustrates Lemma 2 by presenting the reaction functions from (4) for two different cost parameter. If  $c > \bar{c}$ , both countries in equilibrium deploy positive amounts of the technology; for  $c \leq \bar{c}$ , country 1 would actually prefer a negative deployment level to counteract the high deployment level of country 2; being restricted to non-negative

---

<sup>5</sup>Two comments on the free-driver definition are in order. First, the extension of Definition 2 to an  $n$  country setting would involve the choice whether to speak of a free-driver equilibrium when *at least one* country or *all but one* countries do not contribute in equilibrium. Second, the free-driver definition in its current form rests on the impossibility of negative deployment. In the context of Solar Radiation Management, the possibility of counter-geoengineering (Barrett et al. 2014) has been raised. In this case negative contributions would be possible. The analysis of this possibility is left for future research.

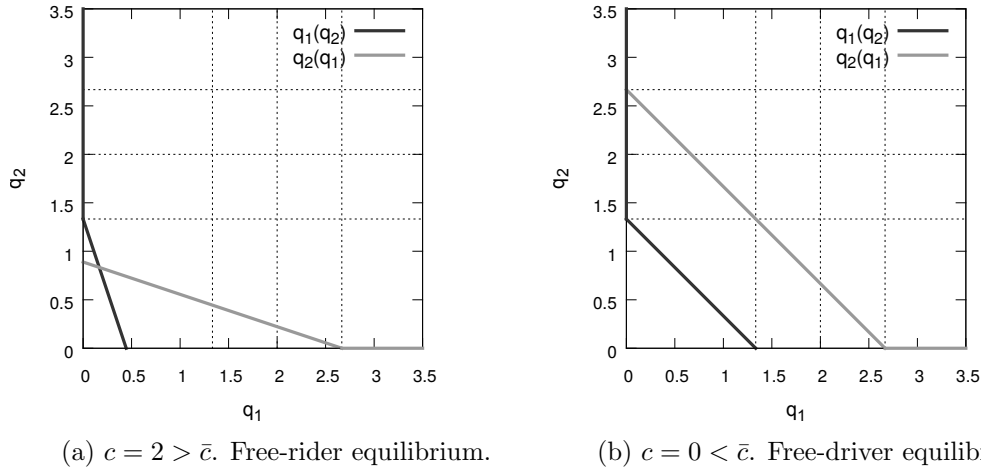


Figure 2: The reaction functions in the asymmetric case ( $a_\Delta > 0$ ) for two different cost parameter  $c$ . The parameter settings are  $b = 1$ ,  $\bar{a} = 2$ ,  $a_\Delta = 2/3$  so that  $\bar{c} = 1$ . With  $c = 2$  (left panel), the equilibrium  $(1/6, 5/6)$  is of the free-rider type. With  $c = 0$  (right panel), the equilibrium  $(0, 8/3)$  is a free-driver equilibrium.

levels country 1's best response is not to deploy. Table 1 summarizes the different outcomes.

Table 1: Nash outcomes of the deployment game in period 2 for asymmetric countries ( $a_\Delta > 0$ ) specifying technology deployment of country 1 ( $q_1^*$ ), country 2 ( $q_2^*$ ), total deployment ( $Q^*$ ) and social optimal total deployment ( $Q^{**}$ ).

Type	Free-rider	Free-driver
$c$	$c > \bar{c}$	$c \leq \bar{c}$
$q_1^*$	$\frac{b}{c+2b}\bar{a} - \frac{b}{c}a_\Delta$	0
$q_2^*$	$\frac{b}{c+2b}\bar{a} + \frac{b}{c}a_\Delta$	$\frac{b}{b+c}(\bar{a} + a_\Delta)$
$Q^*$	$\frac{2b}{c+2b}\bar{a}$	$\frac{b}{b+c}(\bar{a} + a_\Delta)$
$Q^{**}$	$\frac{4b}{c+4b}\bar{a}$	$\frac{4b}{c+4b}\bar{a}$

In terms of the working example CE technologies the above findings are highly relevant. In light of the sharp difference in deployment costs, we can expect CDR technologies like direct air capture to be prone to very different strategic incentives as the low cost SRM technology stratospheric aerosol injection. Table 1 also demonstrates that the framework outlined in section 2 is capable of reproducing the free-driver behavior, established in Weitzman (2012), in a standard smooth public good setting.

The following two subsections will elaborate and compare the deployment and resulting R&D characteristics for free-rider (4.2) and free-driver (4.3) equilibria.

## 4.2 Free-rider equilibria, $c > \bar{c}$

The first part of this section specifies the free-rider deployment equilibrium (the second column in Table 1) and is thus the counterpart of Proposition 1 for asymmetric countries

$a_\Delta > 0$ .

**Proposition 3** (Free-rider deployment equilibrium with asymmetric countries). *The unique free-rider equilibrium has the following properties:*

- (i) *The contributions  $q_1^*$  and  $q_2^*$  are positive with limits  $\lim_{c \rightarrow \bar{c}} q_1^* = 0$ ,  $\lim_{c \rightarrow \infty} q_1^* = 0$ ,  $\lim_{c \rightarrow \bar{c}} q_2^* = a_1$ , and  $\lim_{c \rightarrow \infty} q_2^* = 0$ . While  $q_2^*$  is monotonically decreasing in  $c$ ,  $q_1^*$  has a maximum in  $(\bar{c}, \infty)$ . While  $q_1^*$  increases in the asymmetry parameter  $a_\Delta$ ,  $q_2^*$  decreases in  $a_\Delta$ .*
- (ii) *The sum of contributions  $Q^*$  is smaller than the social optimal amount  $Q^{**}$ .*
- (iii) *The equilibrium  $(q_1^*, q_2^*)$  is not Pareto optimal. Local Pareto improvements consist of  $(q_1^* + \delta\alpha, q_2^* + \delta\beta)$ , where  $\alpha, \beta, \delta > 0$ . While the social optimal configuration  $(q_1^{**}, q_2^{**})$  is an improvement to  $(q_1^*, q_2^*)$  for country 2, this is in general not true for country 1: For  $a_\Delta > 0$ ,  $(q_1^{**}, q_2^{**})$  is not an improvement to  $(q_1^*, q_2^*)$  at least in a neighborhood of  $\bar{c}$ .*

*Proof.* See Appendix C.1. □

Country 2's technology deployment is very similar to the symmetric benchmark. That increases in  $a_\Delta$  drive up  $q_2^*$  is simply because this implies a higher technology preference,  $a_2 = \bar{a} + a_\Delta$ .

Quite different is the non-monotonic pattern of country 1. The reason that a decrease in costs close to the separating cost level  $\bar{c}$  reduces the contribution of country 1 is free-riding of the extremest form: At  $\bar{c}$ , country 2 contributes  $\bar{a} - a_\Delta = a_1$ , exactly country 1's preferred level; as a consequence, country 1 does not need to make own costly contributions. Not surprising, see part (ii), is that this leads overall to an underprovision of the technology.

What is interesting is the look at the Pareto improvements in (iii): Small increases in both contributions would make both countries better off, a simple consequence of the underprovision due to free-riding; the social optimal configuration however is often, and definitely when costs are low, *not* an improvement for country 1. The reason is again that, close to  $\bar{c}$ , country 1 is free-riding in an extreme way on country 2's contribution such that there is no room for better outcomes. As the appendix C.1 shows, a similar situation can also occur for high cost parameter values if the asymmetry  $a_\Delta$  is high enough.

What are the implications for R&D incentives? The following Proposition, the counterpart of Proposition 2 for asymmetric countries, gives the answer.

**Proposition 4** (WTP for R&D with asymmetric countries. Free-rider). *The WTP for free-rider technologies fulfills:*

- (i) *The individual WTP for R&D  $R_1$ ,  $R_2$ , and hence also the total WTP  $R$  are positive and decreasing in  $c$ . We have  $R_1 < R_2$  and  $\lim_{c \rightarrow \bar{c}} \frac{dR_1}{dc} = 0$ .*



- (ii) Increases in the preference asymmetry  $a_\Delta$  decrease  $R_1$  and increase  $R_2$ .
- (iii) The total WTP  $R$  falls short of the social optimal amount  $R^{**}$ . The difference between them increases in the preference asymmetry  $a_\Delta$ .

*Proof.* See Appendix C.1. □

Proposition 4 features strong similarities with the corresponding results in the symmetric benchmark, see Proposition 2. Most importantly, the total WTP falls short of the social optimum (part (iii)), implying the existence of constellations in which a beneficial technology is not developed for strategic reasons. The interpretation for  $R < R^{**}$  is, again, that the prospect of an underprovision of the technology directly reduces R&D incentives.

Besides the intuitive finding that an increase in  $a_\Delta$  drives up country 2's WTP and decreases country 1's WTP (part (ii)), a fact worth mentioning is that country 1's WTP  $R_1$  is strictly positive and flat at  $\bar{c}$ . Both are direct consequences of the fact that, at  $c = \bar{c}$ , the technological deployment profile is perfect from the country 1's viewpoint: No private costs,  $q_1^* = 0$ , but the optimal total technology deployment  $Q = a_1$  provided by country 2.

### 4.3 Free-driver equilibria, $c \leq \bar{c}$

This section demonstrates that free-driver equilibria are very different from their free-rider counterparts, both with regard to the deployment patterns, but also the resulting R&D incentives.

The first part of this section specifies characteristics of the free-driver deployment equilibrium (cf. third column in Table 1).

**Proposition 5** (Free-driver deployment equilibrium with asymmetric countries). *The unique free-driver equilibrium has the following properties:*

- (i) Country 1 does not contribute,  $q_1^* = 0$ , while  $q_2^*$  is positive, monotonically decreasing in  $c$ , and taking on the values  $\bar{a} + a_\Delta = a_2$  and  $\bar{a} - a_\Delta = a_1$  at the boundaries  $c = 0$  and  $c = \bar{c}$ , respectively.
- (ii) The sum of contributions  $Q^*$  is higher than the social optimal amount  $Q^{**}$  if and only if  $c < 4ba_\Delta / (3\bar{a} - a_\Delta) < \bar{c}$ .
- (iii) The equilibrium  $(q_1^*, q_2^*)$  is not Pareto optimal except for  $c = \bar{c}$ . For  $c < \bar{c}$ , local Pareto improvements consist of  $(q_1^* + \delta\alpha, q_2^* + \delta\beta)$ , where  $\alpha, \delta > 0$  and  $\beta < -\alpha$ . At  $c = 0$ , the social optimal configuration is an improvement for country 1 but not for country 2; at  $c = \bar{c}$ , it is the other way round. There always exists an inner range where  $(q_1^{**}, q_2^{**})$  is a Pareto improvement to  $(q_1^*, q_2^*)$ .

*Proof.* See Appendix C.2. □

By definition, in the entire free-driver region  $c \leq \bar{c}$  country 1 does not deploy the technology,  $q_1^* = 0$ . The total technology deployment level thus equals country 2's contribution  $q_2^*$ . It is not surprising to observe, and according to earlier findings, that this level goes up as costs  $c$  decrease. It also makes sense that  $q_2^* = a_2$  at  $c = 0$ : If there are no costs of deployment, country 2 chooses its preferred technology level.

Part (ii) shows that the total deployment level  $Q^*$  can exceed the social optimal amount. This finding, unknown in the standard public good literature, is another indication of the highly unusual and fascinating nature of free-driver equilibria. One has to keep in mind however that the comparison of the total values  $Q^*$  and  $Q^{**}$  is only one dimension of gauging the gap between equilibrium outcomes and social optimal reference points. Due to non-linear deployment costs, the distribution of deployment levels across countries always matters.

Part (iii) describes the possible Pareto improvements. The constellation at  $c = \bar{c}$  must be Pareto optimal as country 1, by definition, cannot do any better: The overall deployment level is at the optimal level  $Q^* = a_1$  without own costly contributions. Once  $c < \bar{c}$ , however, the Nash outcome is not Pareto optimal, and local Pareto improvements consist of country 1 contributing positive amounts, which already suffices to make country 2 better off, while country 2, in order to make it an improvement for country 1, would overproportionally decrease its deployment. Interestingly, the social optimal configuration  $(q_1^{**}, q_2^{**})$  can only be a Pareto improvement if costs are in a certain interior range; when costs get extreme, either to  $c = 0$  or  $c = \bar{c}$ , the Nash outcome favors one country too strongly as if the uniform social optimal deployment profile is attractive for both.

Connecting these findings with the respective part for free-rider equilibria of Proposition 3 gives a comprehensive picture. For country 2, the social optimal configuration is always an improvement except for very small  $c$ . For  $c > \bar{c}$  this is because it overcomes the free-rider problem, for levels below but close to  $\bar{c}$ , it saves deployment costs because  $(q_1^{**}, q_2^{**})$  involves positive contributions by country 1. For cost parameter values close to 0, however, the benefits of being able to afford technology levels close to optimal outweigh the costs of being the only contributor. For country 1, things are slightly more complex. Starting at the low cost end with costs close to zero, it is clear that the social optimal configuration would be an improvement; the reason is that here the free-driver behavior of country 2 is extreme, implying a significant divergence between preferred,  $Q = a_1$ , and actual deployment. Also clear is that for values close to, and on both sides, of the separating cost parameter  $\bar{c}$ , country 1 is better off under the Nash outcome because the overall deployment level is close to  $a_1$  and own provision costs are low, if not zero. Ambiguous however is the case of higher cost parameters: As explained in 4.2,  $(q_1^{**}, q_2^{**})$  is a Pareto improvement only if the asymmetry  $a_\Delta$  is not too high.

The main result in this section is concerned with the implications of free-driving behavior for R&D incentives.

**Proposition 6** (WTP for R&D with asymmetric countries. Free-driver). *The WTP for free-rider technologies fulfills:*

- (i) *Country 2's WTP  $R_2$  is positive and decreasing in the cost parameter  $c$ . Country 1's WTP  $R_1$  is increasing in  $c$  with  $\lim_{c \rightarrow \bar{c}} \frac{dR_1}{dc} = 0$ ;  $R_1$  gets negative for small  $c$  if  $a_\Delta > \bar{a}/3$ .*
- (ii) *Increases in the asymmetry  $a_\Delta$  decrease  $R_1$  and increase  $R_2$ .*
- (iii) *The total WTP  $R = \max\{R_1, 0\} + \max\{R_2, 0\}$  is positive and decreasing in  $c$ . If  $a_\Delta > (\sqrt{2} - 1)\bar{a}$ , the total WTP  $R$  at low values of  $c$  is higher than the the social optimum  $R^{**}$ . If this is the case, country 1's WTP  $R_1$  is necessarily negative.*

*Proof.* See Appendix C.2. □

Figure 3 gives a graphical illustration of Proposition 6.




Figure 3: The WTP of country 1 (dot-dashed line), country 2 (dot-dot-dashed line), total WTP (solid line) and social optimal WTP (dashed line) as a function of the cost parameter  $c$  for different asymmetry levels  $a_\Delta$ . The vertical dotted line is at  $\bar{c}$ , separating free-driver equilibria to the left from free-rider equilibria to the right. Note that  $R = \max(0, R_1) + \max(0, R_2)$ . The parameter settings are  $\bar{a} = 2$  and  $b = 1$  as before.

It is not surprising that country 2's WTP  $R_2$  is positive and decreasing in costs. Also, that  $a_\Delta$  drives  $R_1$  up while it decreases  $R_2$  has been found before and follows the same intuition here.

Quite unusual however is the behavior of country 1's WTP  $R_1$ . In contrast to the free-riding case, see Proposition 4, and sharply different from country 2, here  $R_1$  *decreases as costs go down*. This new finding however makes sense in light of the perfect constellation country 1 has at  $\bar{c}$ . Relative to that, lower  $c$  values drive country 2 into ever higher deployment levels and thus away from country 1's preferred overall technology level  $a_1$ .

Related and also interesting to note is that in the free-driving equilibrium, country 1, although not deploying the technology, is in general ready to sacrifice means to get

the technology. The reason for this positive WTP is that, despite excessive technology deployment by country 2, country 1 is often better off with the technology deployment pattern of country 2 than without any deployment.

This, however, can change for low cost levels if the technology preference across countries is high. Then, country 2's deployment strongly exceeds country 1's optimal deployment level so that country 1 would be better off without the technology. The negative WTP  $R_1$  in this case can be interpreted as the willingness to undertake counter-R&D. This demonstrates that low cost technologies like stratospheric aerosol injection that give rise to free-driver behavior not only suffer from strategic conflicts at the deployment stage in some future, but that the anticipation of those future conflicts may directly give rise to conflicts in the present.

Another peculiarity with free-driver equilibria is that R&D incentives can be *too strong*. The total WTP  $R$ , when defined based on non-negative R&D contributions  $r_i \geq 0$  (for the other case see Remark 1), can exceed  $R^{**}$  when the cost parameter  $c$  is low and the preference asymmetry between countries  $a_\Delta$  is high. In Figure 3, the asymmetry  $a_\Delta$  in the right figure,  $a_\Delta = 1$ , is large enough to feature  $R > R^{**}$  for small  $c$ , while this is not the case in the left figure. If  $R > R^{**}$ , the consequence is that technology R&D takes place (undertaken by country 2 alone) that should, from a societal point, not proceed.

**Remark 1.** *The analysis in Proposition 6 (iii) is correct if R&D contributions are non-negative,  $r_i \geq 0$ . Alternatively, one may consider the possibility of counter-R&D. For simplicity assume that one unit of counter-R&D just cancels one unit of R&D. In terms of possible equilibria, note that with  $R_1 < 0$  and  $R_2 > 0$  there is a unique R&D equilibrium, namely either  $(0, \bar{R})$  or  $(0, 0)$ . The relevant total WTP that decides about successful R&D is  $R = R_1 + R_2$ . It is easy to show (see C.2) that this measure is positive and always falls short of the social optimum; in other words the R&D incentives with counter-R&D are never too strong as in the case discussed before. But also the counter-R&D scenario gives rise to a peculiar outcome: For significant asymmetry between the countries,  $a_\Delta > \bar{a}/3$ , the total WTP  $R$  is non-monotonical in  $c$ .*

## 5 Disentangling Effects

The purpose of this final substantive section is to get a deeper understanding of the R&D incentives. For this purpose, divide the R&D process for a transboundary technology in a thought experiment into two steps. We write country  $i$ 's WTP  $R_i = \pi_i(q_1^*, q_2^*) - \pi_i(0, 0)$  accordingly as

$$R_1 = \underbrace{\pi_1(q_1^*, q_2^*) - \pi_1(q_1^{\text{Priv}}, 0)}_{=: R_1^{\text{Pub}}} + \underbrace{\pi_1(q_1^{\text{Priv}}, 0) - \pi_1(0, 0)}_{=: R_1^{\text{Priv}}} \quad (5)$$

and similar for  $R_2$ . The thought experiment consists of splitting R&D artificially into two processes. First, the development of a fully private technology so that the other country can be excluded from any deployment effects. The optimal technology deployment for such a private technology just balances private benefits and private costs with

$$q_i^{\text{Priv}} = \frac{b}{b+c} a_i \quad , \quad i = 1, 2 . \quad (6)$$

The WTP for such a no-spillover technology is then  $R_i^{\text{Priv}}$ .

The second process is then to transform the private technology to a transboundary technology with perfect spillovers that can be used by both countries. The willingness to open up a private technology in that way, which will result in the Nash contribution pattern  $(q_1^*, q_2^*)$ , is  $R_i^{\text{Pub}}$ .

The first Proposition analyzes the incentives to develop a no-spillover technology.

**Proposition 7** (WTP for private technology). *The incentives to develop a fully private technology are  $R_1^{\text{Priv}} = \frac{b^2}{2(b+c)}(\bar{a} - a_\Delta)^2$  and  $R_2^{\text{Priv}} = \frac{b^2}{2(b+c)}(\bar{a} + a_\Delta)^2$ . Thus:*

- (i) *For both countries, the WTP to develop a private technology is positive and decreasing in  $c$ .*
- (ii) *Increases in the asymmetry  $a_\Delta$  increase  $R_1^{\text{Priv}}$  and decreases  $R_2^{\text{Priv}}$ , and always  $R_1^{\text{Priv}} < R_2^{\text{Priv}}$ .*

*Proof.* Obvious. □

The findings are hardly surprising. In the absence of any interaction effects with other countries' deployment, the lower the costs and the higher the preferred technology level, the higher the WTP to develop such a private technology. Also, due to any interaction effects, the findings of Proposition 7 are valid for any cost level  $c$ .

The simple and obvious nature of the WTP for a private technology cannot explain the rich findings and differences between countries and cost ranges, see Propositions 4 and 6. Before we tackle the asymmetric case, a few words on the symmetric case are helpful to fully appreciate what is to come. In contrast to  $R_i^{\text{Priv}}$ , which is always monotonic in  $c$ , the effect  $R_i^{\text{Pub}}$  is not monotonic, even in the symmetric case. For  $a_\Delta = 0$  we get

$$R_i^{\text{Pub}} = \frac{\bar{a}^2 b^2 c (2c + 3b)}{2(c^3 + 5bc^2 + 8b^2c + 4b^3)} \quad (7)$$

with  $\lim_{c \rightarrow 0} R_i^{\text{Pub}} = \lim_{c \rightarrow \infty} R_i^{\text{Pub}} = 0$ . The reason that  $R_i^{\text{Pub}}$  is not monotonically increasing as costs go down is that opening up a low cost technology has only limited benefits: If costs are low, private deployment levels are already close to optimal; hence there is not much room for benefiting from other's deployment. It is helpful to keep this non-monotonicity in the symmetric case in mind when we now turn to the willingness to open up a technology for interaction for the general, asymmetric case.

**Proposition 8** (WTP for opening up a private technology). *The incentives to open up an existing private technology to a fully spillover technology are as follows:*

- (i) *Country 1. In the free-rider region,  $R_1^{\text{Pub}}$  is positive and non-monotonic in  $c$ . In the free-driver region it increases with  $c$  and is negative for  $c < \frac{4ba_{\Delta}^2}{\bar{a}^2 + 2a_{\Delta}\bar{a} - 3a_{\Delta}^2}$ .*
- (ii) *Country 2. In the free-rider region,  $R_2^{\text{Pub}}$  is positive and non-monotonic in  $c$  with  $\lim_{c \rightarrow \bar{c}} R_2^{\text{Pub}} = 0$ . In the free-driver region,  $R_2^{\text{Pub}} = 0$ .*
- (iii) *The WTP to open the private technology is always higher for country 1,  $R_1^{\text{Pub}} > R_2^{\text{Pub}}$ , except for small  $c$  when  $R_1^{\text{Pub}} < 0$ , see (i).*

*Proof.* See Appendix D. □

Figure 4 gives a graphical representation of both effects.




Figure 4: Disentangling two drivers of the WTP for country 1 (black lines) and country 2 (gray lines). The vertical dotted line is at  $\bar{c}$ , separating free-driver equilibria to the left from free-rider equilibria to the right. The functions depicted are the total WTP  $R_i$  (solid lines), the WTP to develop a no-spillover technology  $R_i^{\text{Priv}}$  (dashed lines), and the WTP to transform a private technology to a perfect spillover technology  $R_i^{\text{Pub}}$  (dot-dashed lines). The parameter settings are  $\bar{a} = 2$ ,  $a_{\Delta} = 2/3$  and  $b = 1$ .

Not surprising is the non-monotonic pattern of the willingness to open up a technology that we find for both countries in the free-rider region. This is essentially the same effect that we already isolated in the symmetric case, see (7).

An interesting feature is that  $R_2^{\text{Pub}} = 0$  over the entire free-driver region. The reason for that is that country 2, with country 1 in the Nash equilibrium not deploying at all, is already using the technology privately.

In terms of interpretation, (iii) shows that over the entire free-rider region  $R_1^{\text{Pub}} > R_2^{\text{Pub}}$ , and the gap increases with the preference asymmetry  $a_{\Delta}$  (see appendix D). This surprising finding can be interpreted as that country 1 is benefiting more from the public

good characteristic of the technology as country 2 does; country 1 is free-riding heavily on country 2's contribution.

Disentangling the different drivers also helps us to develop a more nuanced interpretation of the free-driving zone. At first sight, country 1 could be interpreted as a loser from the technology interaction because it is in the hand of the 'free-driver' and would actually prefer negative own contributions. However, if  $c$  is not too small,  $R_1^{\text{Pub}} > 0$  (and in particular larger than  $R_2^{\text{Pub}}$ ) because country 1 still benefits from country 2's technology deployment. So the story is more subtle. Even though we are in a 'free-driver' equilibrium in which country 1 does not deploy, country 1 is still free-riding on country 2's technology deployment. Only if the cost parameter  $c$  gets very small, country 1 is actually a loser from the spillover character of the technology. The reason that Weitzman (2012) only covers the latter case is that in his model  $c = 0$ . Due to the general public good structure with continuous cost parameter values  $c$ , the present model is capable of generating more complex and subtle effects.

## 6 Conclusions

Technologies constitute a central component in the portfolio of measures against environmental problems. An important example is climate engineering (CE), a set of environmental technologies that has recently received increasing attention. Novel environmental technologies, and CE in particular, however require substantial R&D expenditures to become available. In a decentralized world largely shaped by domestic interests, the incentives for innovation are crucial for successful R&D and thus deserve a thorough analysis.

The present paper focused on a specific problem surrounding R&D incentives. The anticipation of strategic conflicts in their future use, for instance free-riding, can be expected to have repercussions on the willingness to develop these technologies. In fact, the rigorous, yet parsimonious, framework developed in this paper has proved capable of underpinning a narrative of (Popp 2010) holding that the anticipation of free-riding weakens R&D incentives.

Starting from there, the paper also focused on a different strategic effect. In sharp contrast to free-riding, *free-driving* occurs when one country dominates the technology outcome (Weitzman 2012). This novel strategic effect materializes – and the present paper has provided further support for this in a standard smooth public good setting – if preferences for the technology deployment are heterogeneous and deployment costs for the technology are low. Stratospheric aerosol injection, and in general SRM, are expected to exhibit these characteristics and may thus be prone to free-driving outcomes.

This paper demonstrated that the outlook of free-driving has novel and rich impacts on innovation incentives. Most notably, the anticipated future deployment conflict may 'pull forward' to the R&D stage: The free-driver, keen to get the technology, may push

technology development even if this is against the global best. Accordingly, the other country, foreseeing the free-driver's extreme deployment level, may be willing to costly counteract the free-driver's innovation efforts, giving rise to an R&D conflict.

The present paper provides only the first step into the formal and rigorous analysis repercussions of how deployment conflicts impact on R&D incentives. There are a couple of valuable extensions that future research should envisage. The first possible extension is to generalize the two-country setting to  $n$  countries. It is far from clear how to generalize the definitions of 'free-riding' and 'free-driving' to the general case as it involves non-trivial and meaningful choices. This is particularly true in light of the subtleties surrounding free-driving that already emerged in the two country setting: The 'dominated' country, at first sight an obvious loser from the technology interaction, can for many parameter settings be expected to substantially benefit from the technological interaction. The question about winners and losers from the technology in a general setting is a fascinating research question.

A second possible line of research is to modify the framework to cost-reducing R&D. This type of innovation is pervasive in environmental economics (Bosetti et al. 2009; Hall and Helmers 2013), emphasizing the welfare-improving role of low cost technologies. The findings in the present paper suggest that free-driver technologies may be characterized by opposition against cost reducing R&D. The country with the lower preference for the technology, concerned of being worse-off under the free-driver's increasing deployment, may be willing to prevent cost-reducing innovation.

Another important extension of the existing framework is to consider governance structures and treaty formation. The present paper deliberately refrained from discussing these issues. Accordingly, the focus was on non-cooperative behavior and pure Nash outcomes as the relevant solution concept. The motivation to do so was mainly to develop a framework for making positive statements about the R&D incentives – in particular its deviations from global optimal – that we may expect without any governance regime in place. The present framework is thus the ideal starting point to see which, if any, governance structures or treaty options can help to overcome the strategic incentive problems.

Finally, future research should also incorporate an angle from the CE literature by focusing on the interplay of technology and abatement. In particular, Moreno-Cruz (2010) and Goeschl et al. (2013) demonstrated, in very different settings, that abatement can serve as a tool to attenuate strategic conflicts about the use of a technology by shifting their deployment incentives. A promising research question is whether abatement can play this role also in the context of innovation incentives for free-driver technologies.



## References

- Barrett, S. (1994). Self-enforcing international environmental agreements, *Oxford Economic Papers* **46**: 878–894.
- Barrett, S. (2001). International cooperation for sale, *European Economic Review* **45**(10): 1835–1850.
- Barrett, S. (2006). Climate treaties and "breakthrough" technologies, *The American Economic Review* **96**(2): 22–25.
- Barrett, S. (2008). The incredible economics of geoengineering, *Environmental and Resource Economics* **39**(1): 45–54.
- Barrett, S. (2013). Climate treaties and approaching catastrophes, *Journal of Environmental Economics and Management* **66**(2): 235–250.
- Barrett, S., Lenton, T. M., Millner, A., Tavoni, A., Carpenter, S., Anderies, J. M., Chapin III, F. S., Crpin, A.-S., Daily, G., Ehrlich, P., Folke, C., Galaz, V., Hughes, T., Kautsky, N., Lambin, E. F., Naylor, R., Nyborg, K., Polasky, S., Scheffer, M., Wilen, J., Xepapadeas, A. and de Zeeuw, A. (2014). Climate engineering reconsidered, *Nature Climate Change* **4**(7): 527–529.
- Bosetti, V., Carraro, C., Duval, R., Sgobbi, A. and Tavoni, M. (2009). The role of R&D and technology diffusion in climate change mitigation: new perspectives using the WITCH model, *FEEM Working Paper 14*.
- Brander, J. A. and Spencer, B. J. (1983). Strategic commitment with R&D: the symmetric case, *The Bell Journal of Economics* **14**(1): 225–235.
- Caldeira, K., Bala, G. and Cao, L. (2013). The science of geoengineering, *Annual Review of Earth and Planetary Sciences* **41**(1): 231–256.
- Chen, C. and Tavoni, M. (2013). Direct air capture of CO<sub>2</sub> and climate stabilization: A model based assessment, *Climatic Change* **118**(1): 59–72.
- Cozzi, G. (1999). R&D cooperation and growth, *Journal of Economic Theory* **86**(1): 17–49.
- D'Aspremont, C. and Jacquemin, A. (1988). Cooperative and noncooperative R & D in duopoly with spillovers, *The American Economic Review* **78**(5): 1133–1137.
- Diamantoudi, E. and Sartzetakis, E. S. (2006). Stable international environmental agreements: An analytical approach, *Journal of Public Economic Theory* **8**(2): 247–263.
- Finus, M. and Rübbelke, D. T. G. (2013). Public good provision and ancillary benefits: The case of climate agreements, *Environmental and Resource Economics* **56**(2): 211–226.
- Finus, M., Kotsogiannis, C. and McCorriston, S. (2013). The international dimension of climate change policy, *Environmental and Resource Economics* **56**(2): 151–160.
- Goeschl, T. and Perino, G. (2007). Innovation without magic bullets: Stock pollution and R&D sequences, *Journal of Environmental Economics and Management* **54**(2): 146–161.
- Goeschl, T., Heyen, D. and Moreno-Cruz, J. (2013). The intergenerational transfer of solar radiation management capabilities and atmospheric carbon stocks, *Environmental and Resource Economics* **56**(1): 85–104.
- Golombek, R. and Hoel, M. (2011). International cooperation on climate-friendly technologies, *Environmental and Resource Economics* **49**(4): 473–490.
- Hall, B. H. and Helmers, C. (2013). Innovation and diffusion of clean/green technology: Can patent commons help?, *Journal of Environmental Economics and Management* **66**(1): 33–51.
- Harstad, B. (2012). Climate contracts: A game of emissions, investments, negotiations, and renegotiations, *The Review of Economic Studies* **79**(4): 1527–1557.
- Heyen, D., Wiertz, T. and Irvine, P. (2015). Regional disparities in solar radiation management impacts: Limitations to simple assessments and the role of diverging preferences, *IASS Working Paper*.

- Hoel, M. and De Zeeuw, A. (2010). Can a focus on breakthrough technologies improve the performance of international environmental agreements?, *Environmental and Resource Economics* **47**(3): 395–406.
- Hoffert, M. I., Caldeira, K., Benford, G., Criswell, D. R., Green, C., Herzog, H., Jain, A. K., Kheshgi, H. S., Lackner, K. S., Lewis, J. S., Lightfoot, H. D., Manheimer, W., Mankins, J. C., Mauel, M. E., Perkins, L. J., Schlesinger, M. E., Volk, T. and Wigley, T. M. L. (2002). Advanced technology paths to global climate stability: Energy for a greenhouse planet, *Science* **298**(5595): 981–987.
- Kamien, M. I., Muller, E. and Zang, I. (1992). Research joint ventures and R&D cartels, *The American Economic Review* **82**(5): 1293–1306.
- Keith, D. (2013). *A Case for Climate Engineering*, MIT Press.
- Keith, D. W., Ha-Duong, M. and Stolaroff, J. K. (2006). Climate strategy with CO<sub>2</sub> capture from the air, *Climatic Change* **74**(1-3): 17–45.
- Manoussi, V. and Xepapadeas, A. (2014). Cooperation and competition in climate change policies: Mitigation and climate engineering when countries are asymmetric, *SSRN Scholarly Paper ID 2535720*, Social Science Research Network, Rochester, NY.
- McGinty, M. (2006). International environmental agreements among asymmetric nations, *Oxford Economic Papers*.
- Moreno-Cruz, J. (2010). Mitigation and the geoengineering threat, *Unpublished Manuscript*. Available at: <http://works.bepress.com/morenocruz/3>.
- Perino, G. and Requate, T. (2012). Does more stringent environmental regulation induce or reduce technology adoption? When the rate of technology adoption is inverted U-shaped, *Journal of Environmental Economics and Management* **64**(3): 456–467.
- Popp, D. (2006). R&D subsidies and climate policy: is there a free lunch?, *Climatic Change* **77**(3-4): 311–341.
- Popp, D. (2010). Innovation and climate policy, *NBER Working Paper 15673*, National Bureau of Economic Research.
- Porter, J. R., Xie, L., Challinor, A. J., Cochrane, K., Howden, S. M., Iqbal, M. M., Lobell, D. B. and Travasso, M. I. (2014). Food security and food production systems, *Climate Change 2014: Impacts, Adaptation, and Vulnerability. Part A: Global and Sectoral Aspects. Contribution of Working Group II to the Fifth Assessment Report of the Intergovernmental Panel of Climate Change* pp. 485–533.
- Poyago-Theotoky, J. (2007). The organization of R&D and environmental policy, *Journal of Economic Behavior & Organization* **62**(1): 63–75.
- Rosenzweig, C. and Parry, M. L. (1994). Potential impact of climate change on world food supply, *Nature* **367**(6459): 133–138.
- Shepherd, J. (2009). *Geoengineering the climate: science, governance and uncertainty*, Royal Society.
- Stiglitz, J. E. (1999). Knowledge as a global public good, *Global public goods* **1**(9): 308–326.
- Teubal, M. (1978). Threshold R&D levels in sectors of advanced technology, *European Economic Review* **7**(4): 395–402.
- Victor, D. (2008). On the regulation of geoengineering, *Oxford Review of Economic Policy* **24**(2): 322–336.
- Wagner, U. J. (2001). The design of stable international environmental agreements: Economic theory and political economy, *Journal of Economic Surveys* **15**(3): 377–411.
- Weitzman, M. L. (2012). A voting architecture for the governance of free-driver externalities, with application to geoengineering, *NBER Working Paper 18622*, National Bureau of Economic Research.

## Appendix A Toolkit for proving inequalities

There is often the need to prove inequalities of the type  $P < 0$  or  $P > 0$  where  $P$  is a polynomial in  $c$ . In the former case we want to find upper bounds of negative terms to get sufficient conditions for  $P < 0$ . For free-rider technologies,  $c > \bar{c}$ , this is possible by replacing  $c$  by  $\bar{c}$ . If we have to prove  $P > 0$  for free-driver technologies,  $c \leq \bar{c}$ , replacing  $c$  in negative terms by  $\bar{c}$  gives the needed lower bound.

**Remark 2.** For the proof of inequalities it makes sense to deviate from expressing  $a_1$  and  $a_2$  in terms of  $\bar{a}$  and  $a_\Delta$ ; rather,  $a_1$  and  $\delta$  where  $a_2 = a_1 + \delta$  is the best choice. With that,  $\bar{c} = b\delta/a_1$

**Example** If we have to prove  $-3a_1bc^2 + 2\delta b^2c < 0$  for  $c < \bar{c}$ , a sufficient condition is  $-3a_1bc\bar{c} + 2\delta b^2c < 0$ , which is true because  $-3a_1bc\bar{c} = -3\delta b^2c$ . In more complex situations, we will indicate the terms that are combined in that way.

## Appendix B On section 3

**Proof of Proposition 1** The social optimal deployment is  $q_1^{**} = q_2^{**} = \frac{2\bar{a}b}{c+4b}$ , so that  $Q^{**} = \frac{4\bar{a}b}{c+4b}$ .

- (i) This part of the Proposition is obvious.
- (ii) The only fact that needs clarification is the behavior of the fraction  $Q^*/Q^{**} = (c+4b)/(2c+4b)$ : It decreases monotonically from 1 at  $c = 0$  to  $1/2$  in the limit  $c \rightarrow \infty$ .
- (iii) The condition  $\pi_i(q_1^*, q_2^*) < \pi_i(q_1^{**}, q_2^{**})$  is equivalent with  $\frac{a^2b^2c^2}{c^3+8bc^2+20b^2c+16b^3} > 0$  so that  $(q_1^{**}, q_2^{**})$  is a Pareto improvement to  $(q_1^*, q_2^*)$ . The social optimum must be Pareto optimal because otherwise there would exist a constellation that has higher total welfare.

### Proof of Proposition 2

- (i) We have  $R_1 = R_2 = \frac{a^2b^2(3c+4b)}{2(c+2b)^2} > 0$  with  $\frac{dR_1}{dc} = -\frac{a^2b^2(3c+2b)}{2(c+2b)^3} < 0$ .
- (ii) The social optimal reference point is  $R^{**} = \frac{4a^2b^2}{(c+4b)}$ , which obviously declines in  $c$ . The condition  $R < R^{**}$  is equivalent with  $\frac{a^2b^2c^2}{c^3+8bc^2+20b^2c+16b^3} > 0$ .

## Appendix C On section 4

**Proof of Lemma 2** The reaction functions in (4) do not have an interior intersect iff  $c \leq \bar{c}$ . The only statement that remains to be shown is  $\frac{d\bar{c}}{da_\Delta} = \frac{2b\bar{a}}{(\bar{a}-a_\Delta)^2} > 0$ .

### C.1 Free-rider

#### Proof of Proposition 3

- (i) We have  $q_1^* = \frac{b}{c+2b}\bar{a} - \frac{b}{c}a_\Delta > 0$  because  $c > \bar{c}$  and  $q_2^* = \frac{b}{c+2b}\bar{a} + \frac{b}{c}a_\Delta > 0$ . In particular,  $q_1^* = 0$  and  $q_2^* = \bar{a} - a_\Delta = a_1$  at  $c = \bar{c}$ . In terms of the derivatives,

$$\frac{dq_1^*}{dc} = -\frac{b}{(c+2b)^2}\bar{a} + \frac{b}{c^2}a_\Delta \quad , \quad \frac{dq_2^*}{dc} = -\frac{b}{(c+2b)^2}\bar{a} - \frac{b}{c^2}a_\Delta < 0 .$$

Here,  $dq_1^*/dc$  has a root at  $\bar{c} + 2b\sqrt{\bar{a}a_\Delta}/(\bar{a} - a_\Delta) > \bar{c}$ .

- (ii) From (i) we get  $Q^* = \frac{2b}{c+2b}\bar{a}$ . Easy to verify is  $Q^{**} = \frac{4b}{c+4b}\bar{a}$ . Thus,  $Q^* < Q^{**}$  is obvious.

- (iii) The general directional derivative is as follows: Define  $f_i(t) := \pi_i(q_1^* + t\alpha, q_2^* + t\beta)$ . Then  $f'_i(0)$  is the  $(\alpha, \beta)$ -directional derivative of  $\pi_i$  at the Nash equilibrium. In this specific case we get

$$f'_1(0) = \frac{\beta b(\bar{a} - a_\Delta)}{c + 2b}(c - \bar{c}) \quad , \quad f'_2(0) = \frac{\alpha b}{c + 2b}((\bar{a} + a_\Delta)c + 2ba_\Delta)$$

and both expressions are positive for  $\alpha, \beta > 0$ . Whether the social optimal configuration is a Pareto Improvement for country 1 depends on the sign of

$$\begin{aligned} \pi_1(q_1^{**}, q_2^{**}) - \pi_1(q_1^*, q_2^*) = \\ \frac{(\bar{a}^2 - 6a_\Delta\bar{a} + a_\Delta^2)b^2c^3 + (8a_\Delta^2 - 20a_\Delta\bar{a})b^3c^2 + (20a_\Delta^2 - 16a_\Delta\bar{a})b^4c + 16a_\Delta^2b^5}{2c^4 + 16bc^3 + 40b^2c^2 + 32b^3c} \end{aligned}$$

At  $c = \bar{c}$ , this expression reads  $-16b^3a_\Delta^2\bar{a}^2(\bar{a} + a_\Delta)/(\bar{a} - a_\Delta)^3$  and is negative when  $a_\Delta > 0$ . By continuity, this extends to a full neighborhood. For large  $a_\Delta$  the  $c^3$ -coefficient in the numerator gets negative, implying that  $(q_1^{**}, q_2^{**})$  is not a Pareto improvement for large  $c$ . The situation is different for country 2. Here,

$$\begin{aligned} \pi_2(q_1^{**}, q_2^{**}) - \pi_2(q_1^*, q_2^*) = \\ \frac{(\bar{a}^2 + 6a_\Delta\bar{a} + a_\Delta^2)b^2c^3 + (20a_\Delta\bar{a} + 8a_\Delta^2)b^3c^2 + (16a_\Delta\bar{a} + 20a_\Delta^2)b^4c + 16a_\Delta^2b^5}{2c^4 + 16bc^3 + 40b^2c^2 + 32b^3c} \end{aligned}$$

and this is clearly positive.

**Proof of Proposition 4** For the notation used in proofs of inequality, see Appendix A.

$$(i) \quad \frac{dR_1}{dc} = \frac{b^2}{2c^2(c+2b)^3} \left( \underbrace{-2a_1\delta c^3}_{(B)} - \underbrace{3a_1^2c^3}_{(A)} + \underbrace{2\delta^2bc^2}_{(B)} - \underbrace{2a_1^2bc^2}_{(C)} + \underbrace{3b^2\delta^2c}_{(A)} + \underbrace{2b^3\delta^2}_{(C)} \right) < 0 \text{ with equality}$$

$$\text{for } c = \bar{c}. \quad \frac{dR_2}{dc} = \frac{b^2}{2c^2(c+2b)^3} \left( -\delta^2c^3 - \underbrace{4a_1\delta c^3}_{(B)} - \underbrace{3a_1^2c^3}_{(B)} - \underbrace{4a_1b\delta c^2}_{(A)} - \underbrace{2a_1^2bc^2}_{(A)} + \underbrace{3b^2\delta^2c}_{(B)} + \underbrace{2b^3\delta^2}_{(A)} \right)$$

which is strictly negative at  $c = \bar{c}$ .

$$(ii) \quad \frac{dR_1}{da_\Delta} = -\frac{(\bar{a}+a_\Delta)b^2c+2a_\Delta b^3}{c^2+2bc} < 0 \text{ and } \frac{dR_2}{da_\Delta} = \frac{b^2((\bar{a}-a_\Delta)c-2a_\Delta b)}{c^2+2bc} > 0 \text{ (because } c > \bar{c}\text{)}.$$

$$(iii) \quad R^{**} - R = \frac{b^2c^2}{c^3+8bc^2+20b^2c+16b^3}\bar{a}^2 + \frac{b^2}{c}a_\Delta^2$$

## C.2 Free-driver

**Proof of Proposition 5**

- (i) Obvious from  $q_2^* = \frac{b}{b+c}(\bar{a} + a_\Delta)$ .  
 (ii) Simple algebra.  
 (iii) The directional derivatives are

$$f'_1(0) = \frac{(\beta + \alpha)b(\bar{a} - a_\Delta)}{c + b}(c - \bar{c}) \quad , \quad f'_2(0) = \frac{\alpha bc}{c + b}(\bar{a} + a_\Delta)$$

and both expressions are positive when  $\alpha > 0$  and  $\beta < -\alpha$ . Simple algebra shows that  $\pi_2(q_1^{**}, q_2^{**}) - \pi_2(q_1^*, q_2^*) < 0$  iff  $c < 4a_\Delta^2b/(3\bar{a}^2 + 6a_\Delta\bar{a} - a_\Delta^2)$ . Evaluating  $\pi_1(q_1^{**}, q_2^{**}) - \pi_1(q_1^*, q_2^*)$  at this point gives  $ba_\Delta^2 \frac{12\bar{a}^3 + 30a_\Delta\bar{a}^2 + 8a_\Delta^2\bar{a} - 2a_\Delta^3}{9\bar{a}^3 + 36a_\Delta\bar{a}^2 + 45a_\Delta^2\bar{a} + 18a_\Delta^3}$ . The numerator is positive due to  $\bar{a} \geq a_\Delta$ .

**Proof of Proposition 6**

$$(i) \quad R_1 = \frac{b^2}{2(c+b)^2} \left( (2c+b)\bar{a}^2 - 2ba_\Delta\bar{a} - (2c+3b)a_\Delta^2 \right) \text{ with a root at } c_{R_1=0} = \frac{(3a_\Delta - \bar{a})b}{2(\bar{a} - a_\Delta)}$$

$$R_2 = \frac{b^2}{2(c+b)} (\bar{a} + a_\Delta)^2$$

$$\frac{dR_1}{dc} = \frac{b^2}{(c+b)^3} \left( \underbrace{\delta^2 b}_{(A)} + \underbrace{a_1 \delta b}_{(B)} - \underbrace{a_1 \delta c}_{(A)} - \underbrace{a_1^2 c}_{(B)} \right) > 0 \text{ with equality for } c = \bar{c}.$$

(ii) See (i).

$$(iii) \quad \text{Helpful notations: } R_2 = R^{**} \text{ at } c_{R_2=R^{**}} = -\frac{4b(\bar{a}^2 - 2a_\Delta\bar{a} - a_\Delta^2)}{7\bar{a}^2 - 2a_\Delta\bar{a} - a_\Delta^2}. \text{ This is positive iff } a_\Delta >$$

$$(\sqrt{2} - 1)\bar{a}. \text{ For } R = R_1 + R_2, \frac{dR}{dc} = 0 \text{ at } c_{R'=0} = \frac{(3a_\Delta - \bar{a})b}{3\bar{a} - a_\Delta}. \text{ This is positive iff } a_\Delta > \bar{a}/3$$

(same condition as "R<sub>1</sub> has a root"). **If counter-R&D is not possible**, the statement that R is positive and monotonical is justified with two arguments: For all c > c<sub>R<sub>1</sub>=0</sub>, where R = R<sub>1</sub> + R<sub>2</sub> anyways, this results from c<sub>R'=0</sub> < c<sub>R<sub>1</sub>=0</sub> (obvious because these expressions only differ in the denominator). For all other c we have R = R<sub>2</sub> and thus R inherits the characteristics 'positive' and 'decreasing in c'. That at the crossing point R = R<sup>\*\*</sup> necessarily max{R<sub>1</sub>, 0} = 0 results from the proof R < R<sup>\*\*</sup> if R = R<sub>1</sub> + R<sub>2</sub> below. **If counter-R&D is possible**, R =  $\frac{b^2}{2(c+b)^2} ((\bar{a}^2 - a_\Delta^2)(c+2b) + 2\bar{a}(\bar{a} + a_\Delta)) >$

$$0 \text{ with } \frac{dR}{dc} = -\frac{b^2}{2(c+b)^3} \left( (3c+b)\bar{a}^2 - 2(b-c)a_\Delta\bar{a} - (c+3b)a_\Delta^2 \right). \text{ We have } R^{**} - R =$$

$$\frac{b^2}{2c^3 + 12bc^2 + 18b^2c + 8b^3} \left( (5a_1^2 + 4a_1\delta + \delta^2)c^2 + 2a_1^2bc \underbrace{-2a_1\delta bc}_{(A)} + 2\delta^2b^2 \right) > 0.$$

**Appendix D On section 5**

**Proof of Proposition 8** In the free-rider region: With  $C_1 = \frac{b^2}{2c(c^3 + 5bc^2 + 8b^2c + 4b^3)}$ ,

$$R_1^{\text{Pub}} = C_1 \left( \underbrace{2a_1\delta c^3}_{(A)} + \underbrace{2a_1^2c^3}_{(B)} + \underbrace{4a_1\delta bc^2}_{(B)} + \underbrace{3a_1^2bc^2}_{(C)} + 2a_1\delta b^2c \underbrace{-\delta^2bc^2}_{(A)} \underbrace{-2\delta^2b^2c}_{(B)} \underbrace{-\delta^2b^3}_{(C)} \right) > 0$$

The derivative  $\frac{dR_1^{\text{Pub}}}{dc}$  at  $\bar{c}$  is  $a_1^4/(2(a_1 + \delta)^2) > 0$ . Thus, with  $\lim_{c \rightarrow \text{inf ty}} R_1^{\text{Priv}} = 0$ ,  $R_1^{\text{Priv}}$  cannot be monotone. With  $C_2 = \frac{b^2}{2c(c^3 + 5bc^2 + 8b^2c + 4b^3)}$ ,

$$R_2^{\text{Pub}} = C_2 \left( \underbrace{2a_1\delta c^3}_{(A)} + \underbrace{2a_1^2c^3}_{(B)} + \underbrace{2a_1\delta bc^2}_{(B)} + \underbrace{3a_1^2bc^2}_{(C)} \underbrace{-2\delta^2bc^2}_{(A)} \underbrace{-4\delta^2b^2c}_{(B)} \underbrace{-2a_1\delta b^2c}_{(C)} \underbrace{-\delta^2b^3}_{(C)} \right) > 0,$$

getting zero at  $\bar{c}$ . We have  $R_1^{\text{Pub}} - R_2^{\text{Pub}} = \frac{2b^3\bar{a}a_\Delta}{c^2 + 3bc + 2b^2} > 0$ .

In the free-driver region:

$$R_1^{\text{Pub}} = \frac{b^2}{2(c+b)^2} \left( (\bar{a}^2 + 2a_\Delta\bar{a} - 3a_\Delta^2)c - 4a_\Delta^2b \right),$$

which is zero at  $0 < c = \frac{4ba_\Delta^2}{\bar{a}^2 + 2a_\Delta\bar{a} - 3a_\Delta^2} < \bar{c}$ . The c-derivative is, with  $a_2 = a_1 + \delta$ ,

$$\frac{dR_1^{\text{Pub}}}{dc} = \frac{b^2}{2(b+c)^3} \left( \underbrace{2\delta^2b}_{(A)} + \underbrace{2a_1\delta b}_{(B)} + \underbrace{a_1^2b}_{(A)} \underbrace{-2a_1\delta c}_{(A)} \underbrace{-a_1^2c}_{(B)} \right) > 0.$$

Clearly,  $R_2^{\text{Pub}} = 0$ .

# Informativeness of Experiments for MEU – A Recursive Definition <sup>\*†</sup>

Daniel Heyen

Boris Wiesenfarth

## Abstract

The well-known Blackwell's theorem states the equivalence of statistical informativeness and economic valuableness. Çelen (2012) generalizes this theorem, which is well-known for subjective expected utility (SEU), to maxmin expected utility (MEU) preferences. We demonstrate that the underlying definition of the value of information used in Çelen (2012) is in contradiction with the principle of recursively defined utility. As a consequence, Çelen's framework features dynamic inconsistency. Our main contribution consists in the definition of a value of information for MEU preferences that is compatible with recursive utility and thus respects dynamic consistency.

**Keywords:** Blackwell; Value of Information; Maxmin Expected Utility; Recursive Utility.

**JEL Codes:** D81; D83; D84.

## 1 Introduction

For decades economists have been studying the relationship between decision-making under uncertainty and the so-called value of information. A famous and well-known result in this context is Blackwell's theorem (Blackwell 1953) stating that an experiment is more valuable than another if and only if the same experiment is more informative than the latter. In order to obtain this equivalence (e.g. Crémer 1982), a standard assumption has been that decision-makers are subjective expected utility (SEU) maximizers, cf. Savage (1972).

The objective of Çelen (2012) is to extend the Blackwell theorem to MEU preferences, which were axiomatized by Gilboa and Schmeidler (1989). In this note we demonstrate that Çelen's proof relies on a value of information for MEU preferences that is not defined via backward induction and thus is incompatible with the intertemporal extension of MEU by Epstein and Schneider (2003). In particular, optimal strategies in Çelen's framework prescribe decisions conditional on signal realizations that an MEU decision-maker will not find optimal to adhere to once those signal realizations have been observed. In this

---

\*We would like to acknowledge helpful comments by Jürgen Eichberger, Jean-Philippe Lefort and Boğaçhan Çelen. The first author acknowledges the support of the German Science Foundation, grant no. GO 1604/2-1.

†An almost identical version of this article is forthcoming in the *Journal of Mathematical Economics*, doi: 10.1016/j.jmateco.2014.12.002.

sense, Çelen's framework features dynamic inconsistency. Our contribution is to define a value of information that is compatible with the recursive intertemporal formulation of MEU preferences by Epstein and Schneider (2003).

## 2 Framework and definition of the value of information in Çelen (2012)

In the following, we adopt Çelen's framework and notation. Let  $\Omega := \{\omega_1, \dots, \omega_n\}$  be the finite set of *states* and  $X := \{a_1, \dots, a_\chi\}$  the finite set of *actions* available to a decision-maker. Moreover, let  $\Delta(\Omega)$  and  $\Delta(X)$  be the set of all probability distributions defined on  $\Omega$  and  $X$ , respectively. Let further  $u : \Omega \times X \rightarrow \mathbb{R}$  be a *utility function* and  $\mathbf{u}$  with  $u_{ij} = u(\omega_i, a_j)$  the corresponding utility matrix. An SEU decision-maker is characterized by  $(\boldsymbol{\pi}, \mathbf{u})$ , where  $\boldsymbol{\pi} \in \Delta(\Omega)$  is a prior over the states.

An *experiment* is a tuple  $(S, \mathbf{p})$  with the signal space  $S = \{s_1, \dots, s_\sigma\}$  and the Markov matrix  $\mathbf{p}$  with  $p_{ij} = \Pr(s_j|\omega_i)$  for  $s_j \in S$ . Çelen introduces a *strategy* as a vector valued mapping  $f : S \rightarrow \Delta(X)$ , thus characterizing all (mixed) actions the decision maker plans to take after observing certain signal realizations  $s$ . The  $\sigma \times \chi$ -matrix  $\mathbf{f}$  is defined such that  $(f_{i1}, \dots, f_{i\chi}) := f(s_i)$ .

In this framework, Çelen determines the *value of the experiment*  $(S, \mathbf{p})$  for a given strategy  $\mathbf{f}$  as

$$\mathcal{U}_{(\boldsymbol{\pi}, \mathbf{u})}^{\mathbf{f}}(S, \mathbf{p}) = \sum_j \Pr(s_j) \sum_i \Pr(\omega_i|s_j) \sum_k f_{jk} u(\omega_i, a_k) \quad (1)$$

$$= \sum_j \sum_i p_{ij} \pi_i \sum_k f_{jk} u_{ik} \quad (\text{by Bayes' rule}) . \quad (2)$$

With a strategy  $\mathbf{f}^*$  maximizing (2), Çelen defines  $\mathcal{U}_{(\boldsymbol{\pi}, \mathbf{u})}^*(S, \mathbf{p}) = \mathcal{U}_{(\boldsymbol{\pi}, \mathbf{u})}^{\mathbf{f}^*}(S, \mathbf{p})$  as the value of the experiment for an SEU decision-maker.

Building on this, Çelen extends the definition of the value of an experiment to the class of MEU preferences. For that purpose, he characterizes an MEU decision-maker by  $(A, u)$ , where  $A \subset \Delta(\Omega)$  is a convex and compact *set* of priors. As a counterpart of  $\mathcal{U}_{(\boldsymbol{\pi}, \mathbf{u})}^*(S, \mathbf{p})$ , he defines

$$\mathcal{W}_{(A, u)}^* = \max_{\mathbf{f}} \min_{\boldsymbol{\pi} \in A} \mathcal{U}_{(\boldsymbol{\pi}, \mathbf{u})}^{\mathbf{f}}(S, \mathbf{p}) \quad (3)$$

as the value of an experiment  $(S, \mathbf{p})$  for an MEU decision-maker. It is expression (3) that Çelen relies on in his proof of the generalized Blackwell's theorem.

## 3 A recursively defined MEU value of information

It is insightful to note that Çelen's framework essentially constitutes an intertemporal setting with two periods. In the second period, *after* observing a signal realization, the

decision maker takes a (mixed) action. In the first period, *before* observing the signal realization, the value of the experiment  $(S, \mathbf{p})$  is determined. Çelen accounts for the intertemporal structure insofar as he considers strategies, that is complete contingent plans for appropriate play after observing signal realizations.

His formulation, however, is in contrast to the usual intertemporal formulation of MEU preferences that is provided by Epstein and Schneider (2003). One of the main characteristics of their recursive definition of intertemporal MEU is the compatibility with backward induction. We follow the approach in Epstein and Schneider (2003) and develop an alternative definition of the value of information for MEU preferences.

According to backward induction, the first step to define a value of information is to determine the value of the final decision given that the decision maker chooses an optimal action for all possible signal realizations  $s_j$ ,  $j = 1, \dots, \sigma$ . For MEU preferences this value is

$$V(\mathcal{M}_j) = \max_{g \in \Delta(X)} \min_{\mu \in \mathcal{M}_j} \mathbb{E}_\mu[u](g), \quad (4)$$

where  $\mathbb{E}_\mu[u](g) = \sum_{k,i} g_k \mu_i u_{ik}$  denotes the expected utility under action  $g \in \Delta(X)$  and ex-post belief  $\mu$ . The optimal action is determined considering the worst posterior  $\mu \in \mathcal{M}_j$ . Formally, the set of posteriors is  $\mathcal{M}_j = \{\pi(\cdot|s_j) : \pi \in A\}$ , where  $\pi(\cdot|s_j)$  denotes the conditional probability of the prior  $\pi \in \Delta(\Omega)$  given the signal  $s_j$ . We obtain  $\pi(\cdot|s_j)$  via Bayes' rule and update each prior  $\pi$  in this way.<sup>1</sup>

Building on (4), we can define the value of the experiment  $(S, \mathbf{p})$  for MEU preferences as

$$\mathcal{V}_{(A,u)} = \min_{\pi \in A} \sum_{i,j} \pi_i p_{ij} V(\mathcal{M}_j). \quad (5)$$

As usual, the value of information is obtained by taking the expectation over all possible signal realizations. Due to MEU preferences the value of the experiment is the worst of those expectations. This alternative way of defining the value of an experiment is in line with the intertemporal model of recursive utility under multiple priors as introduced in Epstein and Schneider (2003, 2007).

The key characteristic of (5) is that optimal actions are determined with the maxmin rule for each signal realization  $s_j$  individually. In particular, the worst posterior in (4) in general depends on the signal realization  $s_j$ . This is in contrast to (3). By following the derivation of the SEU counterpart, essentially the step from (1) to (2), Çelen silently assumes that the worst prior from the ex-ante perspective coincides with the preimage of *all* worst posteriors, irrespective of the signal realization. For the SEU decision-maker this argumentation is innocent as there is a unique prior, and thus a unique posterior as well. For the MEU decision-maker, however, this argument is in conflict with backward induction.

---

<sup>1</sup>Epstein and Schneider (2007) show that further restrictions on the set  $\mathcal{M}$  can be made. For the sake of simplicity, you may think of full Bayesian updating.



In Appendix A, we demonstrate that the conflict of Çelen’s framework with intertemporal recursive utility can be made even more concrete. We provide an example in which the optimal strategy derived in Çelen’s framework prescribes actions that are different from what an MEU decision-maker will actually do after observing those signals realizations.<sup>2</sup> This supports our claim that the value of information for MEU preferences should be defined by (5). By construction, our definition of the value of information is compatible with dynamic consistency.

## 4 Results and Discussion

We have shown that Çelen’s proof of Blackwell’s theorem only applies to a value of information that is defined in a non-recursive utility framework. We have offered a definition for the value of information derived via backward induction, thus compatible with the dynamic consistent intertemporal axiomatization of Epstein and Schneider (2003). Consequently, we argue that the proof of Blackwell’s theorem should deal with expression (5) as the definition of the value of information for MEU preferences. This proof is still pending.

## References

- Blackwell, D. (1953). Equivalent comparison of experiments, *The Annals of Mathematical Statistics* **24**: 265–272.
- Çelen, B. (2012). Informativeness of experiments for MEU, *Journal of Mathematical Economics* **48**: 404–406.
- Crémer, J. (1982). A simple proof of blackwell’s comparison of experiments theorem, *Journal of Economic Theory* **27**: 439–443.
- Epstein, L. G. and Schneider, M. (2003). Recursive multiple-priors, *Journal of Economic Theory* **113**(1): 1–31.
- Epstein, L. G. and Schneider, M. (2007). Learning under ambiguity, *The Review of Economic Studies* **74**(4): 1275–1303.
- Gilboa, I. and Schmeidler, D. (1989). Maxmin expected utility with non-unique prior, *Journal of Mathematical Economics* **18**(2): 141–153.
- Savage, L. (1972). *The foundations of statistics*, Dover Publications.

---

<sup>2</sup>One could think that the reason we observe this form of dynamic inconsistency is the missing assumption of *rectangularity* of the prior set, a key assumption in Epstein and Schneider (2003) to ensure dynamic consistency within an intertemporal setting of recursive utility. But this is not the case. Even though Çelen’s setting is not fully transferable to the setting of Epstein and Schneider, in particular the analysis in Epstein and Schneider (2007) suggests that rectangularity is no issue in this setting, simply because the learning process is defined via conditional one-step-ahead conditionals, as required by Epstein and Schneider (2003). The reason for the violation of dynamic consistency in Çelen’s framework is that intertemporal utility is defined in a non-recursive way, thus incompatible with dynamic consistency right from the start.

## Appendix A Example demonstrating the dynamic inconsistency of Çelen's framework

We restrict the number of states of the world  $\Omega = \{\omega_1, \omega_2\}$ , actions  $X = \{a_1, a_2\}$  and signal realizations  $S = \{s_1, s_2\}$  to two. Payoffs are specified by  $u_{11} = 1$ ,  $u_{12} = -1$ ,  $u_{22} = 2$  and  $u_{21} = 0$ . This is a simple example of a setting in which the decision maker wants to learn the true  $\omega$  because action  $a_1$  is optimal if  $\omega = \omega_1$  and action  $a_2$  is optimal if  $\omega = \omega_2$ .

For the signal likelihood  $\lambda = p_{11} = p_{22}$  we assume  $1/2 < \lambda < 3/4$  and specify the set of priors as  $A = \{(\pi_1, 1 - \pi_1) : 1/4 \leq \pi_1 \leq 3/4\}$ .

Using equation (2), we obtain that the optimal strategy  $\mathbf{f}^*$  in Çelen's framework is

$$\mathbf{f}^*(s_1) = (1, 0) \quad , \quad \mathbf{f}^*(s_2) = (1/2, 1/2) \quad . \quad (6)$$

In words, the optimal strategy in the Çelen framework consists of taking action  $a_1$  if  $s = s_1$  and mixing over actions  $a_1$  and  $a_2$  with equal weights if  $s = s_2$ .

Next, we demonstrate that an MEU decision-maker operating with the principle of backward induction would deviate from  $\mathbf{f}^*$  as soon as the signal materializes. The decision rule after observing a signal realization  $s_j$  is given in (4), where  $g$  is a randomization over actions  $a_1$  and  $a_2$ , and  $\mathcal{M}_j \subset \Delta(\Omega)$  is the set of posteriors that depends on the set of priors  $A$ , the likelihood  $p$ , and the signal realization  $s_j$  observed. After simple algebra we get

$$g^*(s_1) = (3/4, 1/4) \quad \text{and} \quad g^*(s_2) = (3/4, 1/4) \quad . \quad (7)$$

In words, the optimal action of the MEU decision-maker, both after receiving  $s = s_1$  and  $s = s_2$ , is to mix over actions  $a_1$  and  $a_2$  with the ratio 3 to 1. This is in contrast to the behavior prescribed in (6). Our example thus demonstrates the dynamic inconsistency in Çelen's framework.

# Learning under Ambiguity: A Note on the Belief Dynamics of Epstein and Schneider (2007) \*

Daniel Heyen

## Abstract

Epstein and Schneider (2007) develop a framework of learning under ambiguity, generalizing maxmin preferences of Gilboa and Schmeidler (1989) to intertemporal settings. The specific belief dynamics in Epstein and Schneider (2007) rely on the rejection of initial priors that have become implausible over the learning process. I demonstrate that this feature of ex-post rejection of theories gives rise to choices that are in sharp contradiction with ambiguity aversion. Concretely, the intertemporal maxmin decision-maker equipped with such belief dynamics, under prevalent conditions, prefers a bet in an ambiguous urn over the same bet in a risky urn. I offer two modifications of their framework, each of which is capable of avoiding this anomaly.

**Keywords:** Learning under Ambiguity; Bayesian Updating; Multiple Priors; Maxmin Expected Utility; Ambiguity Aversion.

**JEL Codes:** D81; D83.

## 1 Introduction

Since Ellsberg (1961), ambiguity averse preferences as opposed to subjective expected utility preferences (SEU, Savage 1972) have been a vital economic research topic. One of the most influential axiomatizations of ambiguity aversion is the multiple prior model with maxmin expected utility (MEU) of Gilboa and Schmeidler (1989). The framework of Gilboa and Schmeidler (1989), however, is atemporal and thus not capable of reflecting intertemporal ambiguity aversion. In a recent and influential paper, Epstein and Schneider (2007) (henceforth ES) develop, based on Epstein and Schneider (2003), a tractable framework of intertemporal maxmin preferences and thus open ambiguity aversion to a dynamic learning environment. This framework has already been used extensively in various contexts, including financial markets (Condie and Ganguli 2011; Garlappi et al. 2007; Ju and Miao 2012; Leippold et al. 2008) and real options (Nishimura and Ozaki 2007; Riedel 2009).

---

\*I want to thank Boris Wiesenfarth for discussions in the early stages of this paper and am grateful to Timo Goeschl and Tobias Pfommer for their helpful comments and suggestions. I acknowledge the support of the German Science Foundation, grant no. GO 1604/2-1.

The focus of this paper is a problematic characteristic of the belief dynamics in Epstein and Schneider (2007). The updating process in ES involves the rejection of initial beliefs that have become less plausible given the observed signal history. I demonstrate that the rejection of beliefs renders possible a switch in preferences. To illustrate this, I construct a simple example in which an ambiguity averse decision-maker switches to ambiguity loving behavior after observing one draw from a payoff-relevant urn. The reason for this instability of preferences is found in the rejection of initial beliefs in the updating procedure. It can happen that the uniform distribution, the standard initial prior of the SEU decision-maker, is rejected in the reevaluation procedure of ES. As a consequence, the set of updated beliefs does not include the posterior of the SEU decision-maker, who serves as the usual reference point to define ambiguity preferences. This gives rise to the switch to ambiguity loving behavior.<sup>1</sup> Furthermore, I show that this anomaly of intertemporal beliefs is not just an artifact, but rather a pervasive and general feature of the ES belief dynamics. Finally, I suggest two modifications of the ES setting to ensure stable ambiguity averse preferences over time.

I proceed as follows. In section 2, I recapitulate the basic cornerstones of Epstein and Schneider (2007). The simple example demonstrating the ES anomaly in a standard setting is found in section 3. Section 4 generalizes the example to symmetric urns with an arbitrary number of balls and shows that the ES anomaly may occur under very general conditions. The two parts of the theorem are illustrated graphically in section 5. In section 6, I offer two alternatives to overcome the problems while keeping the general structure of ES. I conclude in section 7.

## 2 Intertemporal maxmin preferences – The setting of Epstein and Schneider (2007)

The main purpose of this section is to recall the basic components of Epstein and Schneider (2007). All readers familiar with the ES framework may thus skip this section.

To reemphasize the motivation of their framework, I consider two urns that are simplifications of the scenarios in ES. Both urns only features risk or ambiguity, but not both kinds of uncertainty at the same time. These urns are used in section 3 to demonstrate the existence of the ES anomaly and are then generalized in section 4 to arbitrary symmetric settings.

### 2.1 Unknown parameter

To motivate the introduction of the parameter space in ES, consider the two urns in Figure 1. Both urns contain exactly three balls (this assumption will be abandoned in

---

<sup>1</sup>ES notes that Ellsberg type behavior in the short run will converge (for a fixed composition urn) to ambiguity neutral behavior in the long run. The possibility of a change in preferences to ambiguity loving choices, however, was not mentioned.




Figure 1

section 4). In both urns there is, apart from a white and a black ball, a third ball that is either black or white. The composition of either urn is thus unknown to the decision-maker but will not change over the course of the experiment.<sup>2</sup> In the language of ES, the ratio of black balls in the urn is the unknown *parameter*  $\theta$  with possible values in the set  $\Theta = \{1/3, 2/3\}$ . The key difference between both urns is that for the risky Urn R, the decision-maker knows that the color of the third ball has been determined via a fair mechanism, e.g. a fair coin. This is the same kind of uncertainty as in scenario 1 in ES. For the ambiguous Urn A however, the decision-maker has no information on the mechanism that determined the color of the third ball. In contrast to scenario 2 in ES, Urn A features pure ambiguity.

## 2.2 State space and recursive utility

In every period and for each urn, one ball is drawn and then put back (sampling with replacement). The period state space is  $S_t = S = \{B, W\}$ , identical for all times. We denote by  $s_t \in S$  the color observed by the agent at time  $t$ . An agent's information at time  $t$  is the history  $s^t = (s_1, \dots, s_t)$ . The natural full state space is  $S^\infty$ .

The agent ranks consumption plans  $c = (c_t)$  according to recursively defined utility,

$$U_t(c; s^t) = \min_{p \in \mathcal{P}_t(s^t)} \mathbb{E}^p [u(c_t) + \beta U_{t+1}(c; s^t, s_{t+1})] , \quad (1)$$

where  $u$  and  $\beta$  have the usual properties. A central component in ES is  $\mathcal{P}_t(s^t)$ . This set of probability measures models beliefs about the next ball observed,  $s_{t+1}$ , given the history  $s^t$ . Such beliefs reflect ambiguity if  $\mathcal{P}_t(s^t)$  is a non-singleton, which is the appropriate description if the draw will be made from Urn A. Beliefs about the next draw from Urn R can be, as usual, described with  $\mathcal{P}_t(s^t)$  being a singleton. ES refer to  $(\mathcal{P}_t)$  as the *process of one-step-ahead beliefs*. Specifying beliefs in this way ensures dynamic consistency of the decision framework, as is shown in Epstein and Schneider (2003). To further clarify this set of beliefs, let us turn to the learning structure in the ES setting.

<sup>2</sup>Thus, we do not consider scenario 3 in Epstein and Schneider (2007) with unknown likelihoods

### 2.3 Learning

By observing data, here a sequence  $(s_t)$  of white and black balls, the decision-maker tries to learn the color of the third ball and thus the unknown parameter  $\theta \in \Theta$ . Ambiguity in initial beliefs about parameters can be represented by a *set*  $\mathcal{M}_0$  of probability measures on  $\Theta$ . The size of  $\mathcal{M}_0$  reflects the decision-maker's (lack of) confidence in the prior information on which initial beliefs are based.

In both urns, Urn R and Urn A, the likelihood of observing a black or a white ball is fully determined by the parameter  $\theta$ , the ratio of black balls in the urn. Obviously,  $l(s = B|\theta) = \theta$ . Throughout this paper, we restrict ourselves to those settings of unique likelihood functions.<sup>3</sup> Beliefs about parameters and the likelihood function jointly determine the process of one-step-ahead beliefs

$$\mathcal{P}_t(s^t) = \left\{ p_t(\cdot) = \int_{\Theta} l(\cdot|\theta) d\mu_t(\theta) : \mu_t \in \mathcal{M}_t(s^t) \right\}, \quad (2)$$

where  $\mathcal{M}_t(s^t)$  is the set of posterior beliefs after observing the data  $s^t$ . This set is basically the priorwise Bayesian update of initial beliefs  $\mu_0 \in \mathcal{M}_0$ . ES, however, incorporate the possibility that the decision-maker rejects at every  $t$  some implausible initial beliefs. The set of posteriors only include updates of initial beliefs that are not rejected in this process. The specific procedure of rejecting beliefs in ES, however, gives rise to the anomaly that is the topic of this paper. In the following subsection, we explain the rejection of beliefs in ES in detail.

### 2.4 Updating and reevaluation

To assess the plausibility of a "theory"  $\mu_0 \in \mathcal{M}_0$  after having observed the history of signals  $s^t$ , ES use the data density evaluated at  $s^t$ . I denote this plausibility of a theory  $\mu_0$  given the data  $s^t = (s_1, \dots, s_t)$  by

$$\text{Plaus}(\mu_0; s^t) = \int_{\Theta} \prod_{j=1}^t l(s_j|\theta) d\mu_0(\theta). \quad (3)$$

With the usual Bayesian updating, recursively defined by

$$d\mu_t(\cdot; s^t, \mu_0) = \frac{l(s_t|\cdot)}{\int_{\Theta} l(s_t|\theta') d\mu_{t-1}(\theta'; s^{t-1}, \mu_0)} d\mu_{t-1}(\cdot; s^{t-1}, \mu_0), \quad (4)$$

---

<sup>3</sup>ES also allow for ambiguity in likelihoods, that is a set of likelihoods  $\mathcal{L}$ . At any point in time, any element of  $\mathcal{L}$  might be relevant for generating the next observation. Multiple likelihoods refer to those components of a decision problem the decision-maker has decided that she will not try to (or is not able to) learn about. The findings of this paper are independent of  $\mathcal{L}$ , which is why I restrict this analysis to the simplest case with a single likelihood function.

ES define the set of posteriors

$$\mathcal{M}_t^\alpha(s^t) = \{\mu_t(\cdot; s^t, \mu_0) : \mu_0 \in \mathfrak{M}_0^\alpha(s^t)\} \quad (5)$$

as the set of prior-by-prior updates of  $\mathfrak{M}_0^\alpha(s^t)$ . Here,

$$\mathfrak{M}_0^\alpha(s^t) = \left\{ \mu_0 \in \mathcal{M}_0 \mid \text{Plaus}(\mu_0; s^t) \geq \alpha \max_{\tilde{\mu}_0 \in \mathcal{M}_0} \text{Plaus}(\tilde{\mu}_0; s^t) \right\} \quad (6)$$

is the set of theories (i.e. initial priors) that are not rejected after having observed the signal history  $s^t$ . Rejected are those initial priors that fail a maximum likelihood test against the most plausible prior, and the parameter  $\alpha$  governs how strict this maximum likelihood test is. ES consider  $0 < \alpha \leq 1$  as possible values for the rejection parameter. Ruling out  $\alpha = 0$  implies that they require the decision-maker to actually use this rejection device.

The likelihood-ratio test is more stringent and the set of posteriors smaller, the greater is  $\alpha$ . In the extreme case  $\alpha = 1$ , only parameters that achieve the maximum likelihood are permitted. If the maximum likelihood estimator is unique, ambiguity about parameters is resolved as soon as the first signal is observed. More generally, we have that  $\alpha > \alpha'$  implies  $\mathcal{M}_t^\alpha \subset \mathcal{M}_t^{\alpha'}$ . It is important that the test is done after every history. In particular, a theory that was disregarded at time  $t$  might look more plausible at a later time and posteriors based on it may again be taken into account.

## 2.5 Subjective expected utility vs. maxmin preferences

Some words on the relation between subjective expected utility (SEU) and intertemporal maxmin (MEU) preferences are in order to round up this section. For Urn R, the initial prior  $(1/2, 1/2)$  that puts equal weight on both parameters is obviously the best description of the decision-maker's state of knowledge, irrespective of whether the decision-rule is SEU or MEU.

For Urn A this is different. An SEU decision-maker – who is the standard Bayesian decision-maker – is by definition characterized by a single belief. As objective knowledge about Urn A is not available, the question how the initial prior is determined is in general not easy to answer (Maskin 1979). Due to the perfect symmetry of this setting, however, it is clear that a Bayesian decision-maker would, according to the *principle of insufficient reason* (see for instance Gilboa 2009), hold the initial prior that assigns equal probability to  $\theta = 1/3$  and  $\theta = 2/3$ . As a consequence, an SEU decision-maker sees no difference between Urn R and Urn A.

### 3 A simple example demonstrating the switch in preferences

In this section I design a simple example to illustrate the key problem in the ES setting. It involves an MEU decision-maker who initially features the usual ambiguity averse preferences. After observing one draw from the urns, however, she switches her preferences and partially exhibits ambiguity loving behavior.

#### 3.1 Preliminaries

The example is constructed within the standard infinite horizon setting with  $S^\infty$ . I will, however, only compare the betting behavior on the color of the next ball before and after a single signal realization  $s$ . That is, I will focus on the betting preferences at  $t = 0$  and  $t = 1$ . For simplicity, I consider exclusively the two bets  $1_B0$  and  $1_W0$ , where the bet  $1_B0$  involves a payment of 1\$ if the color of the ball drawn is black and 0\$ if the ball is white. The bet  $1_W0$  is defined similar.

The example rests on the simple three-ball-urns introduced in section 2, cf. Figure 1. Consequently, the parameter space consists of the two possible ratios of black balls in the urn,  $\Theta = \{1/3, 2/3\}$ . Any prior and posterior over the parameter space has the form  $(\nu, 1 - \nu)$  where the extreme points  $(1, 0)$  and  $(0, 1)$  correspond to full weight on the parameter  $\theta = 1/3$  and  $\theta = 2/3$ , respectively.

For Urn R, both the SEU and the MEU decision-maker hold the uniform distribution as the initial prior. This is also the prior the SEU decision-maker associates with the ambiguous Urn A. The MEU decision-maker, in contrast, holds a set of initial priors regarding Urn A. For simplicity, let the set in this example be the full set of priors  $\mathcal{M}_0 = \{(\nu, 1 - \nu) \mid 0 \leq \nu \leq 1\}$ . Let the rejection parameter  $\alpha$  be  $4/5$ , which means that the MEU decision-maker only updates the initial priors with a plausibility of at least 0.8 of the maximal plausibility.

#### 3.2 Ambiguity aversion before observing the signal

Let us first compare the betting preferences of the SEU and the MEU decision-maker before the signal has been observed. Irrespective of the color to bet on, the SEU is indifferent between the bet on Urn R and Urn A as for her both urns are essentially the same. The MEU decision-maker, however, has – irrespective of the color to bet on – a clear preference for betting on Urn R. The reason is basically expression (1). For bet  $1_B0$ , where  $B$  is the favorable color, the worst scenario is that with the lowest number of black balls in the urn. Thus the MEU decision-maker rests her decision on the prior  $(1, 0)$  that puts full weight on  $\theta = 1/3$ . The associated expected payoffs are  $1/3\$$ . A similar argument for bet  $1_W0$  shows that the worst prior is  $(0, 1)$ , again with the expected payoff  $1/3\$$ . This has to be compared to the expected payoffs for Urn R. Clearly, the prior



$(1/2, 1/2)$  is for both bets associated with expected payoffs of  $1/2\$$ . Thus, the MEU decision-maker strictly prefers either bet in Urn R over Urn A. This is the well-known ambiguity averse behavior.

### 3.3 Switch to ambiguity loving behavior after learning

We now compare the betting behavior after a signal  $s$  has been observed. To make behavior comparable, we restrict to the case that the balls drawn from Urn R and Urn A have the same color. Without loss of generality, say  $s = B$ . This transforms the SEU decision-maker's belief from the initial prior  $(1/2, 1/2)$ , irrespective of the urn, into the posterior  $(1/3, 2/3)$ , reflecting the increased subjective probability for the scenario that the unknown ball is black.

The MEU decision-maker shares this view for Urn R, but naturally has a different take on Urn A. Here, the set of initial priors  $\mathcal{M}_0$  is, according to (5), updated to the set  $\mathcal{M}_1^\alpha(s)$ . To recapitulate the ES procedure explained in section 2.4, the first task is to find the most plausible theory  $\mu_0 \in \mathcal{M}_0$ . This is clearly  $(0, 1)$ . The plausibility of this theory (cf. (3)) is  $2/3$ . With  $\alpha = 4/5$ , the MEU decision maker rejects all theories with a plausibility less than  $4/5 \cdot 2/3$  and thus keeps the set  $\mathfrak{M}_0(s) = \{(\nu, 1 - \nu) \mid 0 \leq \nu \leq 2/5\}$ . Finally, this set is updated to  $\mathcal{M}_1^\alpha(s) = \{(\nu, 1 - \nu) \mid 0 \leq \nu \leq 1/4\}$ .

Let us again compare betting preferences of the SEU and the MEU decision-maker. The SEU remains, for either bet, indifferent between Urn R and Urn A. Turning to the MEU decision-maker, consider first the bet  $1_W0$ . Urn R promises expected payoffs of  $4/9\$$ . For Urn A the worst belief is still  $(0, 1)$ , associated with expected payoffs of  $1/3\$ < 4/9\$$ . Thus, the MEU decision-maker still prefers the bet  $1_W0$  in Urn R over Urn A.

Most intriguing, this is different for the bet  $1_B0$ . Urn R promises an expected payoff of  $5/9\$$ . For Urn A, the worst posterior in  $\mathcal{M}_1^\alpha$  is  $(1/4, 3/4)$ . This translates into expected payoffs of  $7/12\$$ , which is larger than  $5/9\$$ . The maxmin decision-maker thus prefers the bet  $1_B0$  in Urn A over the same bet in Urn R. This very surprising and a clear contradiction to ambiguity averse preferences. The reason for that switch in behavior stems from the fact that in general ES reject too many theories. Here, with a rather high  $\alpha = 4/5$ , the critical ambiguity neutral SEU prior  $(1/2, 1/2)$  is rejected. This prior is critical because with it also all 'pessimistic priors' that would give rise to ambiguity averse choices are rejected. As a consequence, the set of posteriors  $\mathcal{M}_1$  only contains optimistic beliefs that give rise to ambiguity loving choices.

One could argue that this problematic betting behavior can be avoided by adequately choosing  $\alpha$ . In the example, any  $\alpha < 3/4$  would not give rise to the switch in preferences, at least not at  $t = 1$ . However,  $\alpha$  was introduced by ES to be a characteristic of the decision-maker and it is thus unnatural to adjust  $\alpha$  to the specific setting.

The next section will demonstrate that the ES anomaly is pervasive. For every  $\alpha > 0$

there is a similar setting to that considered in the example for which such a disconnect in the behavior of the MEU decision-maker occurs. It is thus not possible to find a rejection parameter  $0 < \alpha \leq 1$  that is not prone to the ES anomaly.

## 4 The general result

In this section, I first generalize the setting from urns with three balls to arbitrary symmetric settings. This provides the framework to demonstrate that the problematic characteristic of the ES updating procedure is not restricted to specific urns. Indeed, we can show that each pair of generalized urns has this property for some rejection parameter  $0 < \alpha \leq 1$ . Even more interesting and less obvious, for each  $0 < \alpha \leq 1$  we can construct a pair of urns and a signal history  $s^t$  such that the MEU decision-maker features the switch to *ambiguity loving* behavior.

### 4.1 Generalized urns

The first step is to define generalizations of Urn R and Urn U. I restrict to symmetric settings in the sense that a priori both colors are interchangeable. The generalized Urn R and Urn A, I denote them by  $U_R(n, k)$  and  $U_A(n, k)$ , respectively, have exactly  $2n + k$  balls. It is known that  $n$  balls are black,  $n$  balls are white and each of the remaining  $k$  balls can either be black or white. As will become clear in what follows, the urns in the example in the previous section correspond to the case  $n = k = 1$ . To generalize the three balls urn example, we assume that the number of black balls within the  $k$  unknown balls in urn  $U_R(n, k)$  is uniformly distributed.<sup>4</sup> The urn  $U_A(n, k)$  is basically the same, but without information about the distribution of the  $k$  unknown balls. In either case, the parameter set is  $\Theta = \{n/(2n + k), \dots, (n + k)/(2n + k)\}$  with  $\theta \in \Theta$  being the true fraction of black balls in the urn, unknown to the decision-maker. The period state space is again  $S$  with the full state space  $S^\infty$ . The likelihood functions are fully specified by  $l(s = B|\theta) = \theta$ .

In terms of initial prior, it is clear that for  $U_R(n, k)$  both SEU and MEU decision-maker have the initial prior  $(1/(k + 1), \dots, 1/(k + 1))$  on  $\Theta$ . This is also the initial prior the SEU decision-maker holds for Urn A. In contrast, the MEU decision-maker operates with a set  $\mathcal{M}_0$  of initial priors, which is by definition a subset of  $\Delta(\{0, \dots, k\})$ . The natural assumption I make is that the uniform distribution is an element of the set of initial priors,  $(1/(k + 1), \dots, 1/(k + 1)) \in \mathcal{M}_0$ .

It is convenient to focus on simple and tractable sets for  $\mathcal{M}_0$ . As a generalization of

---

<sup>4</sup>An alternative generalization would be that of  $k$  independent coin flips. My choice, however, is more intuitive and technically simpler

intervals around  $1/2$  in the case  $k = 1$ , I consider sets of the form

$$\Delta_\epsilon^{(k)} = \left\{ (\nu_0, \dots, \nu_k) \mid \sum_i \nu_i = 1, 0 \leq \epsilon \leq \nu_i \leq 1 - k\epsilon \leq 1 \forall i, \epsilon < \frac{1}{k+1} \right\}. \quad (7)$$

By construction, all these sets contain the uniform distribution. The full set of priors  $\Delta$  is the special case  $\Delta_\epsilon^{(k)}$  with  $\epsilon = 0$ , and for  $\epsilon \rightarrow 1/(k+1)$  the set collapses to a singleton with the uniform distribution as the only element.

## 4.2 Theorem

I now have the toolkit to formulate the main finding. The intertemporal maxmin decision-maker in the ES sense is characterized by the parameter  $\alpha$  that describes to what extent theories are rejected ex-post. I have shown in section 3 that also the uniform distribution, the initial prior of the SEU decision-maker, can be rejected in this process. As a consequence, all pessimistic priors are rejected as well, giving rise to ambiguity loving choices. This anomaly occurs under very general conditions. Concrete,

**Theorem 4.1.** *The ES anomaly can be characterized by two statements.*

- (i) *For any pair  $(n, k) \in \mathbb{N}^2$ , any  $\mathcal{M}_0 = \Delta_\epsilon^{(k)}$  and any bet there is a rejection parameter  $\alpha$  and a signal history  $s^t$  such that after observing  $s^t$  the MEU decision-maker exhibits ambiguity loving behavior, i.e. she prefers the bet in  $U_A(n, k)$  over the same bet in  $U_R(n, k)$ .*
- (ii) *For any rejection parameter  $\alpha \in (0, 1]$  and any bet there is  $(n, k) \in \mathbb{N}^2$ , a set of initial prior  $\mathcal{M}_0 = \Delta_\epsilon^{(k)}$  and a signal history  $s^t$  such that after observing  $s^t$  the MEU decision maker prefers the bet in  $U_A(n, k)$  over the same bet in  $U_R(n, k)$ .*

*Proof.* For both statements of the theorem I have to construct a suited signal history  $s^t$ . Without loss of generality, I consider the bet  $1_B 0$ . Accordingly, the signal history can be constructed as  $s^t = (B, \dots, B)$ . The task then is to construct an appropriate  $t$ . For  $\mathcal{M}_0 = \Delta_\epsilon^{(k)}$ , the maximal plausible prior after observing only black balls is  $(\epsilon, \dots, \epsilon, 1 - k\epsilon)$ . The ES anomaly occurs if the uniform distribution  $(1/(k+1), \dots, 1/(k+1))$  is rejected. This happens if its plausibility is lower than  $\alpha$  times the maximal plausibility. This condition, after multiplying with  $(2n+k)^t$ , reads

$$\sum_{i=0}^k (n+i)^t < (k+1)\alpha \left( \epsilon \sum_{i=0}^{k-1} (n+i)^t + (1-k\epsilon)(n+k)^t \right). \quad (8)$$

For part (i), I have to construct a suited rejection parameter  $0 < \alpha \leq 1$ . In appendix A, I demonstrate that this is equivalent with  $\alpha > \underline{\alpha}$ , where  $\underline{\alpha}$  is a number independent of  $t$ . I show  $\underline{\alpha} < 1$  and thus the existence of a problematic rejection parameter  $\alpha$ .

For part (ii),  $\alpha \in (0, 1]$  is given and the task is to construct an urn  $(n, k)$ , a set of initial prior  $\Delta_\epsilon^{(k)}$  and a signal history  $s^t$  that fulfills (8). With sensible choices, this reduces to finding a number of unknown balls  $k$ . It is obvious that the smaller  $\alpha$ , the larger must be  $k$ . In appendix A, I derive a condition for  $k$  of the form  $k > \underline{k}(\alpha)$  that is sufficient for (8).  $\square$

## 5 Graphical illustration

This section is dedicated to the graphical illustration of the general findings. The first subsection 5.1 illustrates the first part of the theorem by demonstrating the existence of a problematic rejection parameter in any given setting, here  $n = k = 1$ . The second part of the theorem is made more concrete in subsection 5.2: A given rejection parameter that is innocent in some setting ( $k = 1$ ) gets problematic in a higher dimensional setting ( $k = 2$ ).

### 5.1 First part of the theorem

The first part of the theorem can be illustrated with  $n = k = 1$ . In Figure 2 we see the time series of non-rejected initial priors  $\mathfrak{M}_0^\alpha(s^t)$  (top panel), the history of standard SEU beliefs (black line in the bottom panel) and the set of ES beliefs  $\mathcal{M}_t^\alpha(s^t)$  (shaded area in the bottom panel) for different rejection parameter  $\alpha$ . At  $t = 0$ , by definition  $\mathfrak{M}_0 = \mathcal{M}_0$  which is chosen as the full set  $\Delta_0^{(1)}$ . The initial SEU prior is  $(1/2, 1/2)$ .

The signal underlying all subfigures in figure 2 is the iteration of the sequence  $W, B, B$  (six times) and thus evidence for the theory  $\theta = 2/3$ , which corresponds to the belief  $(0, 1)$ . In the subfigures we see how this evidence is processed under different rejection parameter  $\alpha$ . With  $\alpha = 0$ , which is a limiting case not permitted by ES, no theory in  $\mathcal{M}_0$  is ever rejected. Thus  $\mathfrak{M}_0(s^t)$  remains the full set  $\mathcal{M}_0$  over the whole time (top panel). The update of the full set, however, is the full set again. This is why also  $\mathcal{M}_t$  remains constant over time. The update of the SEU decision-maker, of course, is independent of  $\alpha$  and converges to  $(0, 1)$ .

For increasing  $\alpha$  more and more theories are rejected. Note that after observing the first signal,  $s = W$ , the theory  $(1, 0)$  is the most plausible one. This is why theories in the neighborhood of  $(0, 1)$  are rejected for adequately high  $\alpha$  at time  $t = 1$  (cf. top panel) but later, as  $(0, 1)$  has become the most plausible theory, are element of the set  $\mathfrak{M}_0$  again. After observing the first two signal realizations,  $s_1 = W$  and  $s_2 = B$ , all theories are equally likely. This is reflected in the fact that  $\mathfrak{M}_0^\alpha(s^2)$  is the full set, even for the most strict rejection parameter  $\alpha = 1$ . The same effect occurs at  $t = 4$ .

Central in this paper is the question under which conditions the SEU posterior is *not* in the set  $\mathcal{M}_t$  of ES posteriors, potentially causing those problems delineated in section 3 and generally formulated in section 4.2. By definition of  $\mathcal{M}_t$ , this ES anomaly occurs




Figure 2: Illustration of the first part of the theorem. Each subfigure corresponds to a different rejection parameter  $\alpha$  and presents the history of non-rejected initial priors  $\mathfrak{M}_0^\alpha(s^t)$  (top panel), the history of standard SEU beliefs (black line in the bottom panel) and the set of ES beliefs  $\mathcal{M}_t^\alpha(s^t)$  (shaded area in the bottom panel). For  $\alpha > 0.5$ , the ES anomaly occurs at some point in time.

if and only if the initial SEU prior  $(1/2, 1/2)$  is not in  $\mathfrak{M}_0$ . We see this effect in the subfigures with  $\alpha = 0.6$ ,  $\alpha = 0.8$  and  $\alpha = 1.0$ . In the latter case, with the strictest rejection parameter possible, the set of non-rejected theories is either the full set (when the signal history is not conclusive) or the singleton with one of the extreme theories  $(1, 0)$  or  $(0, 1)$ . As a consequence, also the posterior set  $\mathcal{M}_t$  is either the full set or an extreme singleton. This shows the problem of the ES updating in a nutshell: The specific form of ES updating favors extreme beliefs, rather than 'smooth' beliefs around the SEU belief history. With  $\alpha = 1.0$ , the MEU decision-maker is positive that the unknown ball in the urn is white ( $\mathcal{M}_1 = \{(1, 0)\}$ ) after observing the first signal realizations; after the second signal realization, she is clueless ( $\mathcal{M}_2 = \Delta$ ); then positive that the color of the unknown ball is black ( $\mathcal{M}_3 = \{(0, 1)\}$ ); then clueless again before remaining perfectly convinced that the unknown ball is black. This extreme form of reevaluation is due to  $\alpha = 1$ , but even smaller values give rise to a similar behavior.

Figure 2 thus illustrates the first part of the theorem: For a given urn, it states the existence of a signal history  $s^t$  and a rejection parameter  $\alpha$  such that the setting described by those parameters features the ES anomaly. Figure 2 suggests that the ES anomaly does occur for  $\alpha \geq 0.6$ , but not for  $\alpha \leq 0.4$ . The proof of Theorem 4.1 in appendix A indeed shows that  $\underline{\alpha} = 0.5$  is the relevant threshold for  $\alpha$ .

This, in turn, implies that the setting  $n = k = 1$  is innocent for  $\alpha \leq 0.5$ . Indeed, the ES anomaly does not occur even with extreme signal histories. This is demonstrated

in Figure 3. Here, the rejection parameter  $\alpha = 0.5$  is just small enough to ensure that the SEU posterior remains in the set of ES posteriors. Observing only black balls drawn, (0,1) is always the upper bound of the non-rejected initial priors (Figure 3, upper panel) and consequently also the upper bound of the set of posteriors (lower panel). The lower bound of the non-rejected priors  $\mathfrak{M}_0$  converges to  $(1/2, 1/2)$ . Thus, the setting  $\alpha = 0.5$  exhibits no ES anomaly.




Figure 3: Illustration that  $n = k = 1$  is innocent for  $\alpha = 0.5$ .

## 5.2 Second part of the theorem

The second part of the theorem, however, shows that all rejection parameter  $\alpha > 0$  are potentially problematic. An arbitrary  $\alpha \leq 1/2$  might be innocent for the urn characterized by  $k = 1$ ; there exists, however, a generalized urn (characterized by  $n$  and  $k$ ) giving rise to the ES anomaly.

This part of the theorem is illustrated by Figure 4 with  $\alpha = 1/2$ , the signal history  $s^t = (B, B, B, B, B)$  and  $n = k = 2$ .<sup>5</sup> As usual, the beliefs under  $k = 2$  can conveniently be captured in a simplex. Each subfigure in Figure 4 corresponds to one point in time and shows the non-rejected theories  $\mathfrak{M}_0$  with the uniform distribution  $(1/3, 1/3, 1/3)$  marked with a black dot (top panel), SEU posteriors (black dots in the bottom panel) and set of ES posteriors (shaded area in the bottom panel).

We can see in the top panels that the uniform distribution  $(1/3, 1/3, 1/3)$  is not in the set of admissible theories  $\mathfrak{M}_0$  for all  $t \geq 4$ . As a result, the set of posteriors  $\mathcal{M}_t$  does not contain the SEU update. This reflects the general theorem: A rejection parameter  $\alpha$ , here  $\alpha = 0.5$ , might be innocent for certain settings, e.g.  $n = k = 1$  (cf. Figure 3). It is always possible, however, to construct a generalized urn, here  $n = k = 2$ , such that the ES anomaly occurs after observing  $s = B$  for a finite number of times.

The theorem demonstrates the problematic feature of the ES setting: It is not possible to avoid the ES anomaly when the positive rejection parameter  $\alpha > 0$  ought to be

<sup>5</sup>The number of unknown balls,  $k$ , determines the dimension of beliefs and is the relevant number. The number  $n$  only determines likelihoods and is thus of minor importance.

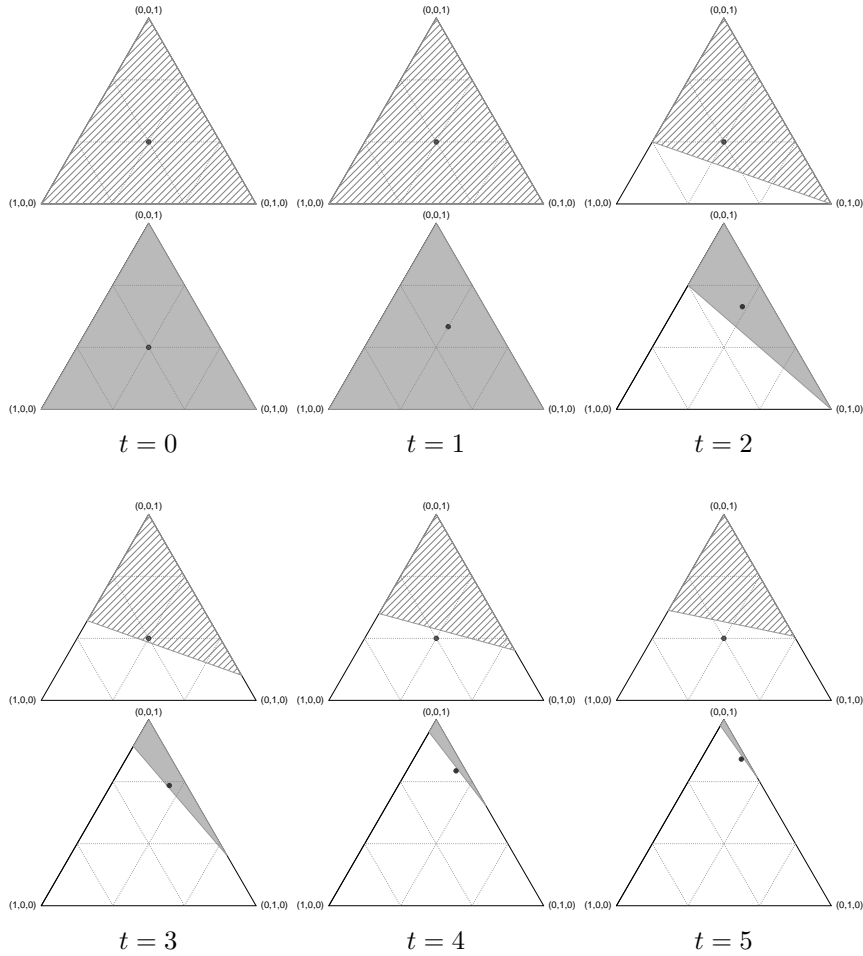


Figure 4: Illustration of the second part of the theorem with  $\alpha = 0.5$  and  $k = 2$ . Each subfigure shows for a certain point in time the set of non-rejected initial priors  $\mathfrak{M}_0^\alpha(s^t)$  (top panel), the standard SEU belief (black dot in the bottom panel) and the set of ES beliefs  $\mathcal{M}_t^\alpha(s^t)$  (shaded area in the bottom panel). The ES anomaly occurs at  $t = 4$  and  $t = 5$  when the uniform distribution is rejected.

*independent* of the decision-problem. In the next section, we offer two modifications of the ES framework in order to avoid the witch in ambiguity preferences.

## 6 Alternatives

In this final substantive section I offer two modifications of the ES framework, each of which is designed to overcome the ES anomaly. The first modification in 6.1 is basically a refinement of the ES approach and thus also involves the rejection of initial priors; the ES anomaly is avoided by defining a set of essential beliefs that are immune to rejection. The second modification, which I consider the preferable fix, has the charm of simplicity. I argue that the rejection of theories is not desirable anyway. Abstaining from the rejection of theories clearly avoids the ES anomaly. As I will demonstrate in 6.2, simple restrictions on the set of initial priors  $\mathcal{M}_0$  suffice to ensure well-behaved learning

dynamics.

### 6.1 Refinement of the rejection of theories

The first modification of the ES framework declares certain theories as unrejectable and can thus avoid the ES anomaly. The decision-maker may feel that a certain set of theories  $\mathcal{M}_0^{\text{ess}}$  is essential and thus should be immune to ex-post rejection, no matter how implausible the essential priors are in light of the signal history.

The modification of (6) is, for  $\mathcal{M}_0^{\text{ess}} \neq \emptyset$ ,<sup>6</sup> given by

$$\mathfrak{M}_0^\alpha(s^t) = \left\{ \mu_0 \in \mathcal{M}_0 \mid \text{Plaus}(\mu_0; s^t) \geq \alpha \min_{\tilde{\mu}_0 \in \mathcal{M}_0^{\text{ess}}} \text{Plaus}(\tilde{\mu}_0; s^t) \right\}. \quad (9)$$

That is, a theory  $\mu_0$  is only rejected if it fails a maximum likelihood test against *all* theories in  $\mathcal{M}_0^{\text{ess}}$ . This ensures that even for  $\alpha = 1$  all theories in  $\mathcal{M}_0^{\text{ess}}$  are updated.

In our setting, a natural candidate for such an essential theory is the uniform distribution held by the SEU decision-maker as the prior distribution. As it might be desirable to ensure that the Bayesian update is an inner point of  $\mathcal{M}_t$ , we choose  $\mathcal{M}_0^{\text{ess}} = \Delta_{\epsilon^{\text{ess}}}^{(k)}$  (cf. 4.1) with some  $0 < \epsilon^{\text{ess}} < 1/(k+1)$ .

In Figure 5 we illustrate the effect of this alternative definition on the dynamics of multiple beliefs. All parameters settings are as in Figure 2. The only difference is that the rejection of initial priors is based on (9) instead of (6).

Figure 5 shows that the ES anomaly is avoided by this alternative definition. Due to (9), the initial prior  $(1/2, 1/2)$  is an element of the set of non-rejected priors for all  $t$  and all rejection parameter  $\alpha$  (upper panels). Thus, the SEU posterior (black line in the bottom panel) is always element of set of modified ES posteriors. As the reference point for theory rejection is not the theory with the maximal plausible anymore, the rejection of theories is less strict even in cases that were not prone to the ES anomaly (compare  $\alpha = 0.2$ ,  $\alpha = 0.4$  and  $\alpha = 0.6$  in Figure 5 to those in Figure 2). The convergence behavior of the modified ES framework is the same as in the initial framework: For  $\alpha > 0$ , the set of posteriors converges to the true distribution.

The appeal of this modified definition is that it is only a slight modification of the ES framework. It avoids the ES anomaly while preserving their general structure of reevaluation of theories. In the next subsection, I argue for an alternative modification that does not rest on the reevaluation of theories ( $\alpha = 0$ ).

### 6.2 A simple alternative to the rejection of theories

If a decision-maker holds the full set of priors  $\Delta_0^{(k)}$  as the initial prior set  $\mathcal{M}_0$ , unrestricted updating (full Bayesian updating,  $\alpha = 0$ ) is not capable of reflecting learning as  $\mathcal{M}_t =$

---

<sup>6</sup>If  $\mathcal{M}_0^{\text{ess}} = \emptyset$  and thus no theory is regarded essential, the ES framework is unchanged. In that case the rejection of theories is defined by (6).






Figure 5: Refinement of the ES theory rejection with essential priors  $\mathcal{M}_0^{\text{ess}}$ . Each subfigure corresponds to a different rejection parameter  $\alpha$ . In each subfigure you find the history of non-rejected initial priors  $\mathfrak{M}_0^\alpha(s^t)$  (top panel), the history of standard SEU beliefs (black line in the bottom panel) and the set of modified ES beliefs  $\mathcal{M}_t^\alpha(s^t)$  (shaded area in the bottom panel). The critical SEU prior, the uniform distribution, is never rejected.

$\mathcal{M}_0$  for all  $t$ , irrespective of the signal history (cf., for example, the first subfigure in Figure 2). As shown by ES, this undesirable feature can be avoided by using a theory rejection parameter  $\alpha > 0$  (cf. Figure 2). This, however, in turn gives rise to a problem I coined the ES anomaly. Theorem 4.1 showed that this anomaly is a pervasive characteristic of the ES framework. The previous subsection introduced a moderate modification of this framework to circumvent this anomaly. In this subsection, I follow a different line of thought.

Reevaluation, that is the rejection of theories after observing a signal history, is not part of the standard Bayesian updating procedure. Clearly, an SEU decision-maker does *not* reevaluate its initial prior. She does not replace her initial prior by a more plausible one to update the latter. Rather, the initial prior (the uniform distribution in this paper) is the prior to be updated for all  $t$  and all conceivable signal histories. The information obtained over the learning process is reflected in the posterior and not used to reevaluate the initial prior.

I argue to keep this characteristic in the multiple prior setting. That is, I argue for defining the set of posteriors  $\mathcal{M}_t$ , under all conditions, as the update of the *full* initial prior set  $\mathcal{M}_0$ . In other words, I argue for extending the ES framework to  $\alpha = 0$  and rule out all  $\alpha > 0$ . Reevaluation of initial priors is not necessary as the information provided by the signal history is reflected in the set of posteriors. This will become apparent below.

The price for this simple fix of the ES anomaly is that I have to avoid the trivial learning dynamics mentioned above. The solution is just the restriction  $\mathcal{M}_0 \neq \Delta_0^{(k)}$ . The price of this restriction, however, is low. Please recall that the axiomatization of maxmin preferences of Gilboa and Schmeidler (1989) regards the set of beliefs as an *endogenous* component of the ambiguity averse preferences. In other words, the set of beliefs, instead of reflecting objective uncertainty, reflects how strong the ambiguity averse preferences of the decision-maker are. With the inflexible and extreme maxmin rule the set of beliefs is basically the only way to express different degrees of ambiguity aversion. In that sense, the full initial prior set  $\Delta_0^{(k)}$  would correspond to the most extreme uncertainty aversion possible. It seems not problematic to rule out this extreme form of ambiguity aversion.

Due to the non-rejection of theories, one might suspect this modification  $\alpha = 0$  to produce very large sets of posteriors. Figure 6, however, demonstrates that this is actually not the case. Even large prior sets narrow down substantially as the signals get more and more informative.




Figure 6: Simple alternative without theory rejection. Initial priors  $\mathfrak{M}_0^\alpha(s^t)$  (top panel), the history of standard SEU beliefs (black line in the bottom panel) and the set of ES beliefs  $\mathcal{M}_t^\alpha(s^t)$  (shaded area in the bottom panel).

Figure 6 shows the set of non-rejected priors  $\mathfrak{M}_0$ , the SEU posteriors and set of ES posteriors for different initial prior sets  $\mathcal{M}_0 \neq \Delta_0^{(k)}$ . The set  $\mathfrak{M}_0$  is, by definition of  $\alpha = 0$ , constant over time. The attractive features of this modification of the ES framework for applications are that (i) the procedure is simple, (ii) that the set of posteriors  $\mathcal{M}_t$  follows a similar trajectory like the SEU posterior and yet (iii) the contraction of the set reflects the increased information about the true parameter  $\theta$ . The belief dynamics in Figure 6 make the impression of being more "smooth" than those in Figure 5, let alone Figure 2. As this impression is further underpinned by the theoretical argument that the rejection of initial priors may be per se problematic, I argue to use this modification of the ES setting in applications.

## 7 Concluding discussion

The intertemporal maxmin framework of Epstein and Schneider (2007) involves the rejection of initial priors that have become implausible in light of the observed signal history. In this note, I have demonstrated that these specific belief dynamics potentially give rise to the so called ES anomaly, namely a problematic switch in ambiguity preferences. Those who apply the framework of Epstein and Schneider (2007) to model intertemporal ambiguity aversion should be aware of the potential switch in preferences and aim to avoid it.

I have offered two modifications of the ES framework. The first solution to the ES anomaly is to declare a set of priors essential and thus immune to rejection. To avoid the switch in ambiguity preferences, the uniform distribution would be such an essential prior. The second alternative, which seems the simpler and more appealing solution, abstains from the rejection of initial priors in any case and thus avoids the preference switch right from the start. As I have demonstrated, this modification leads to well-behaved belief dynamics if the set of initial priors is not the full set.

It is important to note that I have essentially focused in this comment on a reduced version of Epstein and Schneider (2007). Apart from multiple initial priors captured by the set  $\mathcal{M}_0$ , ES also allow for *multiple likelihoods*  $\mathcal{L}$ . The rejection of "theories" in ES does apply to initial priors in  $\mathcal{M}_0$  and likelihoods in  $\mathcal{L}$ . The ES anomaly already occurs in the reduced setting with  $\mathcal{L}$  being a singleton; a similar anomaly, however, may occur with multiple likelihoods when a standard likelihood is rejected over the course of the learning process. Future research ought to isolate the conditions for such a switch in detail. Both modifications of the ES framework presented in this paper, however, seem to be promising candidates to also fix the potential anomalies that may occur due to the reevaluation of multiple likelihoods.

## References

- Condie, S. and Ganguli, J. V. (2011). Ambiguity and rational expectations equilibria, *The Review of Economic Studies* **78**(3): 821–845.
- Ellsberg, D. (1961). Risk, ambiguity, and the savage axioms, *Quarterly Journal of Economics* **75**: 643–669.
- Epstein, L. G. and Schneider, M. (2003). Recursive multiple-priors, *Journal of Economic Theory* **113**(1): 1–31.
- Epstein, L. G. and Schneider, M. (2007). Learning under ambiguity, *The Review of Economic Studies* **74**(4): 1275–1303.
- Garlappi, L., Uppal, R. and Wang, T. (2007). Portfolio selection with parameter and model uncertainty: A multi-prior approach, *Review of Financial Studies* **20**(1): 41–81.
- Gilboa, I. (2009). *Theory of decision under uncertainty*, Cambridge University Press.
- Gilboa, I. and Schmeidler, D. (1989). Maxmin expected utility with non-unique prior, *Journal of Mathematical Economics* **18**(2): 141–153.

- Ju, N. and Miao, J. (2012). Ambiguity, learning, and asset returns, *Econometrica* **80**(2): 559–591.
- Leippold, M., Trojani, F. and Vanini, P. (2008). Learning and asset prices under ambiguous information, *Review of Financial Studies* **21**(6): 2565–2597.
- Maskin, E. (1979). Decision-making under ignorance with implications for social choice, *Theory and Decision* **11**(3): 319–337.
- Nishimura, K. G. and Ozaki, H. (2007). Irreversible investment and Knightian uncertainty, *Journal of Economic Theory* **136**(1): 668–694.
- Riedel, F. (2009). Optimal stopping with multiple priors, *Econometrica* **77**(3): 857–908.
- Savage, L. (1972). *The foundations of statistics*, Dover Publications.

## Appendix A Proof of the theorem

*Proof.* For both statements of the theorem I have to construct a suited signal history  $s^t$ . Without loss of generality, consider the bet  $1_B 0$ . Accordingly, I construct the signal history to generate the ES anomaly in the form of constantly observing  $s = B$ ,  $s^t = (B, \dots, B)$ . The task then is to construct an appropriate  $t$ . For  $\mathcal{M}_0 = \Delta_\epsilon^{(k)}$ , the maximal plausible prior after observing only black signals is  $(\epsilon, \dots, \epsilon, 1 - k\epsilon)$ . The uniform distribution  $(1/(k+1), \dots, 1/(k+1))$  is rejected if its plausibility is lower than  $\alpha$  times the maximal plausibility. This condition, after multiplying with  $(2n+k)^t$ , reads

$$\sum_{i=0}^k (n+i)^t < (k+1)\alpha \left( \epsilon \sum_{i=0}^{k-1} (n+i)^t + (1-k\epsilon)(n+k)^t \right). \quad (10)$$

For part (i), I have to construct, for given  $n$  and  $k$ , a suited rejection parameter  $0 < \alpha \leq 1$ . The condition for the existence of such a rejection parameter is

$$\alpha > \underline{\alpha} := \frac{\sum_{i=0}^k (n+i)^t}{(k+1) \left( \epsilon \sum_{i=0}^{k-1} (n+i)^t + (1-k\epsilon)(n+k)^t \right)}. \quad (11)$$

The existence of a rejection parameter  $\alpha$  giving rise to the ES anomaly is ensured if  $\underline{\alpha} < 1$ . Simple algebra leads to the equivalent condition

$$(1 - (k+1)\epsilon) \sum_{i=0}^{k-1} (n+i)^t < (1 - (k+1)\epsilon) k (n+k)^t. \quad (12)$$

By definition of  $\Delta_\epsilon^{(k)}$ ,  $(k+1)\epsilon < 1$ . Furthermore,  $\sum_{i=0}^{k-1} (n+i)^t < k(n+k)^t$ . This proves  $\underline{\alpha} < 1$ . In particular, there are no further restrictions on  $t$ . This implies that for every urn there is a rejection parameter  $\alpha$  such that the ambiguity loving behavior occurs already after observing *one* signal,  $t = 1$ .

For part (ii),  $\alpha \in (0, 1]$  is given and I have to construct an urn  $(n, k)$ , a set of initial prior  $\Delta_\epsilon^{(k)}$  and signal history  $s^t$  that fulfills (8). As will become clear, it is helpful to choose  $\epsilon = \frac{1}{k(k+1)\alpha}$ , for which  $\epsilon < \frac{1}{k+1}$  if  $k > \frac{1}{\alpha}$ . With that, condition (8) reads

$$\sum_{i=0}^k (n+i)^t < \frac{1}{k} \sum_{i=0}^{k-1} (n+i)^t + ((k+1)\alpha - 1)(n+k)^t. \quad (13)$$

Sufficient for this, by neglecting the first positive expression on the right hand side, is

$$\sum_{i=0}^k \left( \frac{n+i}{n+k} \right)^t < (k+1)\alpha - 1. \quad (14)$$

As  $((n+i)/(n+k))^t$  tends to 0 for  $t \rightarrow \infty$  for all  $0 \leq i < k$ , there is a  $t$  such that  $((n+i)/(n+k))^t < 1/k$  for all  $0 \leq i < k$ . Thus, a sufficient condition for the existence of a parameter  $k$  giving rise to the ES anomaly is the condition  $1 + k \cdot \frac{1}{k} < (k+1)\alpha - 1$ . This is equivalent with

$$k > \underline{k}(\alpha) = \frac{3}{\alpha} - 1. \quad (15)$$

In particular, under this condition also  $k > 1/\alpha$  and thus the  $\epsilon$ -value I defined above is actually feasible.

I have shown the existence of a problematic urn for all rejection parameter  $\alpha$ . It is intuitive that smaller  $\alpha$  make higher  $k$  necessary. It is interesting that there are no restrictions on  $n$ , the number of known black and white balls, respectively.  $\square$

# Information Acquisition under Ambiguity – Why the Precautionary Principle may Keep us Uninformed \*

Daniel Heyen      Timo Goeschl      Boris Wiesenfarth

## Abstract

Agencies charged with regulating complex risks like food safety or novel substances frequently need to take decisions in settings characterized by ambiguity, namely being unable to assign probabilities to possible outcomes of their regulatory actions. For such settings, advocates of the Precautionary Principle (PP) suggest an ambiguity averse decision rule which puts more weight on the adverse scenarios. While there is some literature on the merits of the PP given a certain degree of uncertainty, the question of the optimal level of knowledge if uncertainty can be (costly) reduced has been disregarded. Such active information gathering, however, is a central part of regulatory mandates. The purpose of this paper is to shed light on the implications a precautionary mandate has on the level of investment in active information acquisition. In our parsimonious model, a decision maker can decide on the precision of a signal which provides noisy information on a payoff-relevant parameter. Our key finding is that – contrary to intuition – a mandate of precaution often leads to less investment in active information gathering. In other words, for an important class of regulatory settings, a mandate of PP gives rise to a form of willful ignorance on behalf of the regulator.

**Keywords:** Scientific Uncertainty; Precautionary Principle; Active Information Acquisition; Regulatory Mandates; Maxmin Expected Utility.

**JEL Codes:** D81; D83; Q58.

## 1 Introduction

In the early summer of 2011, the European Union experienced an outbreak of Shiga-toxin producing *Escherichia coli* (STEC).<sup>1</sup> More than 3,100 cases of bloody diarrhea and

---

\*We are grateful to conference participants at the AUROE Young Academics Workshop 2012 in Bern, the Spring Meeting of Young Economists 2012 in Mannheim, the Annual Conference of the Society of Environmental Law and Economics 2012 at the Indiana University in Bloomington, the Summer Conference of the Association of Environmental and Resource Economists (AERE) 2012 in Asheville (NC), the Annual Conference of the European Association of Environmental and Resource Economists (EAERE) 2012 in Prague, the Workshop on Irreversible Choices 2012 in Brescia, seminar participants at the Departmental Seminar in Heidelberg, the FEEM Seminar, the Microcosm Meeting at the IASS, the FZU-ZEW Monthly Brown Bag Seminar, the Heidelberg Geoengineering Forum, and participants of the Harvard Geoengineering Summer School 2013 for helpful comments. Particular thanks goes to Gregor Betz, Valentina Bosetti, Adam Dominiak, Juergen Eichberger, Itzhak Gilboa, Simon Grant, Juan Moreno-Cruz, Joerg Oechssler, Frank Riedel, Christian Traeger, and Nicolas Treich. The first two authors acknowledge the support of the German Science Foundation, grant no. GO 1604/2-1.

<sup>1</sup>The outbreak was first incorrectly classified as EHEC and has become widely known under this label.

more than 850 of haemolytic uremic syndrome (HUS), a serious condition that can lead to kidney failure, were reported during the outbreak. There were 53 confirmed deaths (EFSA 2012). The outbreak of STEC was linked through epidemiological research to the consumption of fresh salad vegetables (BfR 2012). Because research found toxin producing E.coli on cucumbers from Spain, European Commission and German officials issued alerts and effectively required Spanish cucumbers to be withdrawn from the market (BBC 2011; BfR 2012). These events saddled Spanish vegetable growers, among others, with economic damages of several hundred millions of Euros (BBC 2011). After more research commissioned by the European Food Safety Authority (EFSA) and the German Interior Ministry, however, it became clear that officials had in all likelihood misidentified the source: Bean sprouts from a German farm, rather than contaminated Spanish cucumbers, carried the dangerous strain that gave rise to potentially lethal HUS among consumers (BfR 2012). While acknowledged as the most likely source, an entirely conclusive result on the cause of the outbreak has never been established.

The STEC incident from 2011 typifies an important recurring problem for regulators. Here, a regulator was called upon to make decisions of both public mortality and morbidity risk and economic livelihoods on the basis of ambiguous evidence on the source of risk. Not only did the regulator have to decide on which produce to ban in order to eliminate the source of risk, thus imposing damages of several hundred million Euros on its producers. The regulator also had to decide on the scale of the research effort of collecting and screening thousands of samples in order to reduce the risk of erroneously banning harmless produce while allowing harmful produce to continue to be sold. As the European Commission made clear after the episode, the relationship between the regulator's research effort and the quality of the regulatory decision was well understood.<sup>2</sup>

Situations in which a regulator needs to take a highly consequential decision based on poor, but improvable knowledge, are commonplace in today's world (Sunstein 2005a; Randall 2009; Graham 2001). This raises the question of what rules should govern these decisions. Differently put, under what mandate should the EFSA reach its decisions on the appropriate amount of research effort and on product regulation? Two approaches for writing this mandate stand out in the current discussion on regulation. One is the use of a traditional welfare-economic approach based on expected utility theory. This is the approach that underpins most forms of conventional cost-benefit analysis (CBA) (Viscusi et al. 2000). The other is the Precautionary Principle (PP). Despite lacking a clear definition (Asselt et al. 2013) and being criticized on grounds of internally inconsistency (Sunstein 2005b) and logical incoherence (Peterson 2006), the PP has been adopted by the European Commission (European Commission 2000) and gained significance of "a

---

<sup>2</sup>EU Health Commissioner John Dalli said "In future we need to see how the timing of the alerts can be closer to the actual scientific basis and proof" (EUobserver 2011), suggesting that the European Commission would have preferred a better level of information before removing a certain product from the market that might or might not be the cause of a serious health incident.

general principle of EU law” (Recuerda 2006). The PP has several connotations, all of them rooted in the presence of fundamental uncertainties that challenge traditional risk assessments. One of these interpretations of the PP requests the regulator to avoid harm even if the causal chain is subject to scientific uncertainty, and thus to prepare against unfavorable events (Zander 2010; Sunstein 2005c). A related, yet distinct, interpretation of the PP directly targets the level of information under which the regulator has to make her decision. In this widespread view, regulatory mandates based on the PP would lead to ‘more science’ (Tickner 2002) and thus a better informed regulatory decision than conventional CBA mandates (Cranor 2005; Myers and Raffensberger 2005; Martuzzi 2007; Bourg and Whiteside 2009).

How to meaningfully compare the implications of CBA and PP? Over the last decades, the literature on decision-making under uncertainty has offered alternatives to the subjective expected utility framework (Savage 1972) which underpins traditional CBA (Dasgupta and Pearce 1973; Shaw and Woodward 2008). These non-expected utility frameworks provide several coherent ‘rationalizations’ of ambiguity averse preferences (Gilboa and Schmeidler 1989; Klibanoff et al. 2005; Chateauneuf et al. 2007), forming the decision-theoretic foundations of the PP in applications (Asano 2010; Athanassoglou and Xepapadeas 2012; Basili et al. 2008; Lemoine and Traeger 2012; Millner et al. 2013; Treich 2010; Treich et al. 2013; Vardas and Xepapadeas 2010). Despite being important steps into a solid economic foundation of the PP, these contributions are not capable of shedding light on the interplay of regulatory decision and research effort that was at the heart of the EFSA tasks during the STEC outbreak. The reason for this gap is that the present PP literature rests upon a static level of information, leaving disregarded recent decision-theoretic advances in formalizing intertemporal ambiguity averse preferences under learning (Epstein and Schneider 2003, 2007).

The purpose of the present paper is to exploit the insights of these recent advances. By operationalizing the PP as maxmin expected utility preferences (Gilboa and Schmeidler 1989), thus following the main approach in the economic literature on the PP (Asano 2010; Athanassoglou and Xepapadeas 2012; Treich et al. 2013; Vardas and Xepapadeas 2010), we can build on Epstein and Schneider (2003, 2007) to analyze the PP in an intertemporal set-up with learning. In doing so, the paper demonstrates how decision problems of the EFSA type can be formalized within these decision-theoretic frameworks; likewise, it extends these frameworks by showing how decisions about active learning can be incorporated into decision-making under ambiguity. Jointly, these two steps demonstrate that intuition leads in many cases to a mistaken prediction about the relative research efforts expended by a regulator operating under a conventional CBA mandate and one operating under a PP mandate. In many settings conventional CBA, not the PP, leads to a greater effort to understand the true state of the world.

The paper proceeds in three steps. By introducing two specific stylized examples of regulatory decision-making situations, we first illustrate numerically that in one of



them the standard intuition holds, while in the other one it fails. This discrepancy of outcomes, in which the PP sometimes leads to more research, sometimes to less, justifies the development of a simple conceptual model that explains the nature of these effects. We develop such a model and demonstrate that there are two effects at work, a 'Precautionary Learning Effect' that makes a PP regulator value research more highly than a CBA regulator, and a 'Research Pessimism Effect' that has the opposite effect. Which effect dominates depends on the specific features of the decision-making situation. This means that *no* mandate will ensure that the regulatory decision is always better informed, irrespective of the circumstances. The insight that the choice of the mandate does not have a uniform impact on the regulator's level of information is not only important in its own right. The identification of the two countervailing effects also has immediate implications for the design of regulatory institutions. A setting in which the decision on information acquisition is institutionally separate from the regulatory decision can reconcile the PP with its notion of better informed decision-making.

We proceed as follows. Section 2 numerically works out two stylized examples, the STEC example from above and the approval decision for a new pesticide by the Environmental Protection Agency. Despite being structurally very similar, the examples exhibit sharply different ramifications of the PP on the research effort. Section 3 develops the fundamental decision-theoretic model that embraces both examples. The reader more interested in the formal analysis may thus skip section 2 and go straight to section 3. Section 4 demonstrates the existence of two countervailing effects of the PP on information acquisition and, by comparing their relative dependency on the payoff-structure, provides a compelling understanding of the contrarian findings in section 2. Finally, section 5 concludes.

## 2 Two leading examples

In this section we present a stylized numerical version of the STEC example delineated in the introduction. The second example, revolving around the approval decision for a novel pesticide by the Environmental Protection Agency (EPA), is very similar in its structure but features a sharply different effect of the PP on research effort.

### 2.1 Example 1: STEC and the EFSA

At the core of the European Food Safety Authority's (EFSA) decision problem is uncertainty about which vegetable is the actual cause of an STEC outbreak. For the sake of simplicity, say that the set of potential origins of the infection has been narrowed down to cucumbers and sprouts exclusively. Thus, either cucumbers or sprouts should be banned from the market to prevent harm from society.<sup>3</sup> The reason not to ban (or

---

<sup>3</sup>The simultaneous occurrence of the dangerous strain in cucumbers *and* sprouts is extremely unlikely due to separate production lines and can thus be ignored.

prematurely warn for) *both* products is the significant loss in trade value that has serious impacts on the agricultural sector, as was the case for Spanish cucumber farmers. Table 1 gives a stylized specification of societal costs that accrue under the two possible states of the world and the two different actions by the EFSA.

Table 1: Societal outcomes in millions EUR of the EFSA’s ban decision.

	Ban sprouts	Ban cucumbers
Sprouts contaminated	+500	−500
Cucumbers contaminated	−500	+500

A product ban results in losses to the agricultural sector that equal the market value of the product, here assumed to be 500 million EUR for either product. The baseline for calculating the payoffs is the health incident without any intervention. Relative to that, the ban of the contaminated vegetable renders positive payoffs equal to the health damages caused by STEC, which are assumed to be 1000 million EUR. If the EFSA makes the correct decision and bans the contaminated vegetable, final societal payoffs are thus  $+1000 - 500 = +500$  million EUR. In contrast, banning the wrong product just leads to losses in the market value of this product,  $-500$ . The numbers chosen are for illustration purposes only, but are in the order of magnitude of the actual decision problem back in 2011.<sup>4</sup>

Research on the true origin of the outbreak is crucial to increase the chance of making the right decision. It involves taking samples of cucumbers and sprouts from different regions and testing them for the specific dangerous E.coli strain. Such research is always imperfect, as was demonstrated by wrongly suspecting Spanish cucumbers based on positive E.coli tests. Importantly, however, the EFSA is not only passive recipient of research results: It can decide how many resources to invest in research and thus improving the state of knowledge about the source of the infection. The usual economic principles of this learning process are decreasing marginal improvements in the state of knowledge and / or increasing marginal costs. Obviously, the EFSA finds the optimal level of research by balancing costs and benefits of improved knowledge.

Assume that the costs of such a research program are known to the EFSA. The benefits, however, are uncertain for different reasons. The first reason why benefits are uncertain is the imperfectness of research. It is unclear whether results are conclusive, and even if so, results may be misleading. This uncertainty has to be taken into account when assessing the benefits of research. Past experience, however, gives the EFSA a quite thorough understanding of the odds of such a research program. The second source of

---

<sup>4</sup>There are de facto more options for the EFSA, namely a ban of *both* products or of *none*. Both actions would result in payoffs of 0 EUR relative to the baseline, irrespective of the true state of the world. Once the EFSA has observed evidence that, say, sprouts are more likely to be the cause of the outbreak, banning sprouts however strictly dominates the ban of both and the ban of none product. The same obviously holds if there is evidence that cucumbers ought to be banned. As a result, the actions 'ban both' and 'ban none' can be ignored right from the start.

uncertainty is more severe. At the core of the EFSA’s decision problem is uncertainty whether sprouts or cucumbers are the reason for the health emergency. The EFSA initially lacks reliable data on this question and is thus confronted with fundamental scientific uncertainty.

As mentioned in the introduction, there exist two opposed ways for the EFSA to reflect this uncertainty and accordingly make an optimal regulatory decision. The heuristic in standard CBA is to apply subjective expected utility theory and describe the (lack of) initial knowledge with the uniform prior over the two possible states. To be specific, the EFSA’s knowledge before research can be described by  $\rho_0 = 1/2$ , where  $\rho_0 = 0$  ( $\rho_0 = 1$ ) would correspond to perfect knowledge that sprouts (cucumbers) are the contaminated vegetable. Expected benefits are calculated based on this initial prior. Opposed to this standard CBA approach, many have argued to account for the lack of precise data about the problem and use robust decision rules in the precautionary spirit of ‘better be safe than sorry’ (Sunstein 2005c). A prominent example for such precautionary decision making is to describe the initial knowledge by a *set* of probability distributions and to base decision making on the worst probability scenario among them (maxmin expected utility).

Let us first analyze the research decision if the EFSA follows a standard CBA. With reasonable functional specifications on the likelihood of research outcomes and costs of different research precision level (see appendix A), we get the following expected benefits and costs.<sup>5</sup>

Table 2: Research decision by the EFSA with a CBA mandate. Numbers are millions of Euros.

Precision level	Expected benefits	Costs	Net benefits
Low	316	55	261
<b>Medium</b>	432	110	<b>322</b>
High	475	165	310

Marginal expected (gross) benefits are decreasing in the research precision and marginal costs are constant. Expected gross benefits under perfect information would be 500 million EUR since the correct product ban decision could then be taken in any case. But this status of full information is hardly attainable. Research is imperfect so that its benefits, even under high research precision, fall short of the value of perfect information. In our example, the EFSA equipped with the CBA decision rule would choose a medium research level. Increasing research further would increase expected benefits; the higher costs, however, do not justify that.

We now compare this result to the research precision choice if the EFSA followed the PP, modeled as maxmin expected utility (Gilboa and Schmeidler 1989). In contrast to the single probability distribution of the CBA approach, the EFSA equipped with

---

<sup>5</sup>The general framework to calculate these numbers will be developed in section 3

a PP mandate initially holds a *set* of priors (see Vardas and Xepapadeas 2010; Asano 2010). Without information that one state is more likely than the other, it is plausible to assume a symmetric set around the uniform distribution  $\rho_0 = 1/2$ . Let this set be  $\mathcal{M}_0 = [3/8, 5/8]$ . The first consequence of assuming a set of priors is that also ex-post knowledge (after observing research results) is a set of probability distributions (the Bayesian updates of all single priors). The EFSA’s decision under the PP is, by definition of maxmin expected utility, based on the *worst* of these ex-post probability distributions. Likewise, the optimal research decision is found by balancing expected benefits (taking into account the final ban decision for all possible research results) and costs under the worst ex-ante prior.

Table 3: Research decision by the EFSA with a PP mandate. Numbers are in millions of Euros. Expected benefits are calculated based on the worst probability scenario.

Precision level	Expected benefits	Costs	Net benefits
Low	252	55	197
Medium	405	110	295
<b>High</b>	464	165	<b>299</b>

The counterpart of Table 2 is Table 3. Not being subject to uncertainty, costs are the same irrespective of the regulatory mandate. As for the CBA mandate, marginal expected gross benefits are decreasing in the level of precision. We see however that expected benefits under the PP are systematically lower than their CBA counterparts. The simple reason is that expected benefits are calculated based on the worst probability scenario.

The main finding is that optimal research levels under the PP and CBA are different. The PP, in line with the narrative of precautionary learning, increases the research precision choice relative to the CBA mandate. Higher information costs are tolerated to improve the product ban decision. Due to the higher level of information acquisition, the regulatory ban decision is improved, decreasing the odds of further STEC infections. The EFSA operating under a standard CBA, however, considers the information costs for this gain in precision too high.

## 2.2 Example 2: The EPA decides on a novel pesticide

The second example is about the US Environmental Protection Agency (EPA) commissioned to decide whether to approve a novel pesticide. The pesticide relies on a new mechanism against a variety of pests, and suppose its improved efficacy has already been demonstrated. What has not been researched, however, is whether the new pesticide poses any threat to human health. The approval decision of the EPA critically depends on whether this is the case or not and is complicated by the fact that the pesticide builds on a novel mechanism on which no data exists.

The similarities to the EFSA’s task to ban a contaminated vegetable are striking. In both cases there is uncertainty about a payoff relevant parameter with two possible states. The EFSA is uncertain about whether sprouts or cucumbers are responsible for a serious STEC outbreak; the EPA is uncertain whether a novel pesticide has severe side-effects to human health or not. Intimately linked are two possible regulatory actions. The EFSA can either ban sprouts or cucumbers from the market; the EPA can approve or not approve the pesticide. In both examples the appropriate decision depends on the underlying true state of the world. The regulatory thus benefits from information about the true state.

To specify the EPA example, let us assume that non-approval of the pesticide is the baseline and, relative to that, approval gives rise to societal gains of 500 million USD if the substance is innocent and losses of –500 million USD if involves negative health effects.

Table 4: Societal outcomes in millions USD of the EPA’s pesticide approval decision.

	Approval	Non-approval
Pesticide harmless	+500	0
Pesticide has severe side-effects	–500	0

Another similarity of both examples is the option to undertake and shape research efforts to learn about the true state of the world. Pesticides are tested with animals to assess their health impacts on humans, and this testing can take different levels of precision. Suppose the substance can either be tested with mice, rabbits or apes. The more similar to humans the animals are, the higher the costs and the reliability of research. The research precision is a choice variable to the EPA, as it was to the EFSA in the STEC example.

The EPA operates in an uncertain environment. There is uncertainty whether research results will be conclusive and if so, how reliable the results actually are. More severe, the EPA has no data on the novel pesticide, thus facing the same kind of fundamental scientific uncertainty that complicated the EFSA decision. In the previous section we analyzed the EFSA’s research decision under two regulatory mandates, CBA and PP, and found, as expected, that the PP pushed the EFSA to undertake more research relative to CBA. The following tables demonstrate that the impact of the regulatory mandate on the EPA’s research choice is fundamentally different.

Table 5: Research decision by the EPA with a CBA mandate. Numbers are millions of USD.

Precision level	Expected benefits	Costs	Net benefits
Low	158	55	103
<b>Medium</b>	216	110	<b>106</b>
High	238	165	73

Table 5, being the counterpart of Table 2, is based on the same information structure as in the previous example. Again, we observe decreasing marginal expected benefits of research precision. Expected (gross) benefits here are significantly smaller (essentially half the corresponding numbers of the EFSA example) because gains accrue only if the pesticide is harmless, while the EFSA can realize gains in either case. As in the EFSA example, though, trading off benefits and costs of research under the CBA mandate leads to a medium level of research precision.

Table 6: Research decision by the EPA with a PP mandate. Numbers are millions of USD. Expected benefits are calculated based on the worst probability scenario.

Precision level	Expected benefits	Costs	Net benefits
<b>Low</b>	103	55	<b>48</b>
Medium	156	110	46
High	176	165	11

In the EFSA example (cf. Table 3), the PP mandate increased research efforts relative to the CBA decision rule. Table 6 shows that a PP mandate for the EPA has the contrary effect. As before, expected (gross) benefits are increasing in the level of precision with decreasing marginal benefits. Also, the focus on the worst probability scenario gives rise to consistently lower expected benefits compared to the CBA mandate. What is in stark contrast to the previous example, however, is that balancing costs and benefits of research under the PP mandate here leads to a *reduction* in research precision. In other words, the EPA equipped with the maxmin rule would, relative to a CBA mandate, accept an increase in wrong decisions about the pesticide to save information costs. This is in clear contradiction to the notion of precautionary learning.

The EFSA and the EPA examples are very similar in their structure, so why does the PP lead to such disparate information acquisition choices? The following section develops a simple decision-theoretic framework that embraces both examples and thus opens them to a joint analysis. With this framework at hand it will be possible to verify the findings of the two examples in a general algebraic way and gain a deeper understanding of the different and partially counterintuitive findings.

### 3 The general framework

In order to shed light on the interplay of regulatory mandates and information acquisition, this section develops a parsimonious framework combining two building blocks. The first building block is a two states-two actions model with a one-shot noisy signal structure whose precision is a choice variable to the decision-maker. Around this active learning model, we build maxmin expected utility preferences as the second building block. From a decision-theoretic viewpoint, our framework thus analyzes active learning under ambiguity aversion and is, to our knowledge, the first model to do so.

### 3.1 Timeline

Figure 1 presents the time structure of the model. First, nature chooses a payoff relevant parameter  $\theta$  from two possible values,  $\theta_-$  and  $\theta_+$ . The regulator is uncertain which parameter is the true  $\theta$ . In the STEC example in section 2.1, the EFSA is uncertain whether sprouts ( $\theta = \theta_+$ ) or cucumbers ( $\theta = \theta_-$ ) is the origin of the STEC outbreak; in the second example in section 2.2, the EPA is uncertain whether a new pesticide is harmless and beneficial ( $\theta = \theta_+$ ) or has severe side-effects to human health and the environment ( $\theta = \theta_-$ ).

In the second stage, the regulator decides on the precision of research activities. The result of this research realizes in stage three. Research results can be either inconclusive or they might provide evidence for one of the two parameter values being the true  $\theta$ . As usual, however, there is some likelihood that even conclusive research results are wrong. The more resources the regulator invests into research precision (stage 2), the less likely get those erroneous findings. This option to influence the information structure is called *active learning*, also known as *active information acquisition*. The main focus of this paper is to analyze the active information acquisition decision under different regulatory mandates.

After research results from stage 3 have been observed and processed to a better, yet incomplete understanding of the decision problem, the regulator has to make the final regulatory decision  $a$  (stage 4). The regulator chooses among (randomizations over) two actions,  $a^-$  and  $a^+$ , where the first is optimal if  $\theta = \theta_-$  and the second if  $\theta = \theta_+$ . In the examples, the actions available to the EFSA are the ban of sprouts ( $a = a^+$ ) or the ban of cucumbers ( $a = a^-$ ); the EPA can either approve the new pesticide ( $a = a^+$ ) or not approve it ( $a = a^-$ ). In both examples, no action dominates the other: The optimal decision depends on the true state of the world  $\theta$ .




Figure 1: Time structure of the model. The regulator chooses how much to invest into research precision (stage 2) and makes the final regulatory decision (stage 4). Both decisions have to be taken under incomplete knowledge because the regulator is uncertain which payoff relevant parameter  $\theta$  was chosen by nature in stage 1. Research results in stage 3 however provide noisy information about  $\theta$ . The precision of this information depends on the regulator's investment decision in stage 2.

In the following subsections we will explain the components of the model in detail. In the spirit of backward induction, which is the natural tool to solve the model, we start with payoffs (3.2) and competing regulatory mandates (3.3) to analyze the final regulatory decision of stage 4 (3.4). In a next step, we then explain the signal structure that represents the observation of research results in stage 3 (3.5) and finally turn to the choice of research precision in stage 2 (3.6) that is the central topic of this paper.

### 3.2 Payoff structure

Let  $\pi_-^+$  denote the payoff if the true state is  $\theta_-$  and the regulator chooses action  $a^+$ . All other notations accordingly. The assumptions  $\pi_-^- > \pi_-^+$  and  $\pi_+^+ > \pi_+^-$  then reflect that there is no dominant action. We also assume that the difference in payoffs under correct and incorrect decision be independent of the states,  $\pi_+^+ - \pi_+^- = \pi_-^- - \pi_-^+ =: a_\Delta$ , and call  $a_\Delta$  the *error cost*. It measures how large the cost of erroneous decision-making by the regulator is. The second parameter we define is the *payoff asymmetry* parameter  $\pi_\Delta = \pi_+^+ - \pi_-^- = \pi_+^- - \pi_-^+$ . This parameter captures, orthogonal to the interpretation of  $a_\Delta$ , how asymmetric the regulatory problem is in terms of the unknown parameter  $\theta$ . Parameter  $\pi_\Delta$  is non-negative if we assume, without loss of generality, that  $\pi_+^+ \geq \pi_-^-$ . As a further simplification, we normalize payoffs such that  $\pi_-^- + \pi_+^- = \pi_+^+ + \pi_-^+ = 0$ . We will see in section 3.4 that this is equivalent with normalizing the value of no information in the CBA case to 0. The initially four dimensional payoff space is now fully described by the error cost  $a_\Delta$  and the payoff asymmetry  $\pi_\Delta$ . Figure 2 gives a graphical illustration.




Figure 2: The parameter  $a_\Delta$  and  $\pi_\Delta$  fully describe the payoff structure. The error cost  $a_\Delta$  captures how wide the span between correct and wrong decision is. The payoff asymmetry  $\pi_\Delta$  is a measure of how strong payoffs differ across the two realizations of the unknown parameter  $\theta$ .

Table 7 summarizes how the two examples fit into the general framework. The payoff asymmetry parameter  $\pi_\Delta$  vanishes in the STEC example as the final payoffs are symmetric over both parameter values  $\theta_-$  and  $\theta_+$ ; the payoff only depends on whether the EFSA is able to identify the contaminated vegetable. In contrast to that,  $\pi_\Delta$  is positive in the pesticide example because the payoffs associated with a harmless pesticide,  $\theta = \theta_+$ , are consistently higher than those associated with a harmful pesticide,  $\theta = \theta_-$ . We will see in section 4 that  $\pi_\Delta$  plays a crucial role in understanding the effect of the PP on



research incentives.

Table 7: The two examples from section 2 in the language of the general framework.

Parameter	Interpretation and Numbers	
	STEC example (2.1)	Pesticide example (2.2)
$\theta_+$	Sprouts contaminated	Pesticide harmless
$\theta_-$	Cucumbers contaminated	Pesticide harmful
$a^+$	Ban sprouts	Approve pesticide
$a^-$	Ban cucumbers	Not approve pesticide
$\pi_+^+$	+500	+500
$\pi_+^-$	-500	0
$\pi_-^+$	+500	0
$\pi_-^-$	-500	-500
$a_\Delta$	1000	500
$\pi_\Delta$	0	500

### 3.3 Regulatory mandates under static information

Before we can analyze the final regulatory decision, we have to clarify our understanding of the two regulatory mandates under uncertainty we are contrasting in this paper and briefly touch upon their decision-theoretic foundation. The typical decision-problem under uncertainty, with the examples in section 2 as specific applications, is characterized by two components. First, a payoff relevant parameter  $\theta$  is unknown to the decision-maker who only knows that  $\theta \in \Theta$  with some set  $\Theta$ . Secondly, the decision-problem involves the task to choose an action from the set  $A$ . The final payoff  $\pi(\theta, a)$ , we restrict for simplicity to a risk-neutral decision-maker, depends on the true state  $\theta$  and the action taken  $a$ . The decision rules that form the basis of the two regulatory mandates differ in how they deal with the uncertainty about the parameter  $\theta$ .

The basis for the standard CBA mandate is the maximization of subjective expected utility (SEU, Savage 1972). Integral part of this decision rule is the formation of a single belief  $\mu$  over  $\Theta$  that best reflects the decision-maker's knowledge. Based on this belief, the optimal action  $a \in A$  is found by maximizing subjective expected utility,

$$\max_{a \in A} \mathbb{E}_\mu \pi(\theta, a) . \tag{1}$$

Apparent from (1) is the central role of the formation of the belief  $\mu$ . What is relevant for our analysis is that in settings with no prior information about  $\theta$  the belief  $\mu$  is usually modeled as the uniform distribution on  $\Theta$  (Principle of Insufficient Reason, cf. Gilboa 2009).

As mentioned before, standard CBA based on SEU is challenged for two reasons. First, the single probability distribution in (1) pretends a clear knowledge about the

decision-problem that seems arbitrary given the high degree of scientific uncertainty present in problems like the STEC outbreak or the new pesticide on which no data exists. A multiple prior approach is regarded as a possible solution (Athanasoglou and Xepapadeas 2012). Related to that, the second source of criticism of CBA is that in the presence of significant scientific uncertainty a precautionary approach preparing against adverse outcomes may be advised. A well-known axiomatization to accommodate both concerns is maxmin expected utility (MEU, Gilboa and Schmeidler 1989). The subjective uncertainty here is reflected in a *set* of beliefs  $\mathcal{M}$ , and every action available to the decision-maker is assessed based on the *worst* probability distribution in  $\mathcal{M}$ . In other words, the optimal action maximizes the worst expected utility,

$$\max_{a \in A} \min_{\mu \in \mathcal{M}} \mathbb{E}_{\mu} u(\pi(\theta, a)) . \quad (2)$$

Subjective expected utility (1) is the special case of (2) when  $\mathcal{M}$  is a singleton.

Due to its conceptual simplicity and sound axiomatization, MEU has been used in many regulatory settings (for instance Asano 2010; Treich et al. 2013; Vardas and Xepapadeas 2010) to reflect precautionary decision-making. In the present paper we follow this literature and always mean MEU preferences when we write 'PP'. It is important to note that the specific form of the set of beliefs  $\mathcal{M}$  is part of the preferences of the decision-maker (Gilboa and Schmeidler 1989; Etner et al. 2012). The 'size' of  $\mathcal{M}$  can be regarded as the decision-maker's degree of uncertainty aversion and has been associated with the degree of precaution.

The next section contrasts the implications of the competing regulatory mandates for the final regulatory decision.

### 3.4 Final regulatory decision

The final regulatory decision, for instance whether to approve a new pesticide or which vegetable to ban from the market, depends on the research results observed and the regulatory mandate. We start the analysis with the standard CBA approach and then contrast it with the PP.

#### 3.4.1 Standard CBA

The regulator following a CBA has started initially with a single belief about both parameter values (usually 1/2 for each of them), and thus also holds a unique posterior belief after having observed the research results (by Bayesian updating, the precise formulation will be explained in section 3.5). Since the set of states  $\Theta$  has only two elements, the posterior belief can be expressed by the single number  $\rho_1 \in [0, 1]$  that captures the (subjective) probability the regulator holds for  $\theta = \theta_+$ . Then,  $1 - \rho_1$  is the probability

that  $\theta = \theta_-$ .<sup>6</sup> In line with (1), the decision problem under CBA is to maximize expected payoffs,

$$\max_{a \in [0,1]} \rho_1 (a\pi_+^+ + (1-a)\pi_+^-) + (1-\rho_1) (a\pi_-^+ + (1-a)\pi_-^-) . \quad (3)$$

Here  $a$  is a randomization over the two actions  $a^-$  and  $a^+$ , with  $a = 1$  ( $a = 0$ ) corresponding to the pure action  $a^+$  ( $a^-$ ). Based on the payoff assumptions in 3.2, it is easy to show that the regulator strictly prefers the pure action  $a^+$  ( $a^-$ ) if and only if  $\rho_1 > 1/2$  ( $\rho_1 < 1/2$ ); under inconclusive knowledge  $\rho_1 = 1/2$  the regulator regards all randomizations  $a \in [0, 1]$  as equally reasonable.

Based on this profile of optimal actions, we can derive the *value function* that maps the posterior belief to the expected value under the optimal regulatory decision,  $\rho_1 \mapsto V(\rho_1)$ . This value function is a standard tool for the analysis of information acquisition problems (Mirman et al. 1993; Grossman et al. 1977). With the payoff space simplifications that enable us to write all payoffs in terms of the error cost  $a_\Delta$  and the payoff asymmetry  $\pi_\Delta$ , the value function for the CBA regulator reads

$$V(\rho_1) = \begin{cases} (\frac{1}{2} - \rho_1)(a_\Delta - \pi_\Delta) & \rho_1 < 1/2 \\ 0 & \rho_1 = 1/2 \\ (\rho_1 - \frac{1}{2})(a_\Delta + \pi_\Delta) & \rho_1 > 1/2 . \end{cases} \quad (4)$$

As a result of the normalization  $\pi_-^- + \pi_+^- = \pi_+^+ + \pi_-^+ = 0$  (see 3.2), the value of inconclusive knowledge  $V(1/2)$  is zero. When the posterior knowledge  $\rho_1$  approaches subjective certainty  $\rho_1 = 0$  and  $\rho_1 = 1$ , the value  $V(\rho_1)$  converges to the optimal payoffs  $\pi_-^- = (a_\Delta - \pi_\Delta)/2$  and  $\pi_+^+ = (a_\Delta + \pi_\Delta)/2$ , respectively.

Figure 3 depicts the value function for the two examples from section 2. In the STEC example, the expected value, as we move away from the uninformative posterior  $\rho_1 = 1/2$ , rises uniformly on both sides, reflecting the inherent symmetry of this problem. In the pesticide example, the left leg  $0 \leq \rho_1 < 1/2$  is flat because the regulator with such posterior belief does not approve the pesticide, and non-approval of the pesticide yields payoffs of 0 irrespective of the true state  $\theta$  and thus irrespective of the regulator's level of confidence  $\rho_1$ .

### 3.4.2 Precautionary maxmin rule

As explained in 3.3, basically two features set the PP apart from CBA. First, the regulator does not hold a unique belief  $\rho_1$  about the true state  $\theta$  but a *set* of beliefs. We denote the set of ex-post beliefs by  $\mathcal{M}_1 = [\underline{\rho}_1, \bar{\rho}_1]$ . That  $\mathcal{M}_1$  is an interval results from the assumptions we will impose on the set of initial priors  $\mathcal{M}_0$  in section 3.5. The second feature that is in sharp contrast to CBA is that the PP regulator assesses expected

<sup>6</sup>The index '1' refers to the posterior belief after having observed the signal; the index '0' is used for initial priors.

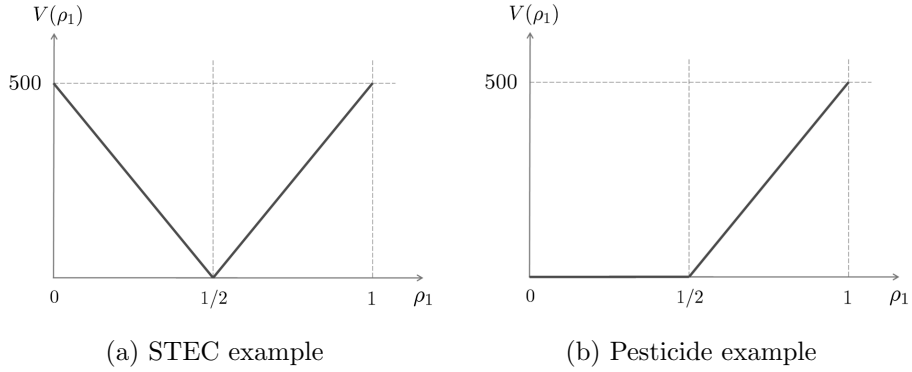


Figure 3: The value functions of the two examples. Every posterior  $\rho_1 \in [0, 1]$  is associated with the expected payoff under the optimal decision given  $\rho_1$ .

benefits of an action  $a \in A$  based on the *worst* posterior, cf. (2),

$$\max_{a \in [0,1]} \min_{\rho_1 \in \mathcal{M}_1} \rho_1 (a\pi_+^+ + (1-a)\pi_+^-) + (1-\rho_1) (a\pi_-^+ + (1-a)\pi_-^-) . \quad (5)$$

The optimal action  $a^*(\mathcal{M}_1)$ , again a randomization over actions  $a^-$  and  $a^+$ , under the PP depends on the set of posteriors  $\mathcal{M}_1$ . Only if *all* posteriors are higher (lower) than  $1/2$ , the PP regulator chooses the pure action  $a^+$  ( $a^-$ ). If however  $1/2 \in \mathcal{M}_1$ , the regulator randomizes over actions in a non-trivial way that depends on the payoff structure: In the STEC example, the PP regulator randomizes over both actions with  $1/2$  while PP regulation in the pesticide example leads to non-approval,  $a^* = 0$  (see appendix B). The clear non-approval under ambiguous knowledge, compared to indifference under CBA regulation, is one indicator of precautionary decision-making.

As before, optimal actions translate into the value function, which is now the worst expected payoff under optimal decision-making. A difference to (4) is that the value depends on the shape of the payoff structure and is a function of the *set* of posteriors  $\mathcal{M}_1 = [\underline{\rho}_1, \bar{\rho}_1]$ . We get (for details see appendix B)

$$\begin{array}{ccc}
 \boxed{a_\Delta \geq \pi_\Delta} & & \boxed{a_\Delta < \pi_\Delta} \\
 \\
 V(\mathcal{M}_1) = \begin{cases} (\frac{1}{2} - \bar{\rho}_1)(a_\Delta - \pi_\Delta) & \bar{\rho}_1 < 1/2 \\ 0 & 1/2 \in \mathcal{M}_1 \\ (\underline{\rho}_1 - \frac{1}{2})(a_\Delta + \pi_\Delta) & \underline{\rho}_1 > 1/2 . \end{cases} & & V(\mathcal{M}_1) = \begin{cases} (\frac{1}{2} - \underline{\rho}_1)(a_\Delta - \pi_\Delta) & \bar{\rho}_1 < 1/2 \\ (\frac{1}{2} - \underline{\rho}_1)(a_\Delta - \pi_\Delta) & 1/2 \in \mathcal{M}_1 \\ (\underline{\rho}_1 - \frac{1}{2})(a_\Delta + \pi_\Delta) & \underline{\rho}_1 > 1/2 . \end{cases} \\
 (6a) & & (6b)
 \end{array}$$

A difference between (6) and (4) is that the value under inconclusive posteriors depends on the payoff structure. Under moderate payoff asymmetry  $\pi_\Delta \leq a_\Delta$ , the value of inconclusive knowledge is zero as in the CBA case. If the payoff asymmetry is larger than the error cost, however, the value depends on the worst posterior  $\underline{\rho}_1$  and gets negative. In particular, the maxmin value function under inconclusive posterior knowledge  $1/2 \in \mathcal{M}_1$

is never larger than the CBA counterpart  $\rho_1 = 1/2$ , a statement that is also true for clearer posterior knowledge,  $1/2 \notin \mathcal{M}_1$ . The reason is that the value of the decision problem is always determined by the worst posterior. Note that what the worst posterior is,  $\underline{\rho}_1$  or  $\bar{\rho}_1$ , depends on the shape of the payoff structure.

### 3.5 Research results - the signal structure

The previous section gave expressions for the value of the final regulatory decision based on the posterior knowledge. This section will shed light on the formation of those posterior beliefs and explain how a priori knowledge is updated in order to process research results. Such research results are modeled as a one-shot noisy signal structure, whose precision is already fixed at this point in time (it is the choice variable of the regulator in stage 2, to be discussed in the next subsection 3.6.) We first explain Bayesian updating in the multiple prior case (3.5.1) and then develop a discrete signal structure (3.5.2) that enables us to derive closed-form solutions.

#### 3.5.1 Bayesian Updating

A signal structure is characterized by a signal space  $S$  and a likelihood function  $l : \Theta \rightarrow \Delta(S)$  that describes how likely the signals  $s \in S$  are if  $\theta$  is the true state. For better readability we write  $l_+(s)$  for  $l(\theta_+)(s)$ . Because  $\Theta$  has only two elements, every belief can be expressed as a single number in the unit interval. Let the regulator initially hold  $\rho_0 \in [0, 1]$ , her ex-ante belief that  $\theta = \theta_+$ . If she observes the signal  $s \in S$ , the prior is updated to the posterior

$$\rho_1(s, \rho_0) = \frac{\rho_0 l_+(s)}{\rho_0 l_+(s) + (1 - \rho_0) l_-(s)} . \quad (7)$$

When the regulator initially has no information about the true  $\theta \in \Theta$ , she would usually follow the *Principle of Insufficient Reason* (see for instance Gilboa 2009) and hold the uniform prior  $\rho_0 = 1/2$ . With that, formula (7) simplifies further:  $\rho_1(s, 1/2) = l_+(s)/(l_+(s) + l_-(s))$ .

The learning process of the PP regulator is similar. Instead of a single prior  $\rho_0$  she holds a set of initial beliefs  $\mathcal{M}_0$ . In analogy to the Principle of Insufficient Reason, let this set be symmetric around the uniform distribution,  $\mathcal{M}_0 = [1/2 - \delta, 1/2 + \delta]$  with the uncertainty parameter  $0 \leq \delta \leq 1/2$ . Importantly, the 'size' of  $\mathcal{M}_0$ , here fully expressed by the uncertainty parameter  $\delta$ , is not exogenous; instead, it is part of the preference structure of the regulator (Gilboa and Schmeidler 1989; Etner et al. 2012). The extreme cases are the CBA regulator, who narrows down the set to a single belief ( $\delta = 0$ ), and the most pessimistic PP regulator who does not exclude any possible prior ( $\delta = 1/2$ ). Further assumptions we are going to make will shortly rule out this extreme pessimistic case.

The learning dynamics of multiple priors follow Epstein and Schneider (2003, 2007). The set of initial beliefs is updated to the set of posteriors  $\mathcal{M}_1$  by updating every single  $\rho_0 \in \mathcal{M}_0$  according to (7),

$$\mathcal{M}_1(s) = \{\rho = \rho_1(s, \rho_0) \mid \rho_0 \in \mathcal{M}_0\} . \quad (8)$$

In other words, we assume full Bayesian updating of multiple priors.<sup>7</sup> This updating process has very intuitive features. For instance, a non-informative signal  $s$  with  $l_-(s) = l_+(s)$  results in no learning,  $\mathcal{M}_1 = \mathcal{M}_0$ . A maximal informative signal structure, on the other extreme, transforms  $\mathcal{M}_1$  into a singleton reflecting subjective certainty; for instance, observing a signal  $s$  that can only be observed in case of  $\theta = \theta_+$ ,  $l_-(s) = 0$ , gives rise to ex-post certainty that  $\theta_+$  must be the payoff relevant parameter,  $\mathcal{M}_1(s) = \{1\}$ .

A standard signal structure widely used in the literature consists of normally distributed likelihoods with some fixed variance and different means for  $\theta_-$  and  $\theta_+$ . The reciprocal of the variance is usually defined as the *precision* of the signal structure (Kihlstrom 1974; Keppo et al. 2008). The main drawback of this approach, however, is its lack of tractability. We thus design a simple discrete signal space. The justification for the discrete signal structure will be given in appendix E where we demonstrate that all main findings of this paper also hold for the continuous structure with normally distributed signals.

### 3.5.2 Discrete signal space

We consider a discrete signal space with three elements,  $S = \{s_-, s_?, s_+\}$ . Their interpretation is straightforward: The signals  $s_-$  and  $s_+$  represent research results that suggest  $\theta = \theta_-$  and  $\theta = \theta_+$  respectively, while  $s_?$  models inconclusive research outcomes. A signal structure requires a specification of the likelihood of all signals under all parameters  $\theta \in \Theta$ . Assuming symmetry over  $\theta_-$  and  $\theta_+$ , the signal structure is fully determined by specifying the likelihood of correct, erroneous and inconclusive research results

$$P_{\text{corr}} = l_-(s_-) = l_+(s_+) , P_{\text{mist}} = l_-(s_+) = l_+(s_-) , P_{\text{inconcl}} = l_-(s_?) = l_+(s_?) . \quad (9)$$

These likelihoods are functions of the signal structure's *precision*  $\tau \in T = [\tau_0, \infty)$  and sum to one for all  $\tau \in T$ .<sup>8</sup> We further assume that  $P_{\text{corr}}$ ,  $P_{\text{mist}}$  and  $P_{\text{inconcl}}$  are twice

---

<sup>7</sup>For simplicity and to ensure dynamic stable preferences (Heyen 2014), we abstain from the ex-post rejection of beliefs that is allowed for in Epstein and Schneider (2007).

<sup>8</sup>There is no natural unit for precision (Chade and Schlee 2002). The assumptions (10) are however invariant under monotone transformations  $\tau \mapsto \tau'$ ; together with the fact that we do not have to make restrictions on the cost function, see section 3.6, this shows that our results are invariant under monotone transformations in the precision parameter  $\tau$ .

continuously differentiable in  $\tau$  and that for all precision level  $\tau$

$$\lim_{\tau \rightarrow \infty} P_{\text{corr}}(\tau) = 1, P'_{\text{corr}}(\tau) > 0 \quad (10a)$$

$$\lim_{\tau \rightarrow \infty} P_{\text{mist}}(\tau) = 0, P'_{\text{mist}}(\tau) < 0 \quad (10b)$$

$$\lim_{\tau \rightarrow \infty} P_{\text{inconcl}}(\tau) = 0, P'_{\text{inconcl}}(\tau) < 0 \quad (10c)$$

$$(1 - 2\delta)P_{\text{corr}}(\tau) > (1 + 2\delta)P_{\text{mist}}(\tau) \quad (10d)$$

$$(P_{\text{corr}}P_{\text{mist}})'(\tau) < 0. \quad (10e)$$

Assumptions (10a) to (10c) are straightforward: The noisiness of the signal structure decreases in the precision of the signal structure and vanishes in the limit of perfect information. Equivalent with (10c) is  $P'_{\text{corr}} + P'_{\text{mist}} > 0$ . Assumption (10d) is equivalent with  $\underline{\rho}_1(s_+) > 1/2$  and  $\bar{\rho}_1(s_-) < 1/2$  and thus ensures that also for PP regulation clear research results lead to clear actions; otherwise information trivially has no value. In particular, assumption (10d) rules out the extreme regulator preferences  $\delta = 1/2$  right from the start. Assumption (10d) rules out, for given  $0 \leq \delta < 1/2$ , small precision levels. This is also true for the technical assumption (10e) because  $P_{\text{corr}}P_{\text{mist}}$  is non-negative and  $P_{\text{corr}}P_{\text{mist}} \rightarrow 0$  for  $\tau \rightarrow \infty$ . Thus, assumptions (10d) and (10e) basically translate into assumptions for the smallest feasible precision level  $\tau_0$ .<sup>9</sup>

With the restriction on the discrete signal structure we can give simple expressions for the updating formula (7). Let  $\rho_0$  be an initial prior. From the inconclusive research result  $s_?$  no information can be gained,  $\rho_1(s_?, \rho_0) = \rho_0$ . Observation of  $s_-$  transforms the initial prior to  $\rho_1(s_-, \rho_0) = \rho_0 P_{\text{mist}} / (\rho_0 P_{\text{mist}} + (1 - \rho_0) P_{\text{corr}})$ . This push towards  $\rho_1 = 0$  is stronger the higher is the research precision  $\tau$ . Similar, observing  $s_+$  pushes  $\rho_1(s_+, \rho_0) = \rho_0 P_{\text{corr}} / (\rho_0 P_{\text{corr}} + (1 - \rho_0) P_{\text{mist}})$  to the right. These formulas are of use both in the CBA and the PP analysis. For CBA,  $\rho_0 = 1/2$ . For the PP, every  $\rho_0 \in \mathcal{M}_0$  is updated in that way, together forming the set of posteriors  $\mathcal{M}_1$ .

### 3.6 The research precision choice

We have now all components at hand to analyze the regulator's research precision choice. The optimal research precision level is found when marginal benefits of research equal marginal costs. Two factors make the cost part very simple so that it need no further attention in our analysis: Firstly, costs are not prone to uncertainty, so that the different regulatory stance has no implications for the cost assessment; secondly, the statements we are going to make about the (marginal) benefits of research will hold for all precision levels  $\tau > 0$  and thus make further specifications of the cost function unnecessary. Sufficient are the standard assumptions of positive and non-decreasing marginal costs.

<sup>9</sup>In particular,  $\tau_0$  usually does not correspond with the "uninformative" signal structure of Radner and Stiglitz (1984).

Being thus disburdened from a sophisticated analysis of costs, we can turn our full attention to a derivation of the marginal benefits of research  $MB(\delta)$ . Here,  $\delta$  is the uncertainty parameter that determines how cautious the regulator is in the description of the initial knowledge.<sup>10</sup> The CBA regulator is the special case  $\delta = 0$ , so that statements about changes in the research behavior under different regulatory decision rules can be made by analyzing the function  $\delta \mapsto MB(\delta)$ . If, for instance,  $MB(\delta) > MB(0)$ , we can unambiguously conclude that the PP regulator chooses a higher research precision level than the CBA regulator.




Figure 4: A reduced form of the timeline. Compared to Figure 1, here the final regulatory decision is replaced by its value under optimal play, and the uncertainty under which the research precision choice has to be made is not explicit. The focus here is on the possible research results  $s \in S$ . Expected benefits of research are the sum of the values  $V(s, \delta)$  over all  $s \in S$ , weighted with their likelihoods of occurrence. The values  $V(s, \delta)$  and likelihoods  $l(s, \delta)$  both depend on the precision  $\tau$ .

The main step for the derivation of  $MB(\delta)$  is to determine the benefits  $B(\delta)$ . Figure 4 gives a reduced form of the timeline (cf. Figure 1), designed to draw the attention to the research precision choice at which the benefits of research have to be assessed. For this purpose, and as is usual in backward induction, the final regulatory decision stage has been condensed into the value of the final decision under optimal play. Accordingly, every possible research result  $s \in S$  is associated with some value  $V(s, \delta)$ . The formula for these values was given in (6): The research result  $s$ , the precision  $\tau$ , and the uncertainty parameter  $\delta$  through the initial prior set  $\mathcal{M}_0 = [1/2 - \delta, 1/2 + \delta]$  jointly determine the posterior set  $\mathcal{M}_1$ . The expected benefits of research are the sum of these values  $V(s, \delta)$  over all possible research results  $s \in S$ , weighting the different signals  $s$  with their likelihood  $l(s)$ . Because the PP regulator consistently follows MEU preferences not only in the final regulatory decision but also in the research precision choice, also the likelihoods depend on the uncertainty parameter  $\delta$ . We get for the MEU benefits of research

$$B(\delta) = l(s_+, \delta)V(s_+, \delta) + l(s_?, \delta)V(s_?, \delta) + l(s_-, \delta)V(s_-, \delta) . \quad (11)$$

Here, the likelihood  $l(s, \delta)$  for the occurrence of the research result  $s \in S$ , assessed from the perspective of the PP regulator, reads  $l(s, \delta) = (1/2 - \delta)l_+(s, \delta) + (1/2 + \delta)l_-(s, \delta)$ .

<sup>10</sup>To keep the notation simple, we do suppress as often as possible the benefit's dependency on the precision level  $\tau$ .



(Recall that  $l_+(s) = l(s|\theta = \theta_+)$  and similar for  $l_-$ ). The reason for less weight on  $l_+$  and more weight on  $l_-$  is  $V(s_+, \delta) \geq V(s_-, \delta)$ , a simple consequence of our assumption that  $\theta_+$  is the more favorable state. Together with  $l_+(s_+) = l_-(s_-) = P_{\text{corr}} > P_{\text{mist}} = l_+(s_-) = l_-(s_+)$  this implies that  $\rho_0 = 1/2 - \delta$  is the prior that minimizes the expected benefits  $\sum_{s \in S} (\rho_0 l_+(s) + (1 - \rho_0) l_-(s)) V(s, \delta)$  and is thus the worst prior relevant for PP regulation. The value in (11) conforms with the value of information under MEU developed in Heyen and Wiesenfarth (2015).

Central for the research precision choice are the marginal benefits  $\text{MB}(\delta)$ ,

$$\text{MB}(\delta) = B'(\delta) . \quad (12)$$

Note that the prime always means the derivative with respect to  $\tau$ , not with respect to  $\delta$ . The benefit function  $B(\delta)$ , as a sum of products of differentiable functions, is obviously differentiable in  $\tau$ . What makes the analysis of  $\text{MB}(\delta)$  intricate is that both the values  $V(s, \delta)$  and likelihoods  $l(s, \delta)$  depend on  $\tau$ .

The expression simplifies if we focus on the specific settings from the regulatory examples in section 2. The partially surprising effects we found there can now be formally validated. The convenient tool for comparing marginal benefits under CBA and PP is to analyze the derivative  $d/d\delta \text{MB}$ , in particular at  $\delta = 0$ .

### 3.6.1 Research precision choice in the STEC example

Due to the symmetry of the STEC example,  $V(s_-, \delta) = V(s_+, \delta)$  and  $V(s_?, \delta) = 0$ , cf. (6). Together with  $l(s_+, \delta) + l(s_-, \delta) = P_{\text{corr}} + P_{\text{mist}}$  we get that the marginal benefits of research are  $\text{MB}(\tau, \delta) = a_\Delta ((P_{\text{corr}} + P_{\text{mist}})(\rho_1(s_+, \delta) - 1/2))'$  and this reads  $a_\Delta \left( \frac{P_{\text{corr}} + P_{\text{mist}}}{2} \frac{(1/2 - \delta)P_{\text{corr}} - (1/2 + \delta)P_{\text{mist}}}{(1/2 - \delta)P_{\text{corr}} + (1/2 + \delta)P_{\text{mist}}} \right)'$ . From that we infer for the  $\delta$ -derivative at  $\delta = 0$

$$\frac{d}{d\delta} \text{MB}(\delta)|_{\delta=0} = -\frac{4a_\Delta}{(P_{\text{corr}} + P_{\text{mist}})^2} (P_{\text{corr}}^2 P'_{\text{mist}} + P_{\text{mist}}^2 P'_{\text{corr}}) . \quad (13)$$

Due to assumption (10e) this expression is positive. By continuity, the positive sign extends to a full neighborhood of  $\delta = 0$ . Thus, compared to a CBA regulator with  $\delta = 0$ , the PP regulator characterized by the uncertainty parameter  $\delta$  (as long as  $\delta$  is not too large) prefers a *higher* amount of research precision. This is the formal verification of the effect that we found in the example in section 2.1.

### 3.6.2 Research precision choice in the pesticide example

The pesticide example is characterized by  $V(s_-, \delta) = V(s_?, \delta) = 0$  because any evidence that is not clear in favor of the pesticide leads to its non-approval by the EPA. We thus get  $\text{MB}(\tau, \delta) = 2a_\Delta (l(s_+, \delta)(\rho_1(s_+, \delta) - 1/2))'$ , where we used  $\pi_\Delta = a_\Delta$ . This is

equivalent with  $\text{MB}(\tau, \delta) = a_{\Delta}((1/2 - \delta)P_{\text{corr}} - (1/2 + \delta)P_{\text{mist}})'$ . We get

$$\frac{d}{d\delta}\text{MB}(\delta) = -a_{\Delta}(P'_{\text{corr}} + P'_{\text{mist}}). \quad (14)$$

This expression is negative because  $P'_{\text{corr}} + P'_{\text{mist}} > 0$  by (10c). Interestingly, the  $\delta$ -derivative – and its sign in particular – does not depend on the uncertainty parameter  $\delta$ . Thus, and in contrast to the STEC example, the decline in the research precision in the pesticide example (cf. section 2.2) caused by the PP regulation is unambiguous.

Throughout this section we have developed a framework for the analysis of research incentives under different regulatory mandates under uncertainty. As demonstrated at the end of the section, this framework is capable of verifying the opposed effects of the PP on research behavior illustrated in the examples. Still unclear, however, is what drives the diverging results. This will be answered in the next section.

## 4 Two countervailing effects

In this section we explain how the intricate implications of the PP on research incentives can be disentangled into two effects, the *Precautionary Learning Effect* and the *Research Pessimism Effect*. While the former drives demand for research precision up relative to CBA, the latter has the opposite effect. The specific way in which both effects depend on the payoff asymmetry  $\pi_{\Delta}$  gives us the final explanation why we found so different net effects of the PP regulation on the demand for research precision in the STEC and the pesticide example (cf. section 2 as well as (13) and (14)).

What is an important step for disentangling the different drivers is the observation that the benefits  $B(\delta)$  in (11), and thus also the marginal benefits  $\text{MB}(\delta)$ , depend on the maxmin rule in two different ways, reflecting that the regulator makes two temporally separated decisions under MEU. Let us formally distinguish the role of the uncertainty parameter in those two decisions and write<sup>11</sup>

$$B(\delta_1, \delta_2) = l(s_+, \delta_1)V(s_+, \delta_2) + l(s_?, \delta_1)V(s_?, \delta_2) + l(s_-, \delta_1)V(s_-, \delta_2). \quad (15)$$

The parameter  $\delta_2$  belongs to the final regulatory decision and is closely connected to the size of the posterior set, determining through this channel the final stage value  $V(s, \delta_2)$ . In contrast to this, the parameter  $\delta_1$  describes the size of the initial prior set  $\mathcal{M}_0$  relevant for the likelihood assessment of the possible research results.

Basic calculus gives us a convenient tool to disentangle the impact of  $\delta_1$  and  $\delta_2$ . In 3.6.1 and 3.6.2, we have seen that the derivative  $d/d\delta \text{MB}(\delta)$  is a key tool for analyzing the implications of the PP on research incentives. We can write the net effect of the PP

---

<sup>11</sup>To avoid excessive notation we use the same symbol  $B$  irrespective of whether we consider the benefits as a function of one or two arguments. In that sense,  $B(\delta) = B(\delta, \delta)$ .

regulation as the sum of the partial effects,

$$\frac{d}{d\delta}\text{MB}(\delta) = \frac{\partial}{\partial\delta_1}\text{MB}(\delta, \delta) + \frac{\partial}{\partial\delta_2}\text{MB}(\delta, \delta) . \quad (16)$$

In the following it will become apparent that the distinction into  $\delta_1$  and  $\delta_2$  is justified: the first partial derivative is negative (4.2) while the second is positive (4.1).

#### 4.1 The Precautionary Learning Effect

We start with the general statement of the theorem and then provide intuition for the results.

**Theorem 4.1** (Precautionary Learning Effect). *The introduction of MEU at the final regulatory decision stage has, irrespective of the payoff structure, a positive impact on research incentives,*

$$\frac{\partial}{\partial\delta_2}\text{MB}(\delta, \delta)|_{\delta=0} > 0 . \quad (17)$$

*By continuity, this extends to a full neighborhood of  $\delta = 0$ .*

*Proof.* The proof is similar to the special case in the STEC example. See appendix C for details.  $\square$

In the remainder of this section we provide intuition for the Precautionary Learning Effect and also show why the positive sign of the  $\delta_2$ -derivative in (17) does not hold for arbitrary  $\delta > 0$ . For the sake of illustration, we restrict to the specific payoff structures of the examples. The general proof is in appendix C.

From 3.6.1 and 3.6.2 we can see that for both examples (up to a factor of 1/2 in the pesticide example) marginal benefits as a function of  $\delta_2$  can be written as

$$\text{MB}(0, \delta_2) = (P_{\text{corr}} + P_{\text{mist}})V'(s_+, \delta_2) + (P'_{\text{corr}} + P'_{\text{mist}})V(s_+, \delta_2) . \quad (18)$$

Expression (18) shows that a higher research precision  $\tau$  is productive for two reasons: Higher precision shifts the posteriors away from the inconclusive middle region, cf. (6), and thus increases the value  $V(s_+, \delta_2)$  (first term); also, higher research precision sharpens the likelihood of correct research results (second term),  $P'_{\text{corr}} + P'_{\text{mist}} > 0$  by (10c). The  $\delta_2$ -derivative of (18) reads

$$\frac{\partial}{\partial\delta_2}\text{MB}(0, \delta_2) = (P_{\text{corr}} + P_{\text{mist}})\frac{\partial}{\partial\delta_2}V'(s_+, \delta_2) + (P'_{\text{corr}} + P'_{\text{mist}})\frac{\partial}{\partial\delta_2}V(s_+, \delta_2) . \quad (19)$$

Theorem 4.1 states that the Precautionary Learning Effect is not clear-cut; the reason is that the second term in (19) is negative. An increase in the final stage uncertainty  $\delta_2$  reduces  $V(s_+, \delta_2)$  and thus dampens the positive marginal effect on the likelihood of correct research results. The first term in (19) however is positive and can thus be

regarded as the key driver behind the Precautionary Learning Effect. The reason for  $\frac{\partial}{\partial \delta_2} V'(s_+, \delta_2) > 0$  is the following: For every precision level  $V(s_+, \delta_2) < V(s_+, 0)$ , but the difference gets smaller and zero in the limit  $\tau \rightarrow \infty$  as both values converge to the value of perfect information. As a result, the increase of  $V(s_+, \delta_2)$  in the precision level  $\tau$  is steeper the higher is  $\delta_2$ . In other words: Research precision is more productive in shifting up the pessimistic value  $V(s_+, \delta_2)$ . This is the Precautionary Learning Effect.

## 4.2 The Research Pessimism Effect

The negative effect of MEU on the demand for research precision holds for all asymmetric payoff structures and, in contrast to the Precautionary Learning Effect, globally.

**Theorem 4.2** (Research Pessimism Effect). *The likelihood assessment of research results with MEU has a negative effect on research incentives,*

$$\frac{\partial}{\partial \delta_1} \text{MB}(\delta, \delta) \leq 0 \quad \text{for all } \delta. \quad (20)$$

*The derivative vanishes if and only if the payoff structure is perfectly symmetric,  $\pi_\Delta = 0$ .*

*Proof.* From (11),

$$\frac{\partial}{\partial \delta_1} \text{MB}(\delta, \delta) = -[(P_{\text{corr}} - P_{\text{mist}})(V(s_+, \delta) - V(s_-, \delta))]'. \quad (21)$$

The first factor,  $P_{\text{corr}} - P_{\text{mist}}$ , is positive by assumption (10d). Its  $\tau$ -derivative is positive by assumptions (10a) and (10b). The second factor,  $V(s_+, \delta) - V(s_-, \delta)$ , is positive by the assumption  $\pi_+^+ \geq \pi_-^-$ . Its derivative is positive for the same reason; if  $V(s_-, \delta)'$  is positive at all, it is bounded by  $V(s_+, \delta)'$ .  $\square$

In the following we will provide intuition for this result. We start from (15) and write

$$\text{MB}(\delta_1, \delta_2) = l(s_+, \delta_1)V(s_+, \delta_2)' + l(s_-, \delta_1)V(s_-, \delta_2)' + \dots \quad (22)$$

In order to focus only on relevant contributions, this omits all terms that involve the inconclusive signal  $s_?$  or marginal effects on the likelihoods. Figure 5 depicts both contributions to marginal benefits. Here, the left bar is lower in height because  $V(s_-, \delta_2)' \leq V(s_+, \delta_2)'$ . The Research Pessimism Effect results from the fact that the maxmin rule shifts likelihood weights: the CBA regulator assesses both signals as equally likely,  $l(s_-, 0) = l(s_+, 0)$ , but the PP regulator assesses the occurrence of  $s_-$  more likely,  $l(s_-, \delta_1) > l(s_+, \delta_1)$ . The simple reason is that signal likelihoods are directly associated with prior beliefs regarding the parameter values  $\theta$ , cf.  $l(s, \delta) = (1/2 - \delta)l_+(s, \delta) + (1/2 + \delta)l_-(s, \delta)$ . As a result of that shift in likelihoods, marginal benefits of research precision decrease under the PP. The Research Pessimism Effect thus deserves its name: Due to the pessimistic maxmin rule, the (marginal) value

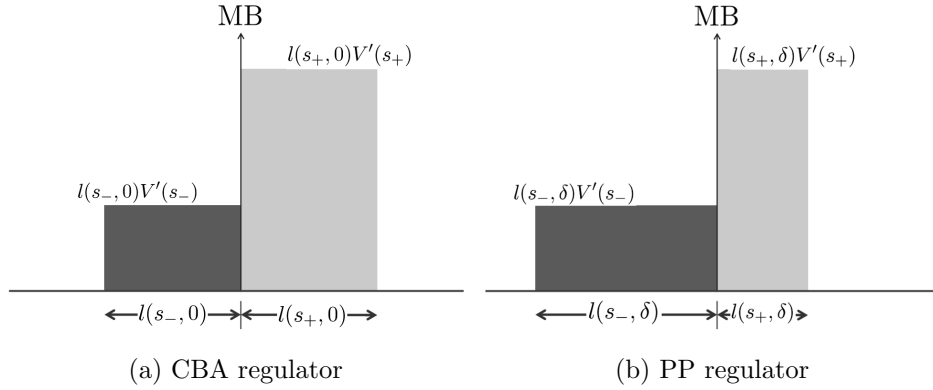


Figure 5: The intuition behind the Research Pessimism Effect. Important contributions to marginal benefits MB are  $l(s_-, \delta_1)V(s_-, \delta_2)'$  (dark shaded area) and  $l(s_+, \delta_1)V(s_+, \delta_2)'$  (light shaded area). Under the CBA regulation both contributions are equally weighted (left subfigure). The PP regulation however implies a shift in the likelihoods and thus more weight on the left bar (right subfigure). This explains why the maxmin rule in assessing the likelihoods of research results decreases the marginal benefits of research precision.

of information is assessed lower compared to a CBA regulation. This reduces the demand for research precision.

### 4.3 The net effect

For any given regulatory problem, both effects explained in the previous sections, the Research Pessimism Effect and the Precautionary Learning Effect, are present and jointly form the net effect of the PP on information acquisition. In this section we analyze how these countervailing effects depend on the payoff asymmetry  $\pi_\Delta$  of the regulatory problem. This will help us to eventually understand why we found a positive net effect of the PP on research incentives in the STEC example, in contrast to the opposite result for the pesticide regulation. To keep the analysis tractable we restrict the evaluation of the derivatives underlying the effects to the most important case  $\delta = 0$ . We show findings for some  $\delta > 0$  in appendix D.

We start with the Research Pessimism Effect. Starting from (21) we get with (6) that

$$\frac{\partial}{\partial \delta 1} \text{MB}(\delta, \delta) = \begin{cases} -\pi_\Delta [(P_{\text{corr}} - P_{\text{mist}})(\rho_1(s_+, \delta) - \bar{\rho}_1(s_-, \delta))] & a_\Delta \geq \pi_\Delta \\ -\pi_\Delta [(P_{\text{corr}} - P_{\text{mist}})(\rho_1(s_+, \delta) - \rho_1(s_-, \delta))] & a_\Delta < \pi_\Delta . \end{cases} \quad (23)$$

We already know that the Research Pessimism Effect vanishes for  $\pi_\Delta = 0$ . From there, its strength increases (piecewise) linearly in the payoff asymmetry  $\pi_\Delta$  because in both cases the two factors in the bracket and their  $\tau$ -derivatives are positive. There is a change of the slope at  $a_\Delta = \pi_\Delta$ . When evaluating the Research Pessimism Effect at  $\delta = 0$  the two slopes coincide.

The payoff asymmetry dependency of the Precautionary Learning Effect is slightly

more complex. From appendix C we get

$$\frac{\partial}{\partial \delta^2} \text{MB}(0, 0) = \begin{cases} \alpha a_\Delta & a_\Delta \geq \pi_\Delta \\ \beta_0(\pi_\Delta - a_\Delta) + \beta_1 \pi_\Delta & a_\Delta < \pi_\Delta \end{cases} \quad (24)$$

with  $\pi_\Delta$ -free positive coefficients  $\alpha$ ,  $\beta_0$  and  $\beta_1$ . Thus, the Precautionary Learning Effect is positive at  $\pi_\Delta = 0$  and remains constant until  $\pi_\Delta = a_\Delta$ . From there on the Precautionary Learning Effect increases linearly in  $\pi_\Delta$ . Figure 6 gives a graphical illustration of all effects (with the specific likelihood assumptions from appendix A and  $a_\Delta = 1000$ ).




Figure 6: The disintangled effects on research incentives and the net effect as a function of the payoff asymmetry  $\pi_\Delta$ .

What we know so far already determines the net effect of the PP mandate on information acquisition for all payoff structures with  $0 \leq \pi_\Delta \leq a_\Delta$ , with the boundaries given by the examples from section 2. At  $\pi_\Delta = 0$ , the STEEC example, the net effect is positive; at  $\pi_\Delta = a_\Delta$ , the class of the pesticide example<sup>12</sup>, the net effect is negative. In between, as the Research Pessimism Effect gets stronger, the net effect decreases and turns negative at some point  $\pi_\Delta < a_\Delta$ .

With the Precautionary Learning Effect and the Research Pessimism Effect both increasing for  $\pi_\Delta > a_\Delta$ , the net effect in this region is yet unclear. One obvious way to proceed would be to compare the slopes of both effects. It however proves easier and more general to write down  $\text{MB}(\delta)$  in this region and determine the overall derivative  $\frac{d}{d\delta} \text{MB}(\delta)$ . With  $l(s_-)V(s_-, \delta) = ((1/2 + \delta)P_{\text{corr}} - (1/2 - \delta)P_{\text{mist}})(a_\Delta - \pi_\Delta)/2$ ,  $l(s_?)V(s_?, \delta) = (1 - P_{\text{corr}} - P_{\text{mist}})\delta(a_\Delta - \pi_\Delta)$  and  $l(s_+)V(s_+, \delta) = ((1/2 - \delta)P_{\text{corr}} - (1/2 + \delta)P_{\text{mist}})(a_\Delta + \pi_\Delta)/2$  we get

$$\text{MB}(\delta) = a_\Delta((1/2 - \delta)P'_{\text{corr}} - (1/2 + \delta)P'_{\text{mist}}) \quad \text{for } a_\Delta < \pi_\Delta \quad (25)$$

and thus  $d/d\delta \text{MB}(\delta) = -a_\Delta(P'_{\text{corr}} + P'_{\text{mist}})$ . This demonstrates that for  $\pi_\Delta > a_\Delta$  the net effect, for any  $\delta \geq 0$ , is independent of the payoff asymmetry  $\pi_\Delta$ . The consequence

<sup>12</sup>The pesticide example was constructed with  $a_\Delta = 500$  so that the point  $\pi_\Delta = a_\Delta = 1000$  in Figure 6 is a scaled version of the pesticide example.

is that the overall effect of the PP on research incentives remains at the same negative level for all  $\pi_{\Delta} \geq a_{\Delta}$  (cf. Figure 6 and also its counterparts with  $\delta > 0$  in appendix D).

## 5 Concluding Discussion

The regulation of complex risks like food safety and novel substances is characterized by far-reaching consequences of erroneous decisions and, at the same time, a poor informational basis for making those decisions. In light of these challenges and with the intention to prevent harm from society, the Precautionary Principle (PP) has recently gained significant importance in the regulatory practice.

What has not received adequate attention in the literature, despite being a central task in the regulation of complex risks, is the regulator's possibility of undertaking research and thus managing her state of knowledge. Most notably, the interplay of regulatory mandates like the PP and the incentives for active information acquisition has so far gone unnoticed. The present paper sheds light on this interplay with a parsimonious decision-theoretic setting of active learning under maxmin expected utility (MEU) preferences. The latter is a common operationalization of the PP, while active learning reflects the regulator's option to choose her preferred state of information.

We find a non-trivial impact of the PP on research incentives. On the one hand, and in line with common narratives about the PP, we find the existence of a 'Precautionary Learning Effect' that induces a regulator following the PP to improve here state of knowledge relative to a standard CBA mandate. On the other hand, however, we also demonstrate the existence of a research dampening 'Research Pessimism Effect'. The total effect of the PP on information acquisition is not clear-cut and depends on the characteristics of the regulatory problem, in particular its payoff structure. The significance of this finding is that no mandate, neither CBA nor PP, always leads to better informed decision-maker. If such a better informed decision-making is regarded as a desirable and crucial feature of the regulation of complex risks, then writing an appropriate mandates is not possible without paying attention to the specific regulatory problem.

Our framework can only be the starting point for researching the interplay of the PP and research incentives. One possible extension is about the decision-theoretic foundation. Although being a standard formulation of the PP in the theoretical regulation literature, MEU is not the only definition of ambiguity averse preferences that has been suggested as a precautionary decision-rule. Even though we expect alternative approaches (Klibanoff et al. 2005; Chateauneuf et al. 2007) to give rise to similar effects on information acquisition, future research ought to clarify these issues with similar and equally tractable frameworks.

Another direction for future research is the significance of our findings for the regulatory practice. Our model is the first step towards informing the regulatory practice about institutional set-ups surrounding PP and information acquisition. Our findings

suggest that an institutional separation of the regulatory decision and research is possible, and that this separation might reconcile the PP with the notion of precautionary learning. Future research can leap from there and, by carefully analyzing real-world examples and their subtleties, make specific suggestions for the design of risk regulation.

## References

- Asano, T. (2010). Precautionary principle and the optimal timing of environmental policy under ambiguity, *Environmental and Resource Economics* **47**(2): 173–196.
- Asselt, M. B. A. v., Versluis, E. and Vos, E. (2013). *Balancing Between Trade and Risk: Integrating Legal and Social Science Perspectives*, Routledge.
- Athanassoglou, S. and Xepapadeas, A. (2012). Pollution control with uncertain stock dynamics: When, and how, to be precautionary, *Journal of Environmental Economics and Management* **63**(3): 304–320.
- Basili, M., Chateauneuf, A. and Fontini, F. (2008). Precautionary principle as a rule of choice with optimism on windfall gains and pessimism on catastrophic losses, *Ecological Economics* **67**(3): 485–491.
- BBC (2011). Spanish anger at cucumber blame, *BBC News Europe*.
- BfR (2012). EHEC outbreak 2011: Investigation of the outbreak along the food chain, *The German Federal Institute for Risk Assessment (BfR)*.
- Bourg, D. and Whiteside, K. H. (2009). Precaution and science-based environmental risk management: Complementary not contradictory, *Building Safer Communities: Risk Governance, Spatial Planning and Responses to Natural Hazards*, Landsdale, pp. 88–104.
- Chade, H. and Schlee, E. (2002). Another look at the Radner-Stiglitz nonconcavity in the value of information, *Journal of Economic Theory* **107**(2): 421–452.
- Chateauneuf, A., Eichberger, J. and Grant, S. (2007). Choice under uncertainty with the best and worst in mind: Neo-additive capacities, *Journal of Economic Theory* **137**(1): 538–567.
- Cranor, C. F. (2005). Some legal implications of the precautionary principle: Improving information-generation and legal protections, *Human and Ecological Risk Assessment: An International Journal* **11**(1): 29–52.
- Dasgupta, A. K. and Pearce, D. W. (1973). *Cost-benefit analysis*, Macmillan.
- EFSA (2012). E.coli: Rapid response in a crisis, *EFSA Features*.
- Epstein, L. G. and Schneider, M. (2003). Recursive multiple-priors, *Journal of Economic Theory* **113**(1): 1–31.
- Epstein, L. G. and Schneider, M. (2007). Learning under ambiguity, *The Review of Economic Studies* **74**(4): 1275–1303.
- Etner, J., Jeleva, M. and Tallon, J.-M. (2012). Decision theory under ambiguity, *Journal of Economic Surveys* **26**(2): 234–270.
- EUobserver (2011). Anger over EU alert system as E. coli scare hits producers, *EUobserver*.
- European Commission (2000). Communication from the Commission on the Precautionary Principle.
- Gilboa, I. (2009). *Theory of decision under uncertainty*, Cambridge University Press.
- Gilboa, I. and Schmeidler, D. (1989). Maxmin expected utility with non-unique prior, *Journal of Mathematical Economics* **18**(2): 141–153.
- Graham, J. (2001). Decision-analytic refinements of the precautionary principle, *Journal of Risk Research* **4**(2): 127–141.



- Grossman, S. J., Kihlstrom, R. E. and Mirman, L. J. (1977). A bayesian approach to the production of information and learning by doing, *The Review of Economic Studies* **44**(3): 533–547.
- Heyen, D. (2014). Learning under Ambiguity – A Note on the Belief Dynamics of Epstein and Schneider (2007), *AWI Discussion Paper 573*.
- Heyen, D. and Wiesenfarth, B. R. (2015). Informativeness of experiments for MEU – a recursive definition, *Journal of Mathematical Economics*, forthcoming.
- Keppo, J., Moscarini, G. and Smith, L. (2008). The demand for information: More heat than light, *Journal of Economic Theory* **138**(1): 21–50.
- Kihlstrom, R. (1974). A general theory of demand for information about product quality, *Journal of Economic Theory* **8**(4): 413–439.
- Klibanoff, P., Marinacci, M. and Mukerji, S. (2005). A smooth model of decision making under ambiguity, *Econometrica* **73**(6): 1849–1892.
- Lemoine, D. M. and Traeger, C. P. (2012). Tipping points and ambiguity in the economics of climate change, *NBER Working Paper 18230*, National Bureau of Economic Research.
- Martuzzi, M. (2007). The precautionary principle: in action for public health, *Occupational and environmental medicine* **64**(9): 569–570.
- Millner, A., Dietz, S. and Heal, G. (2013). Scientific ambiguity and climate policy, *Environmental and Resource Economics* **55**(1): 21–46.
- Mirman, L., Samuelson, L. and Urbano, A. (1993). Monopoly experimentation, *International Economic Review* **34**(3): 549–563.
- Myers, N. and Raffensberger, C. (2005). *Precautionary Tools for Reshaping Environmental Policy*, MIT Press.
- Peterson, M. (2006). The precautionary principle is incoherent, *Risk Analysis* **26**(3): 595–601.
- Radner, R. and Stiglitz, J. (1984). A nonconcavity in the value of information, *Bayesian models in economic theory* **5**: 33–52.
- Randall, A. (2009). We already have risk management — do we really need the precautionary principle?, *International Review of Environmental and Resource Economics* **3**(1): 39–74.
- Recuerda, M. A. (2006). Risk and reason in the European Union Law, *European Food and Feed Law Review*.
- Savage, L. (1972). *The foundations of statistics*, Dover Publications.
- Shaw, W. D. and Woodward, R. T. (2008). Why environmental and resource economists should care about non-expected utility models, *Resource and Energy Economics* **30**(1): 66–89.
- Sunstein, C. R. (2005a). Cost-benefit analysis and the environment, *Ethics* **115**(2): 351–385.
- Sunstein, C. R. (2005b). *Laws of fear: Beyond the precautionary principle*, Cambridge University Press.
- Sunstein, C. R. (2005c). The precautionary principle as a basis for decision making, *The Economists' Voice*.
- Tickner, J. (2002). *Precaution, Environmental Science, and Preventive Public Policy*, Island Press.
- Treich, N. (2010). The value of a statistical life under ambiguity aversion, *Journal of Environmental Economics and Management* **59**(1): 15–26.
- Treich, N., Rey, B. and Courbage, C. (2013). Prevention and precaution, *TSE Working Paper 445*.
- Vardas, G. and Xepapadeas, A. (2010). Model uncertainty, ambiguity and the precautionary principle: Implications for biodiversity management, *Environmental and Resource Economics* **45**(3): 379–404.

Viscusi, W., Vernon, J. M. and Harrington (2000). *Economics of Regulation and Antitrust*, MIT Press.

Zander, J. (2010). *The Application of the Precautionary Principle in Practice: Comparative Dimensions*, Cambridge University Press.

## Appendix A Functional specifications for the example

$P_{\text{corr}}(\tau) := 1 - 2/3 \cdot \exp(-\tau)$ ,  $P_{\text{mist}}(\tau) := 1/3 \cdot \exp(-\tau)$ . Low, medium and high precision level corresponds with  $\tau = 1, 2, 3$ . Costs are linear in precision  $\tau$ ,  $c(\tau) = 55\tau$ . All effects in Figure 6 are evaluated at  $\tau = 2 \ln(4/3)$ .

## Appendix B Optimal actions and the value function in the general case for the maxmin regulator

The expected value of the decision problem is

$$\rho_1(a\pi_+^+ + (1-a)\pi_+^-) + (1-\rho_1)(a\pi_-^+ + (1-a)\pi_-^-). \quad (26)$$

We define the (potentially negative)  $\hat{a} := 1/2 - \pi_\Delta/(2a_\Delta)$ . With that we find that the worst belief is, depending on which action  $a \in [0, 1]$  the regulator considers,

$$\rho_1^{\text{worst}} = \begin{cases} \underline{\rho}_1 & a > \hat{a} \\ \bar{\rho}_1 & a < \hat{a} \end{cases}. \quad (27)$$

After plugging this into the expected value (26) we can determine the optimal action. Three cases are relevant:  $\underline{\rho}_1 > 1/2$ ,  $1/2 \in \mathcal{M}_1$  and  $\bar{\rho}_1 < 1/2$ . Consider  $\underline{\rho}_1 > 1/2$ . Among the actions  $a > \hat{a}$  then clearly  $a = 1$  is best; among  $a < \hat{a}$  (if  $\hat{a} \geq 0$ ) the best action is  $\hat{a}$ . Comparing payoffs under  $a = 1$  and  $\hat{a}$  shows that the former is strictly better. Similarly in the two other cases. Together we get

$$a^*(\mathcal{M}_1) = \begin{cases} 0 & \bar{\rho}_1 < 1/2 \\ \max(0, \hat{a}) & 1/2 \in \mathcal{M}_1 \\ 1 & \underline{\rho}_1 > 1/2 \end{cases}. \quad (28)$$

The corresponding value function (6) is obtained by plugging in  $a^*(\mathcal{M}_1)$ . If  $\max(0, \hat{a}) = 0$ , the value function still depends on the worst belief  $\underline{\rho}_1$ .

## Appendix C Full formula of $\delta_2$ -derivatives

The general  $\delta_2$ -derivative is

$$\frac{\partial}{\partial \delta_2} \text{MB}(\tau, \delta, \delta) = \begin{cases} \frac{1}{X(\delta)}(A_0 a_\Delta + P(\delta)) & \pi_\Delta \leq a_\Delta \\ \frac{1}{X(\delta)Y(\delta)}(B_0(\pi_\Delta - a_\Delta) + B_1 \pi_\Delta + Q(\delta)) & \pi_\Delta > a_\Delta \end{cases}. \quad (29)$$

Here,  $X(\delta) = ((1-2\delta)P_{\text{corr}} + (1+2\delta)P_{\text{mist}})^3$  and  $Y(\delta) = ((1+2\delta)P_{\text{corr}} + (1-2\delta)P_{\text{mist}})^3$ . Both expressions are positive,  $X(\delta)$  due to assumption (10d).  $P(\delta)$  and  $Q(\delta)$  are polynomials in  $\delta$  of degree 2 and 6 respectively with  $P(0) = Q(0) = 0$ . The coefficients are

$$\begin{aligned} A_0 &= -4(P_{\text{corr}} + P_{\text{mist}})(P_{\text{corr}}^2 P'_{\text{mist}} + P_{\text{mist}}^2 P'_{\text{corr}}) > 0 \quad (\text{by (10e)}) \\ B_0 &= (P'_{\text{corr}} + P'_{\text{mist}})(P_{\text{corr}} + P_{\text{mist}})^6 > 0 \quad (\text{by (10c)}) \\ B_1 &= -4(P_{\text{corr}} + P_{\text{mist}})^4(P_{\text{corr}}^2 P'_{\text{mist}} + P_{\text{mist}}^2 P'_{\text{corr}}) > 0 \quad (\text{by (10e)}) \end{aligned}$$

Taken together, this proves the Precautionary Learning Effect: The  $\delta_2$ -derivative at  $\delta = 0$ , as stated in (17), is positive. Due to continuity, this extends to a full neighborhood of  $\delta = 0$ . The  $\delta_2$ -derivative can become negative at some  $\delta > 0$  due to the higher order terms  $P(\delta)$  and  $Q(\delta)$ .

## Appendix D The effects at $\delta > 0$

In section 4.3 we discussed all effects and their dependency on the payoff asymmetry  $\pi_\Delta$  when the effects are evaluated at  $\delta = 0$ . Figure 7 gives some intuition how these findings change when  $\delta > 0$ .

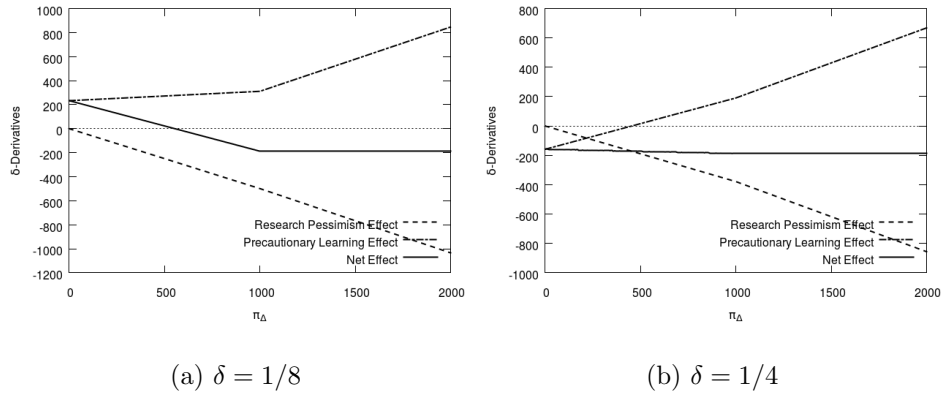


Figure 7: The analog of Figure 6 for  $\delta > 0$ .

The Research Pessimism Effect always vanishes at  $\pi_\Delta = 0$  and gets linearly stronger as  $\pi_\Delta$  increases. The slope at  $\pi_\Delta = a_\Delta$  in general changes, a fact that is obscured in the special case  $\delta = 0$  depicted in Figure 6 in the main text. The Precautionary Learning Effect, when evaluated at  $\delta > 0$ , is in general neither constant nor positive in the range  $0 \leq \pi_\Delta \leq a_\Delta$ ; unambiguous, however, is that it increases for  $\pi_\Delta > a_\Delta$ . The net effect, due to the vanishing Research Pessimism Effect, must equal the Research Pessimism Effect at  $\pi_\Delta = 0$  and can thus be negative for higher  $\delta$ . The net effect at  $\pi_\Delta = a_\Delta$  is unambiguously negative and remains at this  $\delta$ -independent level for all  $\pi_\Delta > a_\Delta$ , as was proven in section 4.3.

## Appendix E Continuous, normally distributed signals

Let the signal space be  $S = \mathbb{R}$  and

$$l_-(s, \tau) = \frac{\tau}{\sqrt{2\pi}} \exp\left(-\frac{1}{2}\tau^2(s+a)^2\right) \quad , \quad l_+(s, \tau) = \frac{\tau}{\sqrt{2\pi}} \exp\left(-\frac{1}{2}\tau^2(s-a)^2\right) \quad (30)$$

the densities. Here,  $\tau = 1/\sigma$  is the usual measure of precision when  $\sigma$  is the variance of the normal distribution. In this signal structure a signal  $s < 0$  ( $s > 0$ ) suggests  $\theta = \theta_-$  ( $\theta = \theta_+$ ). The higher  $|s|$ , the stronger is the signal. The latter feature could not be reflected in the simple discrete structure we use throughout the paper.

As before,  $l(s, \delta) = (1/2 - \delta)l_+(s) + (1/2 + \delta)l_-(s)$ . With that, the benefits of research are  $B(\delta) = \int_S l(s, \delta)V(s, \delta)ds$ , cf. (11). The distinction into  $\delta_1$  and  $\delta_2$  reads  $B(\delta_1, \delta_2) = \int_S l(s, \delta_1)V(s, \delta_2)ds$ .

The signal space splits into three parts. For  $s < \underline{s} \leq 0$ , the signal is strong enough to push all posteriors  $\rho_1 \in \mathcal{M}_1$  below  $1/2$  so that  $V(s, \delta_2)$  is determined by the first line in (6). For  $\underline{s} < s < \bar{s}$ , with  $\bar{s} \geq 0$ , the posterior set contains  $1/2$  so that the middle line in (6) applies. Finally, for  $s > \bar{s}$  all posteriors are above  $1/2$  and  $V(s, \delta)$  always equals  $(\rho_1(s, \delta_2) - 1/2)(a_\Delta + \pi_\Delta)$ . Obviously,  $\underline{s}$  and  $\bar{s}$  depend on  $\tau$  and  $\delta_2$ . For instance  $\underline{s} = \bar{s} = 0$  when  $\delta_2 = 0$ .

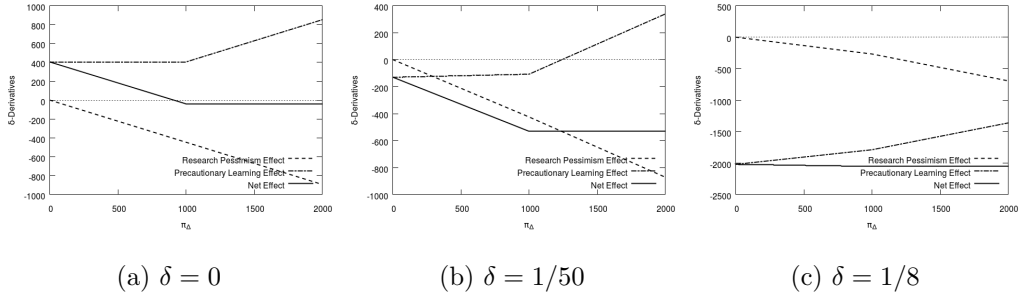


Figure 8: The net effect as a function of the payoff asymmetry  $\pi_\Delta$  for a continuous, normally distributed signal. Compare to Figure 6 and Figure 7.

Figure 8 shows the resulting effects for the continuous signal structure described above. The main figure is the left one with  $\delta = 0$ , the counterpart of Figure 6. The two other figures show the effects for two different positive  $\delta$ -levels.

Focus on the left figure first. It shows that all effects exist with the same qualitative behavior that we found in the discrete signal structure: The Research Pessimism Effect vanishes at  $\pi_\Delta = 0$  and from there evolves linearly. The Precautionary Learning Effect is positive and constant in  $0 \leq \pi_\Delta \leq a_\Delta$  and increases in  $\pi_\Delta$  once  $\pi_\Delta > a_\Delta$ . Most importantly, the net effect is positive at  $\pi_\Delta = 0$ , negative at  $\pi_\Delta = a_\Delta$ , and remains constant for  $\pi_\Delta > a_\Delta$ .

Moreover, the two other figures show that also for the continuous signal structure the implications of  $\delta > 0$  is intricate with the potential to kill the Precautionary Learning Effect while the Research Pessimism Effect is very robust. The reaction to increasing levels of the uncertainty parameter  $\delta$  seems to be more pronounced in the normal distribution structure. Interestingly, however, the net effect is still constant in  $\pi_\Delta$  once  $\pi_\Delta > a_\Delta$ . This constant level, however – and this is the only apparent qualitative difference between discrete and continuous signal structure – is here a function of the point  $\delta$  at which the derivative is evaluated.

In sum, Appendix E has shown that the analytical tractable discrete signal structure captures the relevant features of the model surprisingly well.

## Conclusion

The articles in this dissertation contributed to two important topics in theoretical environmental economics. The first topic is the innovation of environmental technologies with a particular focus on climate engineering. The second topic this dissertation contributed to is regulation under uncertainty and active information acquisition. Information and research plays a crucial role in each of them, thus connecting both parts.

The first two articles focused on deployment conflicts surrounding climate engineering (CE) and identified novel mechanisms within and across generations how technology deployment may impact on innovation incentives. Not only did these articles provide specific insights for ongoing discussions about climate engineering, each of them also delivered parsimonious and powerful frameworks as a starting point for future research.

The dissertation moved on, motivated from solar radiation management (SRM) as a typical high risk and novel technology, to the regulation under uncertainty and the role of the Precautionary Principle (PP). As a contribution to the decision-theoretic literature and valuable for a wide range of applications in environmental economics, two articles established sound concepts of the value of information and the learning dynamics under maxmin expected utility (MEU). Based on this solid foundation, the fifth article developed a framework combining MEU with active learning and argued that this framework is capable of capturing two important regulatory tasks under fundamental uncertainty, namely final regulatory decisions like approval decisions on novel substances as well as research decisions to improve the regulator's state of knowledge. Most importantly, this framework enabled an analysis of the interplay of two notions of the PP: While the formalization of the PP as ambiguity aversion has some appeal, the fifth article identified a conflict with a second, information-based interpretation of the PP.

The importance of research and information is evident in both parts of the thesis. On the one hand, as part of 'research and development' (R&D), research plays an important role in the innovation of new technologies. A typical characteristic of information relevant in this context is its public good nature: once it is in the world, it typically cannot be taken back – neither within nor across generations. On the other hand, and in a normative dimension, information has substantial value. Because a sound knowledge basis is regarded crucial for decision-making under fundamental uncertainty, information acquisition itself has been associated with precaution. Taken together, this dissertation has re-emphasized the remarkable and central role of information in the economic context (Arrow 1996; Stiglitz 2000, 2002).

This dissertation might serve as the starting point for future research. In a first dimension, and related to the first article, the question under what conditions a generation is willing to equip the future with a technology applies to a much broader field than just climate engineering. The interesting feature of technology in this context is that the provision of technology, in sharp contrast to intergenerational transfers in-kind (Bruce and Waldman 1991), opens choices of the recipient. This intergenerational technology provision is particularly interesting when preferences are unstable and uncertain (Krysiak and Krysiak 2006; Heal and Krüström 2002; Le Kama and Schubert 2004).

Another line of future research may be to further improve the analysis of innovation incentives and to account for the uncertainties surrounding climate engineering that the articles on strategic issues in this thesis have abstracted from. On the one hand, there may be a direct effect of uncertainty on innovation (Hart 2009). On the other hand, uncertainty also has ramifications for how technologies will be used; and this, as established by the second article of this dissertation, in turn impacts on R&D incentives. Analyzing the interplay of both uncertainty effects on innovation seems a promising research task. This may be particularly true in light of the quite unusual form of low probability and high impact 'fat tail' uncertainty that characterizes both climate change and SRM (Weitzman 2011; Pindyck 2011; Nordhaus 2011).

This specific fat tail uncertainty also points to a recurrent characteristic of risk regulation. Many regulatory tasks are characterized by so-called 'risk-risk trade-offs' in which the regulator, lacking a safe option, has to make hard choices in the face of high uncertainty (Wiener and Graham 2009; Sunstein 2002; Gray and Hammitt 2000). Risk assessments are substantially challenged if cost-benefit comparisons are questioned on both ends of the spectrum: "How should the bad fat tail of climate uncertainty be compared with the bad fat tails of various proposed solutions such as nuclear power, geoengineering, or carbon sequestration in the ocean floor?" (Weitzman 2009). The specific risk-risk trade-off surrounding SRM has not received adequate attention in this dissertation.

The presence of risk-risk trade-offs about climate engineering, particularly SRM, also point to another related angle for future research. While this thesis relied on SRM to motivate the idea of the PP, it did – in order to keep the second part succinct – not play the insights gained back to inform the regulation of SRM. But there is an obvious gap to be filled. As mentioned before, the idea to regulate SRM under the PP has been raised frequently (Bodansky 1996; Reynolds and Fleurke 2013). Also, and demonstrating that SRM is a typical example for the framework developed in the fifth article, active information acquisition regarding the beneficial and harmful characteristics of SRM is likely to dominate the coming decade (Parson and Keith 2013; Parker 2014; Keith et al. 2014). Assuming the responsibility for SRM deployment and research was given into the hands of one regulating agency, for instance under the roof of the United Nations (Bo-

dansky 2011): How would a precautionary mandate for this agency shape the decisions surrounding SRM?

A solid analysis of this question is left for future research, but preliminary insights can already be sketched. Assuming that the hypothetical agency for SRM regulation were equipped with a PP in the form of MEU to ensure a 'margin of safety' in SRM decisions, then this may lead to problematic effects on the regulator's level of information. The fifth article of this thesis demonstrated that the asymmetry in the payoff structure of the regulatory problem plays a crucial role in shaping the value of research. Accordingly, it seems possible that the SRM agency operating under the PP may undertake more research before making the decision whether marine cloud brightening or stratospheric aerosol injection is the preferable SRM method and thus should receive more R&D funding; this is because in this question, the regulator can be expected to be ex-ante indifferent between both options, precluding the information dampening 'research pessimism effect' established in this thesis. In contrast, however, if the agency had to make the decision on SRM deployment, balancing the clear benefits with uncertain side-effects, the picture may change. The high asymmetry in this setting may induce the regulator to learn less than an ordinary risk management agency would do. And this, again, would not be compatible with the notion of the PP that calls for more and better science in decisions surrounding SRM. Thus SRM can be expected to join the list of examples suggestive of shortcomings in the formalization of the PP.

Hence we should revisit what ought to be an integral part of the Precautionary Principle. The formalization of the PP by means of Knightian uncertainty and ambiguity aversion captures important notions of precaution, but neglects others. The 'margin of safety' notion of the PP can also be read as to prepare against unfavorable surprises. There are plenty examples in which a decision-maker is temporarily ignorant of relevant contingencies, for instance unknown feedback-mechanisms in a complex system (Oppenheimer et al. 2008). Closing the circle to the first sentence of this dissertation, global warming itself is a typical example of temporary ignorance. For decades, mankind filled up the atmospheric stock of greenhouse gases, fully unaware of the negative effects this would bring about.

This suggests an alternative definition of the PP. In a new line of thinking, fundamental uncertainty is not described by ambiguity but, on a deeper level, as state-space uncertainty and thus the possibility of unawareness (Walker 2014; Heifetz et al. 2006; Grant and Quiggin 2013b). Based on these axiomatizations, Grant and Quiggin (2013a) formulate the PP as an heuristic that requires the decision-maker to only consider actions that are not prone to unfavorable surprises. Because the framework in Grant and Quiggin (2013a) also entails considerable restrictions on active information possibilities, and in light of the main message of the fifth article in this thesis, it however remains to be seen how active learning fits into this novel and promising definition of the PP.

This dissertation had its starting point in the danger of climate change. Novel technologies, running under the label of climate engineering, offer remedy; and yet, they pose threats themselves. This is the process of adding ever more layers of complexity to our 'risk society' that the sociologist Ulrich Beck speaks about when he defines risk as "a systematic way of dealing with hazards and insecurities induced and introduced by modernisation itself" (Beck 1992). Maybe society decides that climate engineering is a valuable and controllable tool; maybe it decides to ban it. But the decision on that, whatever it may look like, should be a rational one. This defines the task of modern risk regulation. May the insights derived in this thesis help to build a sound and comprehensive risk management in which sophisticated uncertainty frameworks and CBA's unerring view on trade-offs come together.

## References

- Arrow, K. J. (1996). The economics of information: An exposition, *Empirica* **23**(2): 119–128.
- Beck, U. (1992). *Risk society: Towards a new modernity*, Sage.
- Bodansky, D. (1996). May we engineer the climate?, *Climatic Change* **33**(3): 309–321.
- Bodansky, D. (2011). Governing climate engineering: Scenarios for analysis, *SSRN Scholarly Paper ID 1963397*, Social Science Research Network, Rochester, NY.
- Bruce, N. and Waldman, M. (1991). Transfers in kind: Why they can be efficient and nonpaternalistic, *The American Economic Review* **81**(5): 1345–1351.
- Grant, S. and Quiggin, J. (2013a). Bounded awareness, heuristics and the precautionary principle, *Journal of Economic Behavior & Organization* **93**: 17–31.
- Grant, S. and Quiggin, J. (2013b). Inductive reasoning about unawareness, *Economic Theory* **54**(3): 717–755.
- Gray, G. M. and Hammitt, J. K. (2000). Risk/risk trade-offs in pesticide regulation: An exploratory analysis of the public health effects of a ban on organophosphate and carbamate pesticides, *Risk Analysis* **20**(5): 665–680.
- Hart, R. (2009). Bad eggs, learning-by-doing, and the choice of technology, *Environmental and Resource Economics* **42**(4): 429–450.
- Heal, G. and Kriström, B. (2002). Uncertainty and climate change, *Environmental and Resource Economics* **22**(1-2): 3–39.
- Heifetz, A., Meier, M. and Schipper, B. C. (2006). Interactive unawareness, *Journal of Economic Theory* **130**(1): 78–94.
- Keith, D. W., Duren, R. and MacMartin, D. G. (2014). Field experiments on solar geoengineering: report of a workshop exploring a representative research portfolio, *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* **372**(2031): 20140175.
- Krysiak, F. C. and Krysiak, D. (2006). Sustainability with uncertain future preferences, *Environmental and Resource Economics* **33**(4): 511–531.
- Le Kama, A. A. and Schubert, K. (2004). Growth, environment and uncertain future preferences, *Environmental and Resource Economics* **28**(1): 31–53.
- Nordhaus, W. D. (2011). The economics of tail events with an application to climate change, *Review of Environmental Economics and Policy* **5**(2): 240–257.



- Oppenheimer, M., O'Neill, B. C. and Webster, M. (2008). Negative learning, *Climatic Change* **89**(1-2): 155–172.
- Parker, A. (2014). Governing solar geoengineering research as it leaves the laboratory, *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **372**(2031): 20140173.
- Parson, E. A. and Keith, D. W. (2013). End the deadlock on governance of geoengineering research, *Science* **339**(6125): 1278–1279.
- Pindyck, R. S. (2011). Fat tails, thin tails, and climate change policy, *Review of Environmental Economics and Policy* **5**(2): 258–274.
- Reynolds, J. L. and Fleurke, F. (2013). Climate engineering research: A precautionary response to climate change, *Carbon & Climate Law Review* **2013**(2): 101–107.
- Stiglitz, J. E. (2000). The contributions of the economics of information to twentieth century economics, *Quarterly Journal of Economics* **115**(4): 1441–1478.
- Stiglitz, J. E. (2002). Information and the change in the paradigm in economics, *The American Economic Review* **92**(3): 460–501.
- Sunstein, C. R. (2002). *Risk and reason: Safety, law, and the environment*, Cambridge University Press.
- Walker, O. (2014). Unawareness with “possible” possible worlds, *Mathematical Social Sciences* **70**: 23–33.
- Weitzman, M. L. (2009). On modeling and interpreting the economics of catastrophic climate change, *Review of Economics and Statistics* **91**(1): 1–19.
- Weitzman, M. L. (2011). Fat-tailed uncertainty in the economics of catastrophic climate change, *Review of Environmental Economics and Policy* **5**(2): 275–292.
- Wiener, J. B. and Graham, J. D. (2009). *Risk vs. risk: Tradeoffs in protecting health and the environment*, Harvard University Press.