Norwegian University of Life Sciences
School of Economics and Business

Averting catastrophes in a more complex world

Kine Josefine Aurland-Bredesen

Working Papers No. 6/ 2017

ISSN: 2464-1561

Averting catastrophes in a more complex world

Kine Josefine Aurland-Bredesen*

Abstract

It lies in the nature of a catastrophe that its cost and benefits are non-marginal. This makes policy decision making in the presence of catastrophes a complex task. Previous work has shown that when projects are non-marginal they should not be evaluated in isolation and that standard cost benefit analysis may provide biased results. This paper pursues this topic in three directions. First, I introduce the possibility that policy measures may be dependent. This allows the policy decision maker to exploit synergistic relationships and avoid unwanted and unintended consequences. Second, the existence of real world constraints may effect the optimal policy set. I evaluate the optimal policy under different constraints on model parameters such as cost, willingness to pay and risk. Third, I allow the likelihood of a catastrophe to increase with time. This provides a realistic framework for evaluating catastrophes characterized by accumulation or tipping points. I show how these extensions change the willingness to pay and the optimal policy set. This paper is largely theoretical, but provides policy decision makers with guidance on the existence and nature of possible biases through analytical results and examples.

Keywords: Multiple catastrophes, policy measure dependencies

JEL classification: D61, Q51, Q54

*School of Economics and Business, Norwegian University of Life Sciences. E-mail: kine.josefine.aurland-bredesen@nmbu.no

1 Introduction and background

The answer to the question “How should we evaluate projects?” depends on the nature of the projects we consider. The standard economic tool for evaluating projects is cost benefit analysis (CBA), but if the net benefits of a project are large compared to aggregated consumption, then CBA can cause biased results [Dasgupta et al., 1972]. This creates an interdependency among the projects. Thus, evaluating projects in isolation is not always a good approximation, especially when cost and benefits are non-marginal.

In this paper I argue that the current efforts to consider interdependence have not yet been taken far enough. Using policy measures to avert catastrophes, Martin and Pindyck [2015] are the first to address this interdependence when selecting among a set of large projects. But they fail to consider other dependencies such as direct dependencies among the policy measures or dependencies caused by economic boundaries. Thus, the argument against evaluating projects or policy measures in isolation should be even stronger. I argue that by not taking other dependencies into consideration we get sub-optimal policy recommendations.

Martin and Pindyck [2015] is not the only study that note the bias of CBA. Hoehn and Randall [1989]show how standard CBA is systematically biased when the number of projects is large. The capacity of the economy provides an upper bound on net benefits that is only evident when projects are evaluated together, whereas the standard measure is unbounded. Thus, when the number of projects is large we overestimate net benefits. Dietz and Hepburn [2013] examine the conditions of when CBA is biased when evaluating large projects. They show how using CBA to evaluate non-marginal climate and energy projects can result in sub-optimal solutions and find that the source of the error is the elasticity of marginal utility.

In their paper, Martin and Pindyck [2015][1] show that the standard rule of CBA, positive discounted net welfare, is a necessary, but not sufficient criteria for averting a catastrophe. They find that policies to avert catastrophes should not be evaluated in isolation, but rather in conjunction with each other. The benefits of averting one catastrophe depends positively on the background risk created by the existence of the other catastrophes.

Tsur and Zemel [2017] study intertemporal policies for managing multiple catastrophes where efforts to alleviate a catastrophe can be smoothed out over time and find that background risk can both increase and decrease the benefits of averting a catastrophe. Martin and Pindyck [2015] find that if the total benefits and individual costs are sufficiently small, then the problem can still be approximated by standard CBA. In contrast, Tsur and Zemel [2017] find that this approximation

[1] Other work that note the Martin and Pindyck [2015] study focus mostly on the implication of this on policies to deal with climate change.

may not hold even for marginal projects. This is discussed in more detail in section four.

I use the study by Martin and Pindyck [2015] to illustrate how policy measure dependencies, economic boundaries and time-dependent risk affects the optimal policy for managing catastrophes[2]. To do this I provide three extension that help move the framework closer to reality. All extensions have implications on policy recommendations

First, in the original paper policy measures are assumed to be independent, with the exception of the interdependence caused by background risk. The assumption of independence for both catastrophes and policy measures are likely to be violated in the real world. Policy measures aimed at averting one catastrophe may reduce the likelihood of other catastrophes and the damage distributions of catastrophes are often subjected to tail dependencies. My first extension opens up the framework to allow for dependencies, and I show how this affects both willingness to pay (WTP) for averting catastrophes and which catastrophes it is optimal to avert.

Second, the framework proposed by Martin and Pindyck [2015] does not impose any constraints on benefits, cost or risk. The policy decision maker is likely to face a budget constraint. Considerations such as the limitations of constant relative risk aversion (CRRA) utility [Geweke, 2001] or the overall correlation between a country's economy and WTP, are arguments for binding WTP from above. Introducing such overall constraints makes it more difficult for policy measures to pass the optimal policy criterion. The effect of a constraint on risk for one or more of the catastrophes depends on the nature of the dependencies between catastrophes and policy measures.

Third, in Martin and Pindyck [2015], each catastrophe is assumed to have a constant likelihood of occurring. For catastrophes that are characterized by accumulation or tipping points, the likelihood of the catastrophe occurring increases with time if there is no policy action. One such example is a climate catastrophe. Thus for my third extension I use an inhomogenous Poisson process to illustrate how to easily include one or more catastrophes where risk is time-dependent. If one or more catastrophes have a likelihood of occurring that is non-decreasing in time, it emphasizes the effect of background risk and the importance of discounting.

I show how the extensions affect the optimal policy and I use numerical examples to demonstrate the effect by comparing my results with the results in Martin and Pindyck [2015]. I end the paper by providing a summary of the results and concluding comments.

[2]Martin and Pindyck [2017] provide an interesting extension where they differentiate between catastrophes that cause destruction and death. They show that for catastrophes that kill the benefits of averting is larger. The main results in this paper is likely to hold also for deadly catastrophes because the underlying assumptions of their model does not change.

1.1 Martin and Pindyck's original model

Assume a society is facing N catastrophes. If catastrophe i occurs it causes a permanent drop in log consumption (c_t) equal to the random amount ϕ_i. Log consumption follows a Poisson process,

$$c_t = \log C_t = gt - \sum_{i=k}^{N} \sum_{t=1}^{Q_i(t)} \phi_{i,t}$$

where g is the growth rate and $Q_i(t)$ is the Poisson counting process for catastrophe i with known mean arrival rate λ_i. The cumulant generating function for time t is

$$\kappa(\theta)t = \left\{ g\theta + \sum_{i=1}^{N} \lambda_i (Ee^{-\theta\phi_i} - 1) \right\} t$$

If we do nothing, the discounted present welfare in $t = 0$ is[3]

$$W = \frac{1}{1-\eta} \frac{1}{\delta - \kappa(1-\eta)}$$

Averting catastrophe i is equal to setting $\lambda_i = 0$. The WTP to avert catastrophe i or a set S of catastrophes are given by the expression

$$(1 - w_p)^{1-\eta} = \frac{\delta - \kappa^p(1-\eta)}{\delta - \kappa(1-\eta)}$$

where p can represent both policy measure i or policy set S. The actual cost of averting catastrophe i is a permanent tax on consumption τ_i. Thus, if we choose to avert a set S of catastrophes the cost will be multiplying consumption by $\prod_{i \in S} (1 - \tau_i)$ forever. Martin and Pindyck [2015] defines two variables B_i and K_i ,

$$B_i = (1 - w_i)^{1-\eta} - 1$$

$$K_i = (1 - \tau_i)^{1-\eta} - 1$$

which corresponds to the percentage loss of utility when consumption is reduced by either τ_i% or w_i%. Therefore it is optimal to choose a subset S that solves the problem

$$\max_{S \subseteq \{1,..,N\}} V = \frac{1 + \sum_{i \in S} B_i}{\prod_{i \in S} (1 + K_i(p_i))} \tag{1}$$

[3]For an extensive review of the original model please see Martin and Pindyck [2015].

2 Averting catastrophes when policy measures are dependent

In Martin and Pindyck's [2015] original paper policy measures are, except for the interdependencies caused by non-marginality, assumed to be independent. Catastrophes are also assumed to occur independently of each other. The assumptions of independence for both catastrophes and policy measures are likely to be violated in the real world. Policy measures aimed at averting one catastrophe may reduce the likelihood of other catastrophes. For example, measures to reduce bioterrorism may also reduce other types of terrorism. The risk of a nuclear catastrophe might be reduced by reducing the number of nuclear power plants, but since nuclear power plants produce relative clean energy this may increase the likelihood of a climate catastrophe. Another example is the impact of climate change on the likelihood of natural disasters. If the likelihood of a climate catastrophe is reduced it may also reduce the likelihood of storms, hurricanes, and even catastrophic infectious diseases.

The damage distribution for the individual catastrophes may also be linked. Such dependencies often occur in the extreme values and are referred to as tail dependencies. For example if the damage from catastrophe i exceeds a certain value, then it is more likely for the damage of catastrophe j to also exceed a certain value. Although dependencies are often modeled using a single number, a correlation coefficient, this does not necessary reflect true dependencies.

2.1 Model reformulation

I reformulate the model to allow for policy measures dependencies. Let log consumption follow a spatial two-dimensional Poisson process that is homogeneous with respect to time. The second dimension is a policy vector $\mathbf{p}$ that contains the scaling of all $i = 1, .., N$ policy measures. For all i, the mean arrival rate of catastrophe i depends on $\mathbf{p}$, such that $\lambda_i(\mathbf{p})$. This opens up the model for dependencies between catastrophes and between policy measures. Let

$$\lambda_i(\mathbf{p}) = \lambda_i \{1 - p_i (\mathbf{p})\}$$

where

$$p_i(\mathbf{p}) = p_i + \rho_i(\mathbf{p})$$

and $p_i(\mathbf{p}) \in [0, 1]$ $\forall i$. p_i is the intended decrease in the likelihood of catastrophe i occurring as a result of introducing policy measure i. $\rho_i(\mathbf{p})$ is the unintended decrease or increase in the likelihood of catastrophe i occurring as a result of all other policy measures. This term captures dependencies. With this simple set up the cumulant generating function is

$$\kappa^{\mathbf{p}}(\theta) = g\theta - \sum_{i=1}^{n} \lambda_i(\mathbf{p})(\mathbb{E}e^{-\theta\phi_i} - 1) = g\theta - \sum_{i=1}^{n} \lambda_i \left(1 - p_i\left(\mathbf{p}\right)\right)\left(\mathbb{E}e^{-\theta\phi_i} - 1\right) \tag{2}$$

It does not matter if the origin of the dependencies is the mean arrival rate of catastrophes or the expected damages. Since $p_i\left(\mathbf{p}\right)$ is deterministic I can define

$$\mathbb{E}e^{-\theta\phi_i(\mathbf{p})} = p_i\left(\mathbf{p}\right)\left(1 - \mathbb{E}e^{-\theta\phi_i}\right) + \mathbb{E}e^{-\theta\phi_i}$$

where $p_i\left(\mathbf{p}\right)\left(1 - \mathbb{E}e^{-\theta\phi_i}\right)$ is the consumption we avoid losing when catastrophe i occurs as a consequence of the policy set $\mathbf{p}$. This yields the same cumulant generating function as (2). Thus, it only matters how we define $\rho_i(\mathbf{p})$ and $p_i(\mathbf{p})$ and this definition is determined by the dependencies true nature. Note that this framework does not put any restrictions on the formulation of dependencies as long as $\rho_i(\mathbf{p})$ is such that $p_i(\mathbf{p}) \in [0,1]\ \forall i$. If $\rho_i(\mathbf{p}) = 0$ for all i then all catastrophes and policy measures are assumed to be mutually independent.

I differentiate between evaluating policy measures individually, through a $1 \times N$ single impact policy vector,

$$\mathbf{p}_i = (0, ..., p_i, .., 0)$$

and as a set S, through a $1 \times N$ multiple impact policy vector,

$$\mathbf{p}_s = (p_1, ..., p_i, .., p_s, .., 0)$$

Note that evaluating in isolation does not imply that dependencies among the policy measures (or catastrophes) do not exist, but rather that we can only evaluate the effects of policy measure i on other catastrophes and not how the policy measures depend on each other. The relationship between the willingness to pay for the policy set S and policy measure i follows the same intuition as in Martin and Pindyck [2015], but does now also depend on the complexity of the dependencies. The relationship between the willingness to pay for the policy set S and policy measure i is described in Result 1.

Result 1: The relationship between willingness to pay for $\mathbf{p}_s$ and $\mathbf{p}_i$

The WTP for the total consequences of policy $\mathbf{p}_i$ in isolation is linked to the WTP for the total consequences of a policy set $\mathbf{p}_s = \sum_{i \in S} \mathbf{p}_s$ by the expression

$$(1 - w_s(\mathbf{p}_s))^{1-\eta} - 1 = \tag{3}$$

$$\sum_{i \in S} \left[(1 - w_i(\mathbf{p}_i))^{1-\eta} - 1 \right] + \sum_{i \in S} \left\{ \left[(1 - w_i^c(\mathbf{p}_s))^{1-\eta} - 1 \right] - \sum_{j \in S} \left[(1 - w_i^c(\mathbf{p}_j))^{1-\eta} - 1 \right] \right\}$$

where $w_i^c(\mathbf{p}_s)$ is the WTP for the unintended consequences of policy set S on catastrophe i and $w_i^c(\mathbf{p}_j)$ is the WTP for the unintended consequences of policy $\mathbf{p}_j$ on catastrophe i.

This shows that the original results from Martin and Pindyck [2015],

$$(1 - w_s(\mathbf{p}_s))^{1-\eta} - 1 = \sum_{i \in S} \left[(1 - w_i(\mathbf{p}_i))^{1-\eta} - 1 \right] \tag{4}$$

only holds in the face of dependencies if the last term is equal to zero. The last term is only zero if

$$\sum_{\mathbf{i=1}}^{\mathbf{s}} \rho_j(\mathbf{p_i}) = \rho_j(\sum_{\mathbf{i=1}}^{\mathbf{s}} \mathbf{p_i})$$

Remember that $\rho_j(\sum_{\mathbf{i=1}}^{\mathbf{s}} \mathbf{p_i})$ is equivalent to $\rho_j(\mathbf{p_s})$. The equality above only holds if the functional form of $\rho_j(\mathbf{p})$ is linearly separable in the elements of $\mathbf{p}$. In practice such a restriction means that $\rho_j(\mathbf{p})$ cannot have any interaction with the other elements in $\mathbf{p}$. This implies (4) holds for dependencies modeled using a linear relationship but not for more complex dependencies. From Jensen inequality we know that if the last term is equal to zero then $w_s(\mathbf{p_s}) \leq \sum_{i=1}^{S} w_i(\mathbf{p_i})$, but if the last term is not equal to zero this inequality may not hold.

The relationship between the WTP for a policy measure and a set of policy measures do not say anything about the optimal choice and scaling of $\mathbf{p}$. For easy comparison I follow Martin and Pindyck [2015] and define K_i as

$$K_i(p_i) = (1 - \tau_i(p_i))^{1-\eta} - 1$$

such that the costs of policy measure i depends on the scaling of policy measure, p_i. $B_i(\mathbf{p})$ is defined as

$$B_i(\mathbf{p}) = p_i B_i + \rho_i(\mathbf{p}) B_i$$

where $p_i B_i = \left[(1 - w_i(p_i))^{1-\eta} - 1 \right]$ captures the benefits from the intended consequences and is always non-negative. $\rho_i(\mathbf{p}) B_i = \left[(1 - w_i^c(\mathbf{p}))^{1-\eta} - 1 \right]$ represents the benefits of the unintended consequences which can be both positive and negative. Assuming that the unintended consequences are not both large and also negative, such that $1+\sum_{i=1}^{N} B_i(\mathbf{p}) > 0$, then for all i the policy decision maker should choose a policy vector$\mathbf{p}$ that solves the problem

$$\max_{\mathbf{p}} V = \left\{ \frac{1+\sum_{i=1}^{N} B_i(\mathbf{p})}{\prod_{i=1}^{n}(1+K_i(p_i))} \right\} \text{ s.t. } p_i \in [0,1] \tag{5}$$

Note that there is one important assumption behind this result. The unintended consequences cannot be large and negative. If they are, then $1+\sum_{i=1}^{N} B_i(\mathbf{p}) < 0$ and maximizing V is actually equivalent to minimizing welfare. Thus, the policy recommendations will reduce welfare as much as possible. In the rest of the paper I assume that the unintended consequences are such that they meet the requirement that $1+\sum_{i=1}^{N} B_i(\mathbf{p}) > 0$.

As in Martin and Pindyck [2015], the set of catastrophes can be divided into three groups. There are catastrophes i where discounted net present welfare is strictly increasing in p_i, such that the policy decision maker should do what they can to avert the catastrophe, i.e. setting $p_i = 1$. The cost-benefit trade off for these catastrophes are such that the relative rate of change in benefits is larger than the relative rate of change in costs

$$\frac{\frac{\partial}{\partial p_i}\sum_{i=1}^{n} B_i(\mathbf{p})}{1+\sum_{i=1}^{n} B_i(\mathbf{p})} > \frac{\frac{\partial K_i(p_i)}{\partial p_i}}{1+K_i(p_i)}$$

In addition there are catastrophes that should be partially alleviated. These catastrophes satisfy the criterion

$$\frac{\frac{\partial}{\partial p_i}\sum_{i=1}^{n} B_i(\mathbf{p})}{1+\sum_{i=1}^{n} B_i(\mathbf{p})} = \frac{\frac{\partial K_i(p_i)}{\partial p_i}}{1+K_i(p_i)}$$

This criterion also yields the optimal scaling of p_i when $p_i \in (0,1)$. Finally, for some catastrophes the discounted net present welfare is strictly decreasing in p_i, and the policy decision maker should do nothing. This implies that the relative rate of change in benefits is smaller than the relative rate of change in costs,

$$\frac{\frac{\partial}{\partial p_i}\sum_{i=1}^{n} B_i(\mathbf{p})}{1+\sum_{i=1}^{n} B_i(\mathbf{p})} < \frac{\frac{\partial K_i(p_i)}{\partial p_i}}{1+K_i(p_i)}$$

If there does not exist any dependencies then we will have no unintended consequences. Thus, to make it easier to compare the optimal solution with and without dependencies, I rewrite the benefit side in terms of the intended and unintended consequences

$$\frac{\frac{\partial}{\partial p_i}\sum_{i=1}^{n} B_i(\mathbf{p})}{1+\sum_{i=1}^{n} B_i(\mathbf{p})} = \frac{\frac{\partial p_i B_i}{\partial p_i} + \frac{\partial}{\partial p_i}\sum_{i=1}^{n} \rho_i(\mathbf{p}) B_i}{1+\sum_{i=1}^{n} p_i B_i + \sum_{i=1}^{n} \rho_i(\mathbf{p}) B_i} \tag{6}$$

Remember that when all policy measures are mutually independent, $\frac{\partial}{\partial p_i}\sum_{i=1}^{n}\rho_i(\mathbf{p}) = 0$ and $\sum_{i=1}^{n}\rho_i(\mathbf{p}) = 0$ for all i.

Note that, dependent or not, the optimal solution is characterized by the relationship between the relative rates of change in cost and benefits. The cost of introducing a policy i and scaling it p_i is the same no matter what the dependencies are, thus the cost side provides a fixed point for analyzing the effect of including dependencies in the framework. We can use this to investigate when ignoring dependencies causes us to recommend too much or too little policy action.

Result 2: Dependencies and the optimal set

If the relative rate of change in the benefits from the unintended consequences are larger than the relative rate of change in the benefit from the intended consequence,

$$\frac{\frac{\partial}{\partial p_i}\sum_{i=1}^{n}\rho_i(\mathbf{p})}{\sum_{i=1}^{n}\rho_i(\mathbf{p})} > \frac{\frac{\partial B_i(p_i)}{\partial p_i}}{1+\sum_{i=1}^{n}B_i(p_i)}$$

then it must be that

$$\frac{\frac{\partial}{\partial p_i}\sum_{i=1}^{n}B_i(\mathbf{p})}{1+\sum_{i=1}^{n}B_i(\mathbf{p})} > \frac{\frac{\partial B_i(p_i)}{\partial p_i}}{1+\sum_{i=1}^{n}B_i(p_i)}$$

and we pick $p_i < p_i^*$. The opposite is true when the relative rate of change in the benefits from the intended consequences are smaller than the relative rate of change in the benefits from the intended consequences.

If dependencies exist and we ignore them, then we either underestimate or overestimate the relative rate of change in benefits. Thus, we pick a p_i that is sub-optimal. Given the example above, if the policy decision maker ignores dependencies and find it is optimal to partially alleviate catastrophe i then we know that

$$\frac{\frac{\partial}{\partial p_i}\sum_{i=1}^{n}B_i(\mathbf{p})}{1+\sum_{i=1}^{n}B_i(\mathbf{p})} > \frac{\frac{\partial B_i(p_i)}{\partial p_i}}{1+\sum_{i=1}^{n}B_i(p_i)} = \frac{\frac{\partial K_i(p_i)}{\partial p_i}}{1+K_i(p_i)}$$

The policy decision maker should actually prevent catastrophe i and not partially alleviate it. The driving force behind this is that the unintended effects of p_i on in this case catastrophe j are such that the relative marginal benefit of the unintended effects from i are larger than the relative marginal cost of j. Martin and Pindyck [2015] argue that because projects to avert catastrophes are non-marginal it creates a interdependency among projects such that standard cost benefit analysis may caused biased results. I have shown that this is a simplification and that, unless the additional dependencies among the policy measures are insignificantly small, their framework causes sub-optimal policy recommendations. When including dependencies into the framework it also opens the possibility that the unintended consequences are so large and negative that the approach used in Martin and Pindyck [2015] will result in policy recommendations that minimize welfare, instead of maximizing welfare.

In the next section I illustrate how the optimal policy set changes when we allow for dependencies between a climate catastrophe and natural disasters

2.2 Numerical example

Assume that the likelihood of catastrophic climate change is correlated with the likelihood of storms and hurricanes, such that averting or partially alleviating a climate catastrophe also reduces the likelihood of these two natural disasters. Let's also assume that policies aimed at reducing the likelihood of bioterrorism also reduce the likelihood of nuclear terrorism and mega-virus. I assume a positive linear correlation between them, such that a 5% correlation implies that averting bioterrorism reduces the likelihood of nuclear terrorism and mega-virus by 5%.

In order to make the example similar to the examples in Martin and Pindyck [2015], I use the same parameter values and functional form. An overview of these can be found in the Appendix of this paper.

When the index of relative risk aversion is low ($\eta = 2$) Martin and Pindyck's [2015] numerical example show that it is not optimal to avert a climate catastrophe. In contrast, Figure 1 illustrates that if averting climate catastrophe reduces the likelihood of storms and hurricanes by at least 7,5%, it is optimal to avert a climate catastrophe. When relative risk aversion is low the original decision to not avert a climate catastrophe is not robust in the presence of even small correlations.

The numerical results provided in Martin and Pindyck [2015] show that when relative risk aversion is high ($\eta = 4$) it is not optimal to avert bioterrorism. But, if averting bioterrorism reduces the likelihood of nuclear terrorism and a mega-virus by at least 6%, then it is in fact optimal to avert bioterrorism.

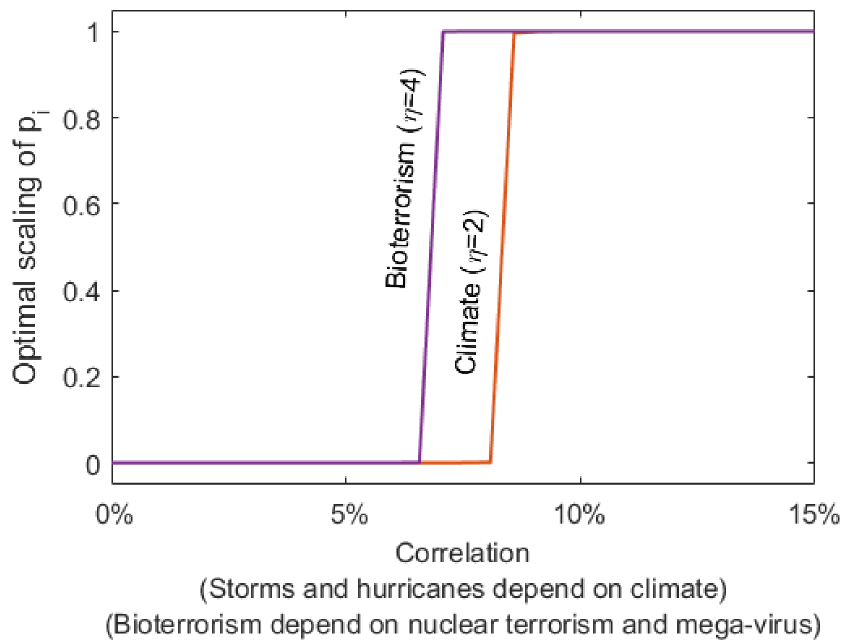

Figure 1: Example of how the optimal policy for bioterrorism and catastrophic climate change changes from not avert to avert if we consider possible correlations.
The likelihood of storms and hurricanes occurring is positively correlated with the likelihood of a climate catastrophe and the likelihood of nuclear terrorism and a mega-virus is positively correlated with the likelihood of bioterrorism. η is the relative rate of risk aversion.

3 Averting catastrophes under economic boundaries

3.1 Binding total costs and benefits from above

When deciding whether or not to avert or partially alleviate a range of catastrophes the decision maker is likely to be facing other non-catastrophic social planning problems that require resources. Thus, there exists a upper bound on the total fraction of consumption the policy decision maker can spend on averting catastrophes. I impose an arbitrary bound called T such that $T \geq \prod_{i=1}^{n} (1 + K_i(p_i))$ where T can be mapped to a upper constraint on the total consumption tax. I attach the multiplier ζ to the constraint and assume the constraint is binding. If the constraint is not binding we have an interior solution and the optimal policy set is the same as for the unconstrained problem. I solve (5) with the constraint imposed and rewrite the expression such that the cost side of the criterion remains the same as before. With a constraint on total cost, the benefit side of the criterion can be written as,

$$\frac{\frac{\partial}{\partial p_i} \sum_{i=1}^{N} B_i(\mathbf{p})}{1 + \sum_{i=1}^{n} B_i(\mathbf{p}) + \zeta T^2}$$

where $T = \prod_{i=1}^{N} (1 + K_i(p_i))$. Since the intuition is the same as for the unconstrained problem, we should partially alleviate catastrophe i if

$$\frac{\frac{\partial}{\partial p_i} \sum_{i=1}^{N} B_i(\mathbf{p})}{1 + \sum_{i=1}^{n} B_i(\mathbf{p}) + \zeta T^2} = \frac{\frac{\partial K_i(p_i)}{\partial p_i}}{1 + K_i(p_i)}$$

Since both $\zeta > 0$ and $T > 0$, we have that

$$\frac{\frac{\partial}{\partial p_i} \sum_{i=1}^{N} B_i(\mathbf{p})}{1 + \sum_{i=1}^{N} B_i(\mathbf{p})} > \frac{\frac{\partial}{\partial p_i} \sum_{i=1}^{N} B_i(\mathbf{p})}{1 + \sum_{i=1}^{N} B_i(\mathbf{p}) + \zeta T^2} = \frac{\frac{\partial K_i(p_i)}{\partial p_i}}{1 + K_i(p_i)} \tag{7}$$

The relative rate of change in costs is the same, but the benefit side is smaller. The constrained criterion provides less slack than the unconstrained case and the optimal policy set shrinks. This is also true if I apply a upper bound on benefits. In theory it is possible that consumers are willing to spend almost all of their consumption on averting catastrophes with small probabilities and large damages, see for example Geweke [2001]. In addition, WTP can be correlated with the overall economy of the country. This is especially true for this framework, where the WTP is given as a fraction of consumption. I impose an arbitrary upper bound on the willingness to pay. This implies the total benefits of any policy set is bounded from above by B. I attach the multiplier β to the constraint and assume the constraint is binding. The benefit side can be written as,

$$\frac{\frac{\partial}{\partial p_i} \sum_{i=1}^{N} B_i(\mathbf{p}) \left(1 - \beta \prod_{i=1}^{n} (1 + K_i(p_i))\right)}{1 + \sum_{i=1}^{N} B_i(\mathbf{p})}$$

Since $\prod_{i=1}^{N} (1 + K_i(p_i)) > 0$ and $0 < \beta < 1$,

$$\frac{\frac{\partial}{\partial p_i} \sum_{i=1}^{N} B_i(\mathbf{p})}{1 + \sum_{i=1}^{N} B_i(\mathbf{p})} > \frac{\frac{\partial}{\partial p_i} \sum_{i=1}^{N} B_i(\mathbf{p}) \left(1 - \beta \prod_{i=1}^{N} (1 + K_i(p_i))\right)}{1 + \sum_{i=1}^{N} B_i(\mathbf{p})} \tag{8}$$

Thus, binding benefits from above also shrinks the optimal policy set. With a more strict criterion it makes sense that we should focus on averting catastrophes with the highest ratio of relative rate of change in benefits to relative rate of change in costs. In other words, averting the catastrophes that give us the most bang for our buck.

Let $RM_i\Delta(\mathbf{p})$ be the ratio of relative rate of change in benefits to relative rate of change in cost. It is defined as

$$RM_i\Delta(\mathbf{p}) = \frac{\frac{\partial}{\partial p_i} \sum_{i=1}^{N} B_i(\mathbf{p})}{1 + \sum_{i=1}^{N} B_i(\mathbf{p})} \frac{1 + K_i(p_i)}{\frac{\partial K_i(p_i)}{\partial p_i}}$$

Let's use this to develop a rule for a sequential approach for choosing among projects. For example, if we avert or partially alleviate only one catastrophe, which one would we pick?

Result 3: Rule for sequential approach

The first catastrophe we should fully or partially alleviate, even when there exists a given a set of constraints on either WTP or tax, is the one with the largest ratio of relative rate of change in benefits relative to relative rate of change in costs above one. Pick catastrophe i such that

$$\{RM_i\Delta(\mathbf{p}) \geq RM_j\Delta(\mathbf{p}) | RM_i\Delta(\mathbf{p})) \geq 0\}\,, \forall j \neq i$$

As we will see, this is exactly what happens in the numeric examples in this paper. Using the parameter values in Martin and Pindyck [2015] I determine which catastrophes should be averted for a range of different constraints on both WTP and consumption tax. Catastrophes that are not optimal to avert or partially alleviated are excluded from Figure 2.

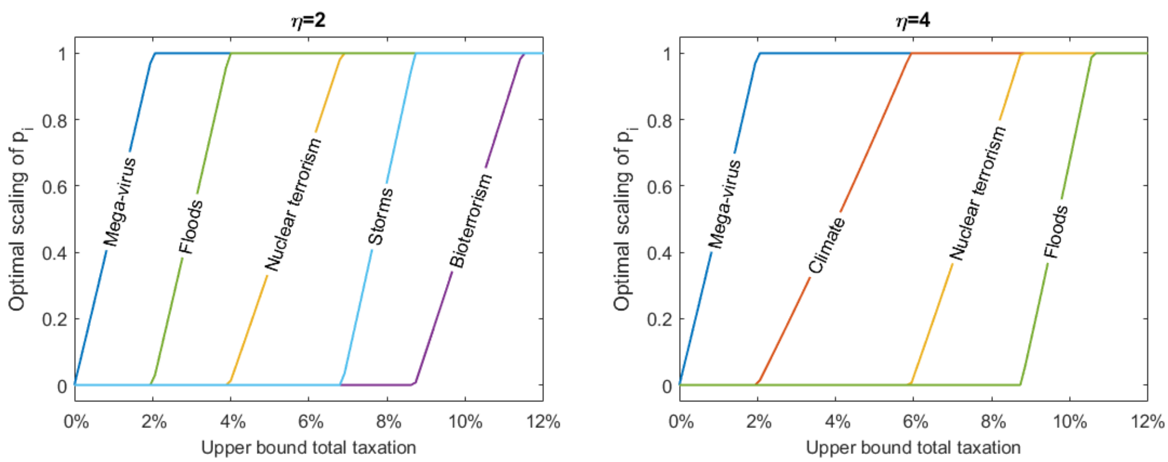

Figure 2: Constraint on taxation and the optimal scaling of policy vector $\mathbf{p}$. η is rate of relative risk aversion

If there is no constraint on taxation and the relative risk aversion is low ($\eta = 2$), it is optimal to avert five catastrophes. Such a policy action requires the equivalent of a 11,5% taxation on consumption. If there is no constraint on taxation and the relative risk aversion is high ($\eta = 4$) it is optimal to avert four catastrophes; this requires a 10,5% tax on consumption. Any boundary set above this has no effect on the optimal policy set.

With a very strict constraint on taxation, then, it is only optimal to avert a mega-virus catastrophe. This result holds for both levels of the relative risk aversion. If society can spend 5% of consumption on averting catastrophes we should avert avert a mega-virus catastrophe for both levels of relative risk aversion. If $\eta = 2$ the society should also avert floods and reduce the likelihood of nuclear terrorism by 1/3 or, if $\eta = 4$, reduce the likelihood of a climate catastrophe by 3/4.

The results of binding WTP from above is similar and can be seen in Figure 3.

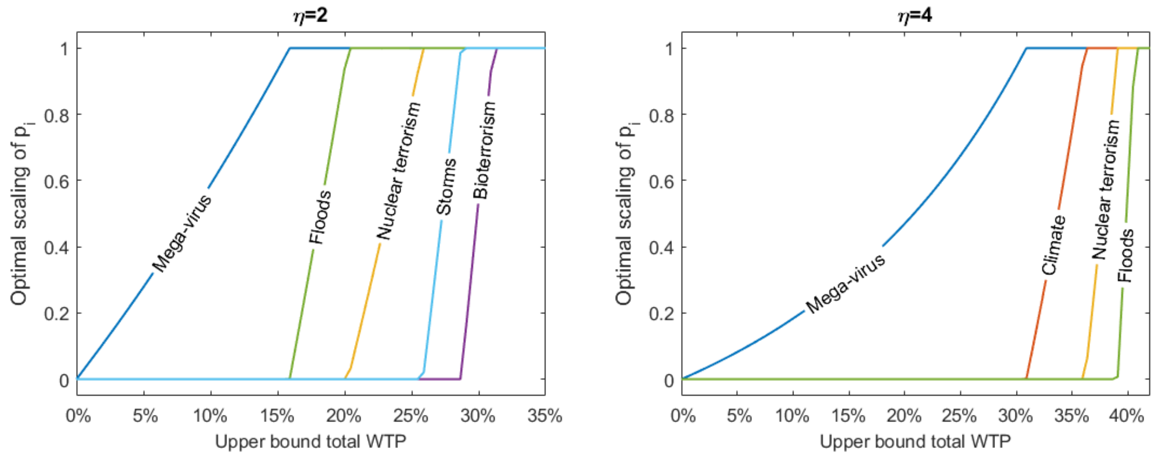

Figure 3: Constraint on WTP and the optimal scaling of policy vector $\mathbf{p}$. Relative risk aversion η

The willingness to pay to avert the optimal set when relative risk aversion is low is 31% of consumption and 41% of consumption when it is high. Any boundary set above this has no effect on the optimal policy set. If the willingness to pay is restricted such that a society is only willing to pay at most 20% of the consumption, the optimal policy bundle shrinks to avert a mega virus and floods when $\eta = 2$ and only a mega-virus if $\eta = 4$.

Introducing constraint on costs or benefits that are binding clearly changes the optimal policy set. The numerical analysis shows that, given these parameter values, the catastrophe the policy decision maker should focus on averting is a mega-virus catastrophe.

3.2 Constraints on risk

In the real world some catastrophes may dominate both media and policy agendas. Such catastrophes are likely to receive more attention from policy decision makers than others. For this reason, policy decision makers may be interested in reducing the likelihood of a given catastrophe so that it falls below some threshold.

Let λ_i^{max} be the maximum likelihood of catastrophe i that society (or the social planner) is willing to accept and assume $0 < \lambda_i^{max} < \lambda_i$. I attach the multiplier π_i and assume the constraint is binding. The cost side remains the same, but the benefit side of the criterion is

$$\frac{\frac{\partial}{\partial p_i}\sum_{i=1}^{N} B_i(\mathbf{p}) + \pi_i \lambda_i p_i \prod_{i=1}^{N}(1 + K_i(p_i))}{1 + \sum_{i=1}^{N} B_i(\mathbf{p})}$$

Since $\pi_i > 0$ and $\lambda_i p_i > 0$,

$$\frac{\frac{\partial}{\partial p_i}\sum_{i=1}^{N} B_i(\mathbf{p}) + \pi_i \lambda_i p_i \prod_{i=1}^{N}(1 + K_i(p_i))}{1 + \sum_{i=1}^{N} B_i(\mathbf{p})} > \frac{\frac{\partial}{\partial p_i}\sum_{i=1}^{N} B_i(\mathbf{p})}{1 + \sum_{i=1}^{N} B_i(\mathbf{p})} \tag{9}$$

This is not surprising, and it basically implies that if we impose an upper constraint on the likelihood of catastrophe i we should increase the scaling of the relevant measure. For catastrophes that are not optimal to avert or partially alleviate, this is a direct result of the constraint forcing us to scale up the policy. What is more interesting is what happens to the optimal scaling of policy measure i if there exist a constraint on the likelihood of another catastrophe, say k. The benefit side of the criterion is now

$$\frac{\frac{\partial}{\partial p_i}\sum_{i=1}^{N} B_i(\mathbf{p}) + \pi_k \lambda_k \left(\frac{\partial p_k(\mathbf{p})}{\partial p_i}\right) \prod_{i=1}^{N} \left(1 + K_i(p_i)\right)}{1 + \sum_{i=1}^{N} B_i(\mathbf{p})} \tag{10}$$

The last expression in (9) depends on the value of $\frac{\partial p_k(\mathbf{p})}{\partial p_i}$. If $\frac{\partial p_k(\mathbf{p})}{\partial p_i} > 0$ an increase in i decreases the likelihood of catastrophe k and vice versa. This implies that

- If $\frac{\partial \rho_k(\mathbf{p})}{\partial p_i} < 0$, the second term is negative, and the optimal scaling is characterized by $p_i^{no\ constraint} > p_i^{constraint}$
- If $\frac{\partial \rho_k(\mathbf{p})}{\partial p_i} > 0$, the second term is positive, and the optimal scaling is characterized by $p_i^{no\ constraint} < p_i^{constraint}$
- If $\frac{\partial \rho_k(\mathbf{p})}{\partial p_i} = 0$, the second term is zero, and the optimal scaling is characterized by $p_i^{no\ constraint} = p_i^{constraint}$

The next numerical example combines a binding upper constraint on the maximum likelihood of a climate catastrophe with different constraints on WTP and taxation. The example can be seen in Figure 4 and illustrates how introducing a binding upper constraint on risk can reduce welfare. If there is no constraint on WTP or taxation, the reduction in welfare compared to the optima is small. Introducing constraints on WTP and taxation yields a larger reduction in welfare. For example, if the policy decision maker can only spend 2.5% of consumption on averting catastrophes, reducing the likelihood of a climate catastrophe by half yields an 11% reduction in welfare compared to the unconstrained case. The point here is that, with these parameter values, there exists other catastrophes with a much larger ratio of relative rate of change between benefits and costs ($RM_i\Delta(\mathbf{p})$), than a climate catastrophe. Introducing an upper constraint on the likelihood of a climate catastrophe forces these catastrophes out of the optimal policy set, and this yields a drop in welfare.

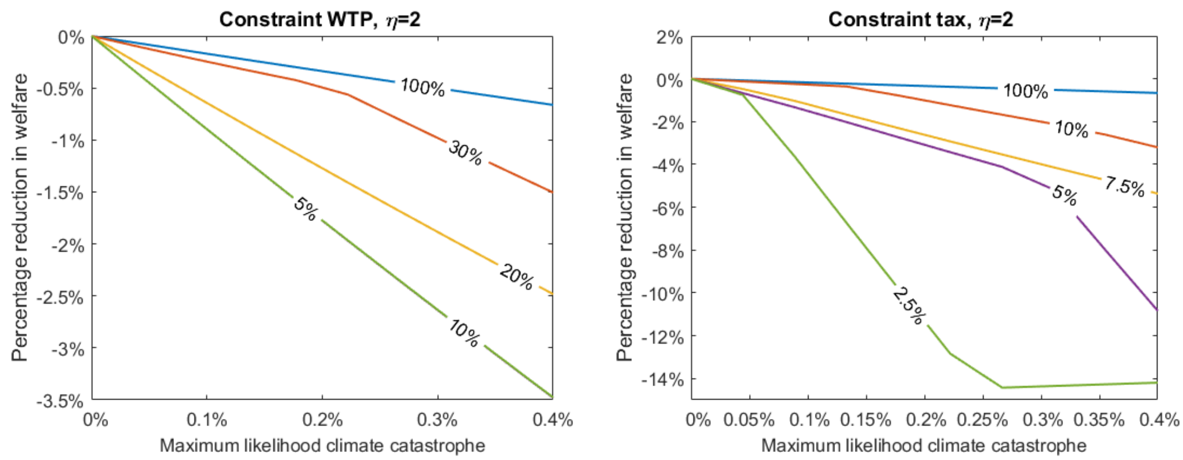

Figure 4: The welfare loss of imposing an upper constraint on the likelihood of a climate catastrophe combined with constraints on WTP and taxation. (10% constraint implies that taxation or WTP cannot exceed 10% of consumption)

4 Averting catastrophes when risk is time-dependent

In Martin and Pindyck [2015], log consumption follows a homogeneous Poisson process and the mean arrival rate for all catastrophes is constant over time. The assumption that the mean arrival rate is constant may not hold for all catastrophes. Take catastrophic climate change for example: If we choose to do nothing and continue the development path we have today is it likely that the probability of a climate catastrophe occurring is the same today and 50 years into the future? Probably not. For catastrophes where accumulation occurs or a tipping point exists, it is reasonable to assume that the mean arrival rate will increase with time.

The WTP to avert catastrophe i or a set S is derived by setting the discounted net present welfare of doing nothing equal to the discounted net present welfare of averting i or set S. A change in the likelihood of one catastrophe occurring, therefore, does not only affect the WTP to avert the catastrophe it self, it also affects the WTP to avert all the other catastrophes. This extension shows how to easily adapt the framework to one or more catastrophes with non-decreasing mean arrival rate and what the implications of this may be.

4.1 Model reformulation

Let log consumption follow an inhomogenous Poisson process. An inhomogeneous Poisson process is similar to an ordinary Poisson process except that the mean arrival rate is time dependent. This allows for seasonal fluctuations in the likelihood of a catastrophe occurring, or for the likelihood to be non-decreasing in time. The mathematical cost of this is that we lose the property of stationary increments, which is one of the properties of Levy processes.

Let $Q_i(t)$ be the counting process of an inhomogeneous Poisson process for catastrophe i with known intensity function $\lambda_i(t)$. Note that $\lambda_i(t)$ is a deterministic function of time t. The cumulative intensity function is $\Lambda_i(t) = \int_0^t \lambda_i(\tau)d\tau$. For the inhomogeneous Poisson process the probability of $Q_i(t)$ being equal to m_i is given by

$$P\left\{Q_i(t) = m_i\right\} = \frac{\Lambda_i(t)^{m_i}}{m_i!} e^{-\Lambda_i(t)}$$

I assume that $\lambda_i(t)$ is a non-decreasing function of time t and define it as

$$\lambda_i(t) = \left\{ \begin{array}{ccc} f_i(t) & if & \tau \in [0, t'_i] \\ \lambda_i^{max} & if & else \end{array} \right\}$$

This implies that the mean arrival rate is increasing in t until it reaches an upper level λ_i^{max} at time $t = t'_i$. The cumulant intensity function is then defined as

$$\Lambda_i(t) = \left\{ \begin{array}{ccc} \int_0^t f_i(t)dt = F_i(t) & if & t \in\in [0, t'_i] \\ \int_0^t \lambda_i^{max} dt = \lambda_i^{max} t & if & else \end{array} \right\}$$

Assume there are N catastrophes, and that catastrophes $i = 1, .., k$ have a mean arrival rate that is non-decreasing in time t. Let

$$r_i(t) = \left\{ \begin{array}{ccc} \Lambda_i(t) & if & i = 1, .., k \\ \lambda_i t & if & else \end{array} \right\}$$

such that the cumulant generating function in time t is

$$\kappa_t(\theta) = \theta g t + \sum_{i=1}^{n} r_i(t) \left(E e^{-\theta \phi_i} - 1\right)$$

Assume $i = 1, .., k$ is ordered such that $t'_{i-1} < t'_i$ for all i and that $t'_0 = 0$. The discounted net present welfare of doing nothing is

$$W_0 = \frac{1}{1-\eta} \int_0^\infty e^{-\delta t} e^{\kappa_t(1-\eta)} dt$$

which can be calculated by splitting the integral

$$W_0 = \frac{1}{1-\eta} \sum_{i=0}^{k-1} \left(\int_{t'_i}^{t'_{i+1}} e^{-\delta t} e^{\kappa_t(1-\eta)} dt \right)$$

Integration is complicated by the functional form of $f_i(t)$. The effect of introducing one or more catastrophes with a non-decreasing mean arrival rate are illustrated easiest through numerical examples. Still, there are some general conclusions that can be drawn. Assume $i = 1$ and let w denote WTP. One possible reasonable simplification is to choose the constant mean arrival rate to be a value in the range between the minimum and maximum, $\lambda \in [\lambda_0, \lambda^{max}]$.

- ***Constant mean arrival rate as a minimum = underestimates WTP and vice versa:*** *Setting $\lambda = \lambda_0$, implies that $w(\lambda) < w(\lambda(t))$. The opposite is true for $\lambda = \lambda^{max}$.*

Another reasonable simplification is to set the constant mean arrival rate to be the average of the non-decreasing mean arrival rate, but this overestimates WTP.

- ***Constant mean arrival rate as an average of the intensity function is overestimating WTP:*** *If we simplify and use the original framework with $\lambda = \overline{\lambda(t)}$, then because $t \to \infty$ we pick $\lambda = \lambda^{max}$. This is equivalent to setting $t' = 0$. Thus $w(\lambda) > w(\lambda(t))$ as long as the discount factor is non-zero.*

There is a clear relationship between the WTP and when the mean arrival rate reaches the maximum value. The earlier in time the maximum value is reached, the higher is the WTP to avert the catastrophe. Because this provides an increase in background risk, it also increases the WTP to avert other catastrophes.

- ***Relationship between WTP and the stop time t':*** *For $t' < t''$, t' reaches λ^{max} faster than t'', then $w(t') > w(t'')$.*

When the mean arrival rate is non-decreasing time plays a role in two different dimensions. It increases the likelihood of the catastrophe occurring, thus also increasing the benefits of averting the catastrophe. In contrast to this, we value the benefits from averting the catastrophe less because the largest benefits of doing so now occur further into the future.

- ***The importance of discounting:*** *Let δ be the discount rate. Because of discounting there exist t' and t'' such that for $t' < t''$ we have $w(t'') = w(t')$. Then $\delta_m > \delta_n$ implies that if $w(t', \delta_n) = w(t'', \delta_n)$ for t' and t'' then $w(t', \delta_m) > w(t'', \delta_m)$. Discounting plays a even more important role when the mean arrival rate is non-decreasing.*

This extension is related to earlier mentioned work by Tsur and Zemel [2017] on intertemporal policies to manage multiple catastrophes. Whereas I simply assume that the likelihood of a catastrophe occurring depends on time, Tsur and Zemel [2017] specify the underlying relationship. They assume the likelihood is increasing in some state variable. In the context of a climate catastrophe this state variable can be the atmospheric GHG concentration, and it evolves over time depending on abatement efforts and emissions output. Tsur and Zemel [2017] allow for partial alleviation (here:

abatement efforts) to be smoothed over time. I assume that the decision is made at time $t = 0$. The aim of my extension is to, in an intuitive way, show how including an intertemporal dimension can change policy recommendations. For example how the different approaches to choosing the constant mean arrival rate affects the WTP, and how ignoring the intertemporal dimension may underestimate the WTP for efforts to avert a climate catastrophe. To illustrate the latter I use a numerical example. This example looks at both the effect on WTP and the optimal policy set when the likelihood of a climate catastrophe is non-decreasing in time.

4.2 Numerical example

Assume that the likelihood of a climate catastrophe ($i = 2$) is non-decreasing in time. The intensity function is defined as

$$f(t) = \lambda_2^0 + \alpha t^\beta$$

such that

$$F(t) = t\left(\lambda_2^0 + \frac{\alpha t^\beta}{\beta + 1}\right)$$

I introduce three different values for λ_2^0 and λ_2^{max} and let $\beta = 0.25$. The rest of the parameter values remain the same as in the original article by Martin and Pindyck [2015]. The minimum and maximum values of the mean arrival rate are symmetrically picked around the original constant mean arrival rate $\lambda_2 = .004$. For the pair $\left\{\lambda_2^0, \lambda_2^{max}\right\} = \{.002, .006\}$ we are underestimating the true WTP for averting a climate catastrophes as long as the mean arrival rate reaches its maximum value less than 175 years into the future.

In the original article it is not optimal to avert a climate catastrophe when relative risk aversion is low ($\eta = 2$). In Figure 5, as long as the stop time t is to the left of the yellow dotted line, it is optimal to avert the catastrophe. This implies that, if the mean arrival rate reaches it maximum value less of .006 less than 130 years into the future, we should avert a climate catastrophe.

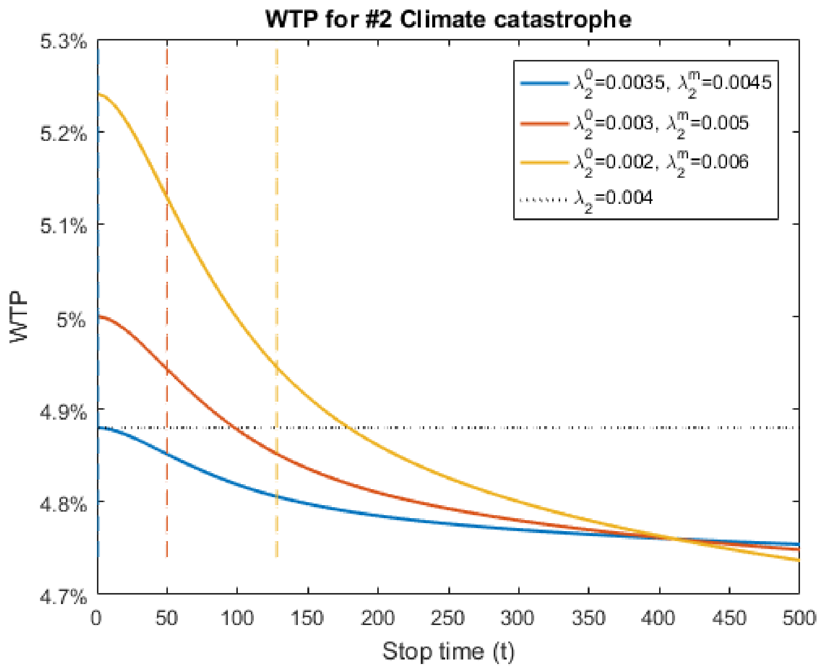

Figure 5: The WTP for averting a climate catastrophe when the mean arrival rate is non-decreasing in time.
$\beta = 0.25$. $\lambda_2(t)$ reaches λ_2^{max} at stop time t'. For each pair of λ-values if the stop time t' is to the left of the colored dotted line it is optimal to avert the catastrophe. Relative risk aversion $\eta = 2$.

For the pair $\{\lambda_2^0, \lambda_2^{max}\} = \{.003, .005\}$ we are underestimating the true WTP for averting a climate catastrophe as long as the mean arrival rate reaches it maximum value less than 100 years into the future. If we reach the maximum mean arrival rate of .005 before 50 years into the future we should, given these parameter values, avert the climate catastrophe. The further into the future stop time t' is, the lower the WTP to avert the catastrophe. The reason for this is that the largest benefits of averting a catastrophe occur when the likelihood is the highest. If this happens too far into the future, these benefits will be "discounted away". I have argued earlier that the assumption that the mean arrival rate of the climate catastrophe is constant is not realistic. Figure 5 shows how we risk underestimating the WTP to avert such a catastrophe and how it causes us to make sub-optimal policy decisions. For low relative risk aversion ($\eta = 2$) we should not avert a climate catastrophe when the mean arrival rate is constant, but Figure 5 shows how this may be subject to change. For example, if we believe that instead of the likelihood being 0.004 it is instead 0.002 today, and that if we do nothing(!) this will increase to 0.006 in the next 100 years, we should take measures to avert a climate catastrophe.

Increasing the likelihood of one catastrophe occurring over time also results in an increase in background risk. This increase causes a increase in the willingness to pay to avert other catastrophes. The further into the future a climate catastrophe reaches it maximum likelihood of occurring, the smaller the effect of the background risk from this catastrophe will be. This is because we discount future net benefits. This is clearly illustrated in Figure 6.

In contrast to Martin and Pindyck [2015], Tsur and Zemel [2017] find ambiguous effects when it comes to the effect of background risk. My work show similar results. We can see in Figure 6 that the WTP to avert other catastrophes is increasing in the beginning, even though this is where the effect of background risk should be the strongest.

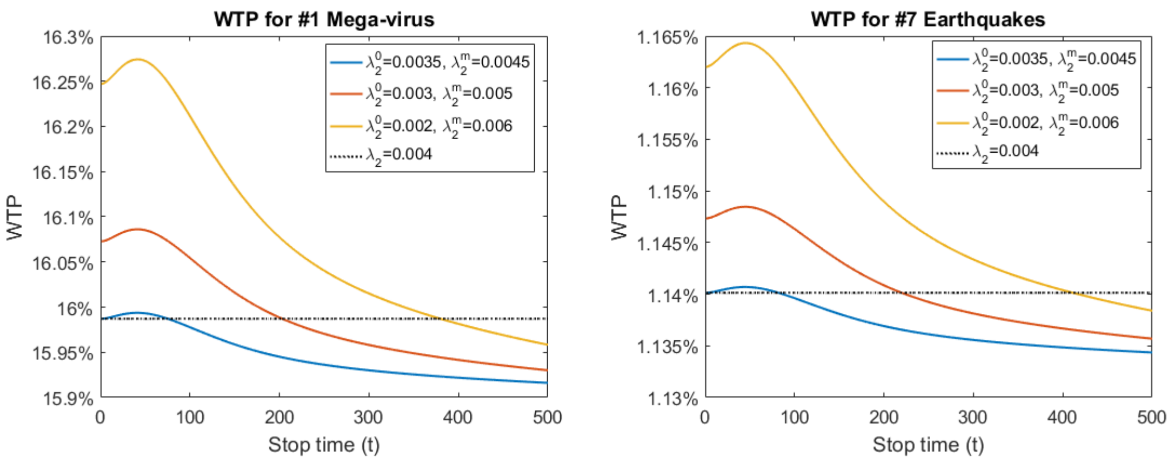

Figure 6: The WTP to avert a mega-virus and earthquakes when the mean arrival rate for a climate catastrophes ($i = 2$) is non-decreasing in time.
$\beta = 0.25.\lambda_2(t)$ reaches λ_2^{max} at stop time t'. For each pair of λ-values if the stop time t' is to the left of the colored dotted line it is optimal to avert the catastrophe. Relative risk aversion $\eta = 2$.

For the pairs $\{.003, .005\}$ and $\{.002, .006\}$ we underestimate the WTP for averting a mega-virus catastrophe and earthquakes if the mean arrival rate of a climate catastrophe reaches it's maximum value before 200 years into the future.

5 Conclusion and summary

The main motivation of this paper is to analyze how different dependencies affect policy recommendations for managing multiple catastrophes. I have demonstrated how the model in Martin and Pindyck [2015] can be extended to include policy measure dependencies, economic boundaries and time-dependent risk. The effect of the extensions are illustrated both using analytical results and numerical examples.

In the original paper Martin and Pindyck [2015] assume there are no dependencies between policy measures or between catastrophes. When this assumption is removed we see how the relationship between the WTP to avert a set of catastrophes and the sum of the individual WTP to avert the catastrophes in the set depends on the nature of dependencies. The relationship introduced in the original paper only holds for simple dependencies such as linear relationships where there is no interaction between terms. Dependencies also have an effect on the optimal choice and scaling of policy. If the relative rate of change in the benefits from the unintended consequences of the policy measure are larger than the relative rate of change in the benefit from the intended consequences of the policy measure, we should scale the policy measure higher than we would if there were no dependencies. The reason for this is that we ignore the possibility that the unintended benefits have a lower marginal cost than the intended benefits. In general if dependencies are overlooked will lead to sub-optimal policy recommendations. In the numerical examples it is obvious that the policy recommendations change for even small correlations between the catastrophes. For example, when relative risk aversion is low we should not avert a climate catastrophe unless averting a climate catastrophe decreases the likelihood of storms and floods by 7,5%.

Binding either WTP or consumption taxation provides a stricter criterion for which catastrophes should be averted or partially alleviated, keeping only the catastrophes with the largest ratio of relative rate of change in benefits to relative rate of change in costs in the optimal policy set. For example, with strict constraints on either costs or benefits, the only catastrophe we should avert is a mega-virus catastrophe. The effect of a constraint on risk depends on the nature of dependencies. If the constraint is on catastrophe k and there are no dependencies between catastrophe i and k, the constraint has no effect on the optimal scaling of the policy action to avert catastrophe i. If there are synergistic relationships, we should scale up the policy measure to avert catastrophe i and the opposite for mitigating relationships. The numerical example also illustrate that such constraints reduce overall welfare.

In the Martin and Pindyck's [2015] original model, all catastrophes have a constant likelihood of occurring. For catastrophes where accumulation occurs or a tipping point exist, the likelihood is likely to increase with time. One such example is a climate catastrophe. If the mean arrival rate of one catastrophe is non-decreasing in time it implies the largest benefits from averting the catastrophe occurs in the future. This emphasizes the importance of the choice of discount rate. The point is illustrated by looking at the example of a climate catastrophe. The minimum and maximum values of the mean arrival rate were symmetrically chosen around the constant mean arrival rate. When relative risk aversion is low and the mean arrival rate is constant it is not optimal to avert a climate catastrophe. In contrast, when the mean arrival rate is non-decreasing, this result is not robust. The numerical examples also show ambiguous effects when it comes to the effect of background risk.

The extensions presented in this paper make the framework more complex and maybe less intuitive, but they also reflect real world issues. The results underline the importance of evaluating policy measures to avert catastrophes as a set. One lesson that we can take from Martin and

Pindyck [2015] is that because of the non-marginal nature of catastrophe we should not evaluate policy measures to avert a climate catastrophe in isolation My argument goes further than Martin and Pindyck's [2015]. In addition to interdependencies caused by non-marginal projects and background risk, we should not evaluate in isolation because of possible policy dependencies, economic boundaries and time-dependent risk. This paper illustrates how policy dependencies cause bias in both WTP and the optimal policy set. The magnitude and size of the bias depends on the nature of the policy dependencies. Thus, more more empirical research on the existence and nature of policy dependencies is needed.

References

Partha Dasgupta, Amartya Sen, Stephen Marglin, et al. Guidelines for project evaluation. *Guidelines for Project Evaluation.*, 1972.

Simon Dietz and Cameron Hepburn. Benefit–cost analysis of non-marginal climate and energy projects. *Energy Economics*, 40:61–71, 2013.

John Geweke. A note on some limitations of crra utility. *Economics letters*, 71(3):341–345, 2001.

John P Hoehn and Alan Randall. Too many proposals pass the benefit cost test. *The American Economic Review*, 79(3):544–551, 1989.

Ian Martin and Robert S Pindyck. Averting catastrophes that kill. Technical report, National Bureau of Economic Research, 2017.

Ian WR Martin and Robert S Pindyck. Averting catastrophes: The strange economics of scylla and charybdis. *The American Economic Review*, 105(10):2947–2985, 2015.

Yacov Tsur and Amos Zemel. Coping with multiple catastrophic threats. *Environmental and Resource Economics*, pages 1–22, 2017.

Appendix

Parameter values

$z_i = e^{-\phi_i}$ is distributed according to the power distribution with parameter $\alpha_i > 0$, such that $b(z_i) = \alpha_i z_i^{\alpha_i - 1}$ with $0 \leq z_i \leq 1$. Growth rate $g = 0.02$ and discount rate $\delta = 0.02$

Potential catastrophe	λ_i	β_i	τ_i
Mega-virus	0.02	5	0.02
Climate	0.004	4	0.04
Nuclear terrorism	0.04	17	0.03
Bioterrorism	0.04	32	0.03
Floods	0.17	100	0.02
Storms	0.14	100	0.02
Earthquakes	0.03	100	0.01

Table 1: Parameter values Martin and Pindyck [2015]

Results from original article

Potential catastrophe	$\eta = 2$		$\eta = 4$	
	w_i	Avert=1	w_i	Avert=1
Mega-virus	0.159	1	0.309	1
Climate	0.048	0	0.180	1
Nuclear terrorism	0.086	1	0.141	1
Bioterrorism	0.047	1	0.079	0
Floods	0.061	1	0.096	1
Storms	0.051	1	0.082	0
Earthquake	0.011	0	0.020	0

Table 2: Results Table 1 Martin and Pindyck [2015]

Proofs

Proof: Result 1

Let$\kappa^S(\theta)$ be the CGF for the intended consequences and $\kappa_c^S(\theta)$ the CGF for the unintended consequences. Because $p_i(\mathbf{p}) = p_i + \rho_i(\mathbf{p})$ the cumulant generating function can be split into two parts such that,

$$\kappa^{\mathbf{S}}(\theta) = \kappa^S(\theta) + \kappa_c^S(\theta) - \kappa(\theta) *$$

and similar for **i**. Then

$$\sum_{\mathbf{i}\in\mathbf{S}} \frac{\delta-\kappa^{\mathbf{i}}(\theta)}{\delta-\kappa(\theta)} = \sum_{\mathbf{i}\in\mathbf{S}} \frac{\delta-\kappa^{i}(\theta)-\kappa_c^i(\theta)+\kappa(\theta)}{\delta-\kappa(\theta)}$$

Remember that $\sum_{\mathbf{i}\in\mathbf{S}} \kappa^i(\theta) = (S-1)\kappa(\theta)+\kappa^S(\theta)$, such that

$$\sum_{\mathbf{i}\in\mathbf{S}} \frac{\delta-\kappa^{\mathbf{i}}(\theta)}{\delta-\kappa(\theta)} = \frac{S\delta+\kappa(\theta)-\kappa^S(\theta)-\sum_{\mathbf{i}\in\mathbf{S}}\kappa_c^i(\theta)}{\delta-\kappa(\theta)}$$

Insert for $\kappa^S(\theta)$ using * $(\kappa^{\mathbf{S}}(\theta)-\kappa_c^S(\theta)+\kappa(\theta)=\kappa^S(\theta))$, this yields

$$\sum_{\mathbf{i}\in\mathbf{S}} \frac{\delta-\kappa^{\mathbf{i}}(\theta)}{\delta-\kappa(\theta)} = \frac{\left(\delta-\kappa^{\mathbf{S}}(\theta)\right)+\left\{\sum_{\mathbf{i}\in\mathbf{S}}\left(\delta-\kappa_c^i(\theta)\right)-\left(\delta-\kappa_c^S(\theta)\right)\right\}}{\delta-\kappa(\theta)} \quad **$$

We have that

$$\frac{\delta-\kappa^{\mathbf{i}}(\theta)}{\delta-\kappa(\theta)}-1=(1-w(\mathbf{p_i}))^{1-\eta}-1$$

and similar for all other. Notation: $w(\mathbf{p_i})$ is the total WTP for policy measure i,$w^c(\mathbf{p_i})$ is the WTP for the unintended consequences caused by policy measure i and$w^c(\mathbf{p_s})$ is the WTP for the unintended consequences caused by policy set S. Then $**$ can be rewritten to

$$(1-w(\mathbf{p_s}))^{1-\eta}-1=\sum_{\mathbf{i}\in\mathbf{S}}\left((1-w(\mathbf{p_i}))^{1-\eta}-1\right)+\left\{\left((1-w^c(\mathbf{p_s}))^{1-\eta}-1\right)-\left(\sum_{\mathbf{i}\in\mathbf{S}}(1-w^c(\mathbf{p_i}))^{1-\eta}-1\right)\right\}***$$

Since (from Martin and Pindyck's [2015] original result)

$$(1-w^c(\mathbf{p_s}))^{1-\eta}-1=\sum_{\mathbf{i}\in\mathbf{S}}\left\{\left((1-w_i^c(\mathbf{p_s}))^{1-\eta}-1\right)\right\}$$

I can rewrite $***$ such that

$$(1-w(\mathbf{p_s}))^{1-\eta}-1=\sum_{\mathbf{i}\in\mathbf{S}}\left((1-w(\mathbf{p_i}))^{1-\eta}-1\right)+\sum_{\mathbf{i}\in\mathbf{S}}\left\{\left((1-w_i^c(\mathbf{p_s}))^{1-\eta}-1\right)-\sum_{\mathbf{j}\in\mathbf{S}}\left((1-w_i^c(\mathbf{p_j}))^{1-\eta}-1\right)\right\}$$

■

Proof: Unconstrained maximization

$$\max_{\mathbf{p}} V=\left\{\frac{1+\sum_{i=1}^n B_i(\mathbf{p})}{\prod_{i=1}^n(1+K_i(p_i))}\right\} \text{ s.t. } p_i\in[0,1]$$

Attach multipliers γ_i to constraint $0 \geq p_i - 1$ and μ_i to $0 \geq -p_i$, the Lagrangian is $\mathcal{L} = V(\mathbf{p}) + \sum_{i=1}^{N} \gamma_i (1 - p_i) + \sum_{i=1}^{N} \mu_i (p_i)$ and the solution is given by $\frac{\partial V}{\partial p_i} = \mu_i - \gamma_i$. For the corner solution $p_i = 0$ I have that $\frac{\partial V}{\partial p_i} = -\mu_i$, which implies

$$\frac{\frac{\partial}{\partial p_i} \sum_{i=1}^{N} B_i(\mathbf{p})}{1 + \sum_{i=1}^{N} B_i(\mathbf{p})} < \frac{\frac{\partial K_i(p_i)}{\partial p_i}}{\prod_{i=1}^{N} (1 + K_i(p_i))}$$

For the corner solution $p_i = 1$ I have that $\frac{\partial V}{\partial p_i} = \gamma_i$, which implies

$$\frac{\frac{\partial}{\partial p_i} \sum_{i=1}^{N} B_i(\mathbf{p})}{1 + \sum_{i=1}^{N} B_i(\mathbf{p})} > \frac{\frac{\partial K_i(p_i)}{\partial p_i}}{\prod_{i=1}^{N} (1 + K_i(p_i))}$$

and finally for the interior solution I have that $\frac{\partial V}{\partial p_i} = 0$, which implies

$$\frac{\frac{\partial}{\partial p_i} \sum_{i=1}^{N} B_i(\mathbf{p})}{1 + \sum_{i=1}^{N} B_i(\mathbf{p})} = \frac{\frac{\partial K_i(p_i)}{\partial p_i}}{\prod_{i=1}^{N} (1 + K_i(p_i))}$$

Proof: Constrained maximization - taxation

$$\max_{\mathbf{p}} V = \left\{ \frac{1 + \sum_{i=1}^{n} B_i(\mathbf{p})}{\prod_{i=1}^{n} (1 + K_i(p_i))} \right\} \text{ s.t. } p_i \in [0, 1]$$

and $T \geq \prod_{i=1}^{N} [1 + K_i(p_i)]$. The Lagrangian is $\mathcal{L} = V(\mathbf{p}) + \sum_{i=1}^{N} \gamma_i (1 - p_i) + \sum_{i=1}^{N} \mu_i (p_i) + \zeta \left(T - \prod_{i=1}^{N} [1 + K_i(p_i)] \right)$. The solution is given by $\frac{\partial V}{\partial p_i} - \zeta \frac{\partial K_i(p_i)}{\partial p_i} = \mu_i - \gamma_i$. I assume the constraint $T \geq \prod_{i=1}^{N} [1 + K_i(p_i)]$is binding. For the corner solution $p_i = 0$ I have that $\frac{\partial V}{\partial p_i} - \zeta \frac{\partial K_i(p_i)}{\partial p_i} = -\mu_i$, which implies

$$\frac{\frac{\partial}{\partial p_i} \sum_{i=1}^{N} B_i(\mathbf{p})}{\left(1 + \sum_{i=1}^{N} B_i(\mathbf{p}) + \zeta T^2\right)} < \frac{\frac{\partial K_i(p_i)}{\partial p_i}}{\prod_{i=1}^{N} (1 + K_i(p_i))}$$

For the corner solution $p_i = 1$ I have that $\frac{\partial V}{\partial p_i} - \zeta \frac{\partial K_i(p_i)}{\partial p_i} = \gamma_i$, which implies

$$\frac{\frac{\partial}{\partial p_i} \sum_{i=1}^{N} B_i(\mathbf{p})}{\left(1 + \sum_{i=1}^{N} B_i(\mathbf{p}) + \zeta T^2\right)} > \frac{\frac{\partial K_i(p_i)}{\partial p_i}}{\prod_{i=1}^{N} (1 + K_i(p_i))}$$

and finally for the interior solution I have that $\frac{\partial V}{\partial p_i} - \zeta \frac{\partial K_i(p_i)}{\partial p_i} = 0$, which implies

$$\frac{\frac{\partial}{\partial p_i} \sum_{i=1}^{N} B_i(\mathbf{p})}{\left(1 + \sum_{i=1}^{N} B_i(\mathbf{p}) + \zeta T^2\right)} = \frac{\frac{\partial K_i(p_i)}{\partial p_i}}{\prod_{i=1}^{N} (1 + K_i(p_i))}$$

For all $p_i \in [0, 1]$ the benefit side of the criterion is given by

$$\frac{\frac{\partial}{\partial p_i} \sum_{i=1}^{N} B_i(\mathbf{p})}{\left(1 + \sum_{i=1}^{N} B_i(\mathbf{p}) + \zeta T^2\right)}$$

while the cost side remains the same.

Proof: Constrained maximization - WTP

$$\max_{\mathbf{p}} V = \left\{ \frac{1 + \sum_{i=1}^{n} B_i(\mathbf{p})}{\prod_{i=1}^{n} (1 + K_i(p_i))} \right\} \text{ s.t. } p_i \in [0,1]$$

and $B \geq \sum_{i=1}^{N} B_i(\mathbf{p})$. The Lagrangian is $\mathcal{L} = V(\mathbf{p}) + \sum_{i=1}^{N} \gamma_i (1 - p_i) + \sum_{i=1}^{N} \mu_i (p_i) + \beta \left(B - \sum_{i=1}^{N} B_i(\mathbf{p} \right)$. The solution is given by $\frac{\partial V}{\partial p_i} - \beta \left(\frac{\partial}{\partial p_i} \sum_{i=1}^{N} B_i(\mathbf{p}) \right) = \mu_i - \gamma_i$. I assume the constraint $B \geq \sum_{i=1}^{N} B_i(\mathbf{p})$ is binding. Following the same procedure as above, this yields the benefit side

$$\frac{\frac{\partial}{\partial p_i} \sum_{i=1}^{N} B_i(\mathbf{p})) \left(1 - \beta \prod_{i=1}^{N} (1 + K_i(p_i)) \right)}{\left(1 + \sum_{i=1}^{N} B_i(\mathbf{p}) \right)}$$

leaving the cost side the same.

Proof: Constrained maximization - Risk

$$\max_{\mathbf{p}} V = \left\{ \frac{1 + \sum_{i=1}^{n} B_i(\mathbf{p})}{\prod_{i=1}^{n} (1 + K_i(p_i))} \right\} \text{ s.t. } p_i \in [0,1]$$

and $\lambda_k^{max} \geq \lambda_k(1 - p_k(\mathbf{p}))$. The Lagrangian is $\mathcal{L} = V(\mathbf{p}) + \sum_{i=1}^{N} \gamma_i (1 - p_i) + \sum_{i=1}^{N} \mu_i (p_i) + \pi_k (\lambda_k^{max} - \lambda_k(1 - k(\mathbf{p})))$. The solution is given by $\frac{\partial V}{\partial p_i} + \pi_k \lambda_k \frac{\partial p_k(\mathbf{p})}{p_i} = \mu_i - \gamma_i$. This yields the benefit side

$$\frac{\frac{\partial}{\partial p_i} \sum_{i=1}^{N} B_i(\mathbf{p}) + \pi_k \lambda_k \frac{\partial p_k(\mathbf{p})}{\partial p_i} \prod_{i=1}^{N} (1 + K_i(p_i))}{(1 + \sum_{i=1}^{n} B_i(\mathbf{p}))}$$

and if $k = i$

$$\frac{\frac{\partial}{\partial p_i} \sum_{i=1}^{N} B_i(\mathbf{p}) + \pi_k \lambda_k p_i \prod_{i=1}^{N} (1 + K_i(p_i))}{(1 + \sum_{i=1}^{n} B_i(\mathbf{p}))}$$

leaving in both cases the cost side the same.