

University of Heidelberg

Department of Economics



Discussion Paper Series | No. 605

**Risk Assessment under Ambiguity:
Precautionary Learning vs. Research Pessimism**

Daniel Heyen, Timo Goeschl, and Boris Wiesenfarth

December 2015

Risk Assessment under Ambiguity: Precautionary Learning vs. Research Pessimism*

Daniel Heyen
Department of Economics
Heidelberg University

Timo Goeschl
Department of Economics
Heidelberg University

Boris Wiesenfarth
Department of Economics
Heidelberg University

November 2, 2015

Abstract

Agencies charged with regulating complex risks such as food safety or novel substances frequently need to take decisions on risk assessment and risk management under conditions of ambiguity, i.e. where probabilities cannot be assigned to possible outcomes of regulatory actions. What mandates should society write for such agencies? Two approaches stand out in the current discussion. One charges the agency to apply welfare economics based on expected utility theory. This approach underpins conventional cost-benefit analysis (CBA). The other requires that an ambiguity-averse decision-rule – of which maxmin expected utility (MEU) is the best known – be applied in order to build a margin of safety in accordance with the Precautionary Principle (PP). The contribution of the present paper is a relative assessment of how a CBA and a PP mandate impact on the regulatory task of risk assessment. In our parsimonious model, a decision maker can decide on the precision of a signal which provides noisy information on a payoff-relevant parameter. We find a complex interplay of MEU on information acquisition shaped by two countervailing forces that we dub 'Precautionary Learning' and 'Research Pessimism'. We find that – contrary to intuition – a mandate of PP rather than CBA will often give rise to a less informed regulator. PP can therefore lead to a higher likelihood of regulatory mistakes, such as the approval of harmful new substances.

Keywords: scientific uncertainty; ambiguity; learning; risk assessment; precautionary principle; active information acquisition; regulatory mandates.

JEL codes: D81, D83, Q58.

*We are grateful to Gregor Betz, Valentina Bosetti, Adam Dominiak, Juergen Eichberger, Itzhak Gilboa, Simon Grant, Juan Moreno-Cruz, Joerg Oechssler, Frank Riedel, Bernard Sinclair-Desgagné, Christian Traeger, Nicolas Treich, and participants at the conferences AUROE 2012 Bern, Spring Meeting of Young Economists 2012 Mannheim, AERE 2012 Asheville, EAERE 2012 Prague, Workshop on Irreversible Choices 2012 in Brescia, the Harvard Geoengineering Summer School 2013 and the seminars at the Department of Economics Heidelberg, FEEM, IASS, ZEW, Heidelberg Geoengineering Forum for their helpful comments and suggestions. The first two authors acknowledge the support of the German Science Foundation, grant no. GO 1604/2.

1 Introduction

In the early summer of 2011, the European Union experienced an outbreak of Shiga-toxin producing *Escherichia coli* (STEC).¹ More than 3,100 cases of bloody diarrhea and 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.²

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

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

²EU 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.

approach that underpins most forms of conventional cost-benefit analysis (CBA) (Viscusi et al. 2000).³ 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 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; Barrieu and Sinclair-Desgagné 2006; Gollier and Treich 2003; Gollier et al. 2000). 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

³Also this traditional approach is prone to subtleties, cf. Martin and Pindyck (2015).

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.⁴ The reason not to ban (or 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.

	Sprouts contaminated	Cucumbers contaminated
Ban sprouts	+500	−500
Ban cucumbers	−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 2011.⁵

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

⁴The simultaneous occurrence of the dangerous strain in cucumbers *and* sprouts is extremely unlikely due to separate production lines and can thus be ignored.

⁵There 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.

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 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.⁶

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
Medium	432	110	322
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 mil-

⁶The general framework to calculate these numbers will be developed in section 3

lion 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 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
High	464	165	299

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. See table 4.

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

	Pesticide harmless	Pesticide harmful
Approval	+500	-500
Non-approval	0	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
Medium	216	110	106
High	238	165	73

Table 5, 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
Low	103	55	48
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 in order 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 θ from two possible values, θ_- and θ_+ . The regulator is uncertain which parameter is the true θ . 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 θ . 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*. It is an example of This decision stage 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 θ .

In a regulatory dimension, stage 2 and stage 4 correspond to *risk assessment* and *risk management* (Sunstein 2002; Haimes 2005), respectively. Correspondingly, the economic theory of uncertainty and information differentiates *informational* and *terminal* moves (Hirshleifer and Riley 1992).

In the following subsections we will explain the components of the model in detail.




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 θ was chosen by nature in stage 1. Research results in stage 3 however provide noisy information about θ . The precision of this information depends on the regulator's investment decision in stage 2.

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 π_-^+ denote the payoff if the true state is θ_- 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_Δ 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_Δ , how asymmetric the regulatory problem is in terms of the unknown parameter θ . Parameter π_Δ 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_Δ and the payoff asymmetry π_Δ . Figure 2 gives a graphical illustration.

Table 7 summarizes how the two examples fit into the general framework. The payoff asymmetry parameter π_Δ vanishes in the STEC example as the final payoffs are symmetric over both parameter values θ_- and θ_+ ; the payoff only depends on whether the EFSA is able to identify the contaminated vegetable. In contrast to that, π_Δ is positive




Figure 2: The parameter a_Δ and π_Δ fully describe the payoff structure. The error cost a_Δ captures how wide the span between correct and wrong decision is. The payoff asymmetry π_Δ is a measure of how strong payoffs differ across the two realizations of the unknown parameter θ .

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 π_Δ 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)
θ_+	Sprouts contaminated	Pesticide harmless
θ_-	Cucumbers contaminated	Pesticide harmful
a^+	Ban sprouts	Approve pesticide
a^-	Ban cucumbers	Not approve pesticide
π_+^+	+500	+500
π_+^-	-500	0
π_-^-	+500	0
π_-^+	-500	-500
a_Δ	1000	500
π_Δ	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 θ is unknown to the decision-maker who only knows that $\theta \in \Theta$ with some set Θ . 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 θ 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 θ .

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 μ over Θ 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,

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

Apparent from (1) is the central role of the formation of the belief μ . What is relevant for our analysis is that in settings with no prior information about θ the belief μ is usually modeled as the uniform distribution on Θ (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 (Athanassoglou 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 which vegetable to ban from the market or whether to approve a new pesticide, 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 Θ 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 $\theta = \theta_-$.⁷ 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_Δ and the payoff asymmetry π_Δ , 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 ρ_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.

⁷The index '1' refers to the posterior belief after having observed the signal; the index '0' is used for initial priors.

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 θ and thus irrespective of the regulator's level of confidence ρ_1 .




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 ρ_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 ρ_1 about the true state θ 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 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)

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

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 θ is the true state. For better readability we write $l_+(s)$ for $l(\theta_+)(s)$. Because Θ 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

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

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 ρ_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 δ , 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.⁸ 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 θ_+ 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 θ_- and θ_+ . 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

⁸For 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).

$\theta = \theta_-$ 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 θ_- and θ_+ , 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 $\tau \in T$.⁹ We further assume that P_{corr} , P_{mist} and P_{inconcl} are twice continuously differentiable in τ and that for all precision level τ

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

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

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

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

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

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 τ_0 .¹⁰

With the restriction on the discrete signal structure we can give simple expressions for the updating formula (7). Let ρ_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 τ . 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 .

⁹There 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 τ .

¹⁰In particular, τ_0 usually does not correspond with the "uninformative" signal structure of Radner and Stiglitz (1984).

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.

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, δ is the uncertainty parameter that determines how cautious the regulator is in the description of the initial knowledge.¹¹ 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 τ .

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 τ , and the uncertainty parameter δ through the initial prior set $\mathcal{M}_0 = [1/2 - \delta, 1/2 + \delta]$ jointly determine

¹¹To keep the notation simple, we do suppress as often as possible the benefit’s dependency on the precision level τ .

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 δ . 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 $l(s, \delta) = (1/2 - \delta)l_+(s, \delta) + (1/2 + \delta)l_-(s, \delta)$. (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 θ_+ 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.

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

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 δ -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 δ (as long as δ 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 δ -derivative – and its sign in particular – does not depend on the uncertainty parameter δ . 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 π_Δ 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 write¹²

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

¹²To 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, $B(\delta) = B(\delta, \delta)$.

The parameter δ_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 δ_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 δ_1 and δ_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 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 δ_1 and δ_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 δ_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 δ_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 τ 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 δ_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 δ_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 τ is steeper the higher is δ_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 τ -derivative is positive by assumptions (10a) and (10b). The second factor, $V(s_+, \delta) - V(s_-, \delta)$, is positive by the assumption $\pi_+^{\dagger} \geq \pi_-^{\dagger}$. 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




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.

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 θ , 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 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 π_Δ 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 π_{Δ} because in both cases the two factors in the bracket and their τ -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 π_{Δ} -free positive coefficients α , β_0 and β_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 π_{Δ} . 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 π_{Δ} .

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 STEC example, the net effect is positive; at $\pi_{\Delta} = a_{\Delta}$, the class of the pesticide example¹³, 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

¹³The 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.

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 π_Δ . The consequence 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 cautious, *Journal of Environmental Economics and Management* **63**(3): 304–320.
- Barriau, P. and Sinclair-Desgagné, B. (2006). On precautionary policies, *Management science* **52**(8): 1145–1154.
- 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 (Lansdale, PA: IOS Press, 2009)* 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.
- 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. (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.
- Haimes, Y. Y. (2005). *Risk modeling, assessment, and management*, Vol. 40, John Wiley & Sons.
- Heyen, D. (2014). Learning under Ambiguity – A Note on the Belief Dynamics of Epstein and Schneider (2007), *AWI Discussion Paper 573*.
- Hirshleifer, J. and Riley, J. G. (1992). *The Analytics of Uncertainty and Information*, Cambridge University Press.
- 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, *Working Paper 18230*, National Bureau of Economic Research.
- Martin, I. W. R. and Pindyck, R. S. (2015). Averting Catastrophes: The Strange Economics of Scylla and Charybdis, *American Economic Review* **105**(10): 2947–2985.
- 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. (2002). *Risk and reason: Safety, law, and the environment*, Cambridge University Press.
- 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.

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 τ , $c(\tau) = 55\tau$. All effects in Figure 6 are evaluated at $\tau = 2 \ln(4/3)$.

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 $\underline{\rho}_1$.

C Full formula of δ_2 -derivatives

The general δ_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 δ 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 δ_2 -derivative at $\delta = 0$, as stated in (17), is positive. Due to continuity, this extends to a full neighborhood of $\delta = 0$. The δ_2 -derivative can become negative at some $\delta > 0$ due to the higher order terms $P(\delta)$ and $Q(\delta)$.

D The effects at $\delta > 0$

In section 4.3 we discussed all effects and their dependency on the payoff asymmetry π_Δ 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 π_Δ 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 $le\pi_\Delta \leq a_\Delta$; 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 δ . The net effect at $\pi_\Delta = a_\Delta$ is unambiguously negative and remains at this δ -independent level for all $\pi_\Delta > a_\Delta$, as was proven in section 4.3.

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 σ 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 δ_1 and δ_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 τ and δ_2 . For instance $\underline{s} = \bar{s} = 0$ when $\delta_2 = 0$.




Figure 8: The net effect as a function of the payoff asymmetry π_Δ 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 δ -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 π_Δ 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 δ seems to be more pronounced in the normal distribution structure. Interestingly, however, the net effect is still constant in π_Δ 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 δ 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.