Journal of Environmental Economics and Policy

ISSN: 2160-6544 (Print) 2160-6552 (Online) Journal homepage: http://www.tandfonline.com/loi/teep20

When can the environmental profile and emissions reduction be optimised independently of the pollutant level?

Nils Chr. Framstad

Journal of Environmental Economics and Policy, 2014
Vol. 3, No. 1, 25–45, http://dx.doi.org/10.1080/21606544.2013.856353

When can the environmental profile and emissions reduction be optimised independently of the pollutant level?

Nils Chr. Framstad*†

Department of Economics, University of Oslo, Box 1095 Blindern, NO-0317 Oslo, Norway

(*Received 15 July 2013; accepted 14 October 2013*)

Consider a model for optimal timing of a policy measure which changes the emission rate, e.g. trading off the cost of reduction against the time-additive aggregate of environmental damage, the disutility from the pollutant stock $M(t)$ the infrastructure contributes to. Intuitively, the optimal timing for an *infinitesimal* pollution source should reasonably not depend on its historical contribution to the stock, as this is negligible. Dropping the size assumption, we show how to reduce the minimisation problem to one not depending on the history of M, under linear evolution and suitable linearity or additivity conditions on the damage functional. We employ a functional analysis framework which allows for delay equations, non-Markovian driving noise, a choice between discrete and continuous time, and a menu of integral concepts covering stochastic calculi less frequently used in resource and environmental economics. Examples are given under the common (Markovian Itô) stochastic analysis framework.

Keywords: optimal control; optimal stopping; environmental policy; emissions reduction; linear model; Banach space; stochastic differential equations

1. Introduction

Consider an economic activity that leads to emissions of a stock pollutant, which again leads to a certain damage per time unit, depending on the stock. We can at any time intervene in the environmental profile – for example, cease the emissions by shutting down (or possibly retrofitting with clean technology) at a given cost k, which may incorporate direct costs, scrap values and/or the loss of utility from the services the economic activity provides. The stock of a pollutant will trend upwards as long as the polluting activity persists, and then will trend downwards towards zero contribution (or at most stay put) from the time of implementation. Consider then the problem of finding the stopping (i.e. non-clairvoyant) time which minimises the expected discounted total cost and damage.

A formal model for such a problem was treated in Pindyck (2000, 2002), slightly more general than described in Section 2 below. Therein, the running cost was specified as the bilinear form $\Theta(t)M(t)$, where $M(t)$ models the pollution from the activity (as stock units) at time t, evolving through a linear differential equation, and $\Theta(t)$ models the impact cost factor, the damage per unit this (marginal) activity causes, also assumed to obey linear dynamics, more specifically a geometric Brownian motion (gBm). Below we shall review the solution of the minimisation problem, which turns out to admit reformulation into a problem which does not depend on the stock level M.

*Email: n.c.framstad+research@econ.uio.no
†Also affiliated with The Financial Supervisory Authority of Norway. The usual disclaimer applies.

This property is the motivation for this paper, as indicated in the title: when can we carry out the optimisation without paying attention to the pollutant stock level? To be a bit more specific, this paper is not about those problems which after the analysis is carried out turn out to have, e.g. the extreme optimal *strategy* of cleaning up everything immediately (no matter M) or nothing ever; the fully linear model treated in Pindyck (2000) and Pindyck (2002) can be rewritten in a way where the pollution stock is taken outside the *optimisation operation.* The question this paper sets out to answer is: without affecting this property, which model assumptions can be relaxed or removed? To what extent do we need (i) the running damage being linear (and the total damage functional affine!) in the process M? (ii) linear dynamics of M – and with our control not affecting the first-order term? (Pindyck 2002 calls for a generalisation of his specification at this point); (iii) the analogous linearity conditions on Θ? (iv) the restriction to irreversible once-and-for-all policy measures? (both Pindyck 2000 and Pindyck 2002 call for richer models at this point); (v) the dimensionality: one cost and one pollution stock? (vi) the Markovian Itô calculus framework? (vii) the (functional and/or stochastic) independence between M and Θ?

To sum up this paper's contribution, consider these questions in reverse order: possible weakening of the latter assumption will only be briefly discussed in the closing remarks, but all the assumptions that do not involve the linearity of/in M will be disposed of. We shall employ a Banach space framework which allows fairly general models including, e.g. the dispersion of a pollutant through (physical!) diffusion, and we shall allow for flexible choices of strategies and opportunity sets, as long as we assume that damage is an affine functional of the process M, which in turn is described by an affine functional equation where our control enters additively, i.e. in the zero-order term only. These restrictions – i.e. questions (i) and (ii) above – are known to be essential: more specifically, the optimal policy will depend upon the stock if the running damage is replaced by ΘM^2. There has been some confusion on this in the literature: Pindyck (2000) suggests it does not, while Pindyck (2002, Equations (43)–(45)) suggests the opposite, although both through reasoning that fails to hold true; Framstad (2011) rectifies only the former error, and, finally, Balikcioglu, Fackler, and Pindyck (2011) do actually offer a (numerical) solution. Of course, replacing «ΘM» by «ΘM^2» could equivalently be done by replacing the dynamics, so the linear evolution assumption on M cannot be disposed of.

This paper thus fills a gap on characterisation of problems and their solutions, without actually finding these solutions except in Example 5.3. The consequences of the characterisations given herein – assuming of course validity of the model structure – could be summed up as follows:

- Arguably, it is not uncommon to approach such optimisation problems making the approximation by an infinitesimal agent. We establish some robustness of such approximations: under the assumed model structure, size is not an issue.
- As the optimal policy does not depend on the pollution stock, it can also be implemented without this information. One could think that the current pollution stock, or historical records, could be easier to measure than, e.g. future damage per unit, but the optimal policy will not depend on the *future* stock level either, thus not on others' emissions. The apparent trade-off between precaution and the value of running business-as-usual while waiting for information does not at all apply to the uncertainty over future stock levels, as they have no impact on the optimal adaptation.
- The optimisation problem could be easier to actually solve. In high-dimensional systems, halving the dimensionality could make the problem tractable and dispose

of the need for other (more objectionable) simplifications to the model. It should be mentioned that once one has established the non-dependence by rewriting as in this paper, it is not at all clear that the rewritten problem is any easier to solve – the knowledge of the structure could maybe simplify the solution of the original formulation.

On the other hand, if the model is not a valid description of the real world, this paper outlines clear-cut implications which could help falsifying not merely a single model, but a model structure.

The paper is organised as follows. Section 2 reviews the simplest model, and sketches a solution method different from Pindyck (2000) using the linear structure, and takes note (Remark 2.2) of what properties are actually used to perform the operation. By formulating the model in a Banach space (i.e. a normed vector space, complete in the Cauchy sense), we shall see in Section 3 that the property of non-dependence upon history (as well as exogenous future emissions) will carry over when, in terms of the above model, «everything about M» is linear; however, it shall also turn out that there are nonlinear cases which could be of interest (see Section 4). Our approach will not be restricted to the optimal stopping problem, but will cover continuous and discontinuous optimisation over the environmental profile – increasing or decreasing emissions. Before that, we shall, however, present the solution of the above model, and point out why linearity yields the non-dependence property. The paper will return to Itô stochastic calculus in Section 5, which offers a longer example generalising the problem given in Proposition 2.1. Then, alternative stochastic calculi are briefly discussed. The exposition involves somewhat specialised terminology, but the basic idea is simple: suppose M is given by a first-order differential equation with an initial condition $M(0) = m$. Integrating the differential equation, we get an expression of the form

$$M = X + \boldsymbol{\Xi}(M),$$

where $\boldsymbol{\Xi}$ is the operator that integrates up (over time) the coefficient of the differential equation; if the differential equation is linear, then we can take $\boldsymbol{\Xi}$ as a linear operator («linear» without a constant term, as can be incorporated in X). The functional analysis framework does provide tools for computing as if $\boldsymbol{\Xi}$ were a matrix Ξ – in which case one could solve for M, uniquely, under the invertibility of $(I - \Xi)$. The Banach fixed-point theorem expresses the inverse as a geometric series, and states conditions for the existence through the convergence of this series. Taken this machinery for given, the reader can appreciate the main content of the paper using merely linear algebra rules.

2. The simple optimisation problem

This section will briefly review a simplified version of the Pindyck (2000) model, henceforth the «simple model». Suppose that the aggregate discounted economical and ecological cost associated with a given activity to be closed down at time τ at cost k is given by

$$\int_0^\infty e^{-rt}\Theta(t)M(t)\mathrm{d}t + ke^{-r\tau}, \tag{1}$$

where the integral term aggregates the damages from stock M of the pollutant, assumed to evolve according to the linear differential equation

$$\mathrm{d}M = (\beta\eta(t) - \delta M)\mathrm{d}t, \qquad M(0) = m, \tag{2}$$

where $\eta(t)$ could be interpreted as the emission rate and β is the fraction that finds its way to the environment. η will in the simple model be restricted to being a positive constant until implementation, and 0 thereafter (with M continuous at the implementation time). The initial state $M(0) = m$ is assumed non-negative.

The optimisation can be solved by dynamic programming. Leaving a formal setup for Section 3, we shall now restrict ourselves to the result for the problem where $\eta(t)$ is a given constant up to the intervention time and 0 thereafter, henceforth a «one-shot problem».

Proposition 2.1 (Pindyck 2000): *Consider M following the dynamics* (2) *with $\beta = 1$ and $\eta(t)$ restricted to the form $\eta(t) = \overline{\eta}\cdot 1_{t\in[0,\tau]}$, with M continuous at τ and C^1 elsewhere. Suppose furthermore that Θ obeys the (Itô) stochastic differential equation*

$$\mathrm{d}\Theta(t) = \Theta(t)\cdot(\alpha\mathrm{d}t + \sigma\mathrm{d}Z(t)), \qquad \Theta(0) = \theta > 0, \tag{3}$$

where Z is standard Brownian, and that $r - \alpha > 0$, $\delta > 0$ and $k = k(\overline{\eta}) > 0$. Then, the problem of minimising the expected value of (1) *over all stopping times τ is solved by stopping first time Θ exceeds*

$$\theta^* = \frac{\gamma k(\overline{\eta})}{(\gamma-1)\overline{\eta}}(r-\alpha)(r+\delta-\alpha) \qquad \textit{with} \quad \gamma = \frac{1}{2} - \frac{\alpha}{\sigma^2} + \sqrt{\left(\frac{\alpha}{\sigma^2} - \frac{1}{2}\right)^2 + \frac{2r}{\sigma^2}}, \tag{4}$$

and the value function is

$$\frac{\theta m}{r+\delta-\alpha} + k(\overline{\eta})\cdot\begin{cases} \left[\gamma\dfrac{\theta}{\theta^*} - \left(\dfrac{\theta}{\theta^*}\right)^{\gamma}\right]\Big/(\gamma-1) & \text{if} \quad \theta < \theta^* \\ 1 & \text{if} \quad \theta \geq \theta^*. \end{cases} \tag{5}$$

In particular, the optimal rule does not depend on m.

2.1. Removing the stock from the optimisation

Notice in the solution of the simple model that the m-dependent first term is the damage that would incur even without the installation ($\tau = k = 0$). The term $(\theta/\theta^*)k\gamma/(\gamma-1) = \theta\overline{\eta}/\big((r-\alpha)(r+\delta-\alpha)\big)$ is the damage from running the installation forever from now, and the subtractive element represents the value of the option to stop. The latter two do not depend on m. Let us give an argument for this property without using dynamic programming nor the form (5) directly – although we can later use dynamic programming on the one-dimensional problem we reduce the problem to. The solution for $M(t)$ is

$$M(t) = e^{-\delta t}m + e^{-\delta t}\int_0^t e^{\delta s}\eta(s)\,\mathrm{d}s \tag{6}$$

so that when we restrict controls to $\eta(s) = \overline{\eta}1_{s\leq\tau}$ as above, with only $\tau \geq 0$ to choose, damage and intervention cost add up to

$$m\int_0^\infty e^{-(r+\delta)t}\,\Theta(t)\,\mathrm{d}t + \overline{\eta}\int_0^\infty e^{-(r+\delta)t}\,\Theta(t)\left(\int_0^{\min\{t,\tau\}} e^{\delta s}\mathrm{d}s\right)\mathrm{d}t + ke^{-r\tau}. \tag{7}$$

Already before applying mathematical expectation, we observe that the first term – the damage obtained by running forever – does not depend on τ, while the two others do not depend on m.

We can actually compute the distribution and expectation of (7) without dynamic programming, looking up the distributional properties of the gBm and its stopping times in, e.g. Borodin and Salminen (1996). However, we could just as well reformulate in terms of a minimisation problem which does not depend on m, and then guess and verify by means of the Bellman equation. Let us actually write out this, reorganising (7), as

$$\left(m - \frac{\overline{\eta}}{\delta}\right)\int_0^\infty e^{-(r+\delta)t}\,\Theta(t)\,\mathrm{d}t + \int_0^\infty e^{-rt}\,\underbrace{\Theta(t)e^{\delta\min\{0,\tau-t\}}\cdot\overline{\eta}/\delta}_{=:F(t)}\,\mathrm{d}t + ke^{-r\tau}.$$

We then have that F satisfies

$$\mathrm{d}F(t) = F(t)\cdot\left[(\alpha - \delta 1_{t\in[0,\tau]})\,\mathrm{d}t + \sigma\,\mathrm{d}Z(t)\right], \qquad F(0) = \theta\overline{\eta}/\delta\,.$$

For the minimisation of the expectation of the latter two terms, the Bellman equation takes the form

$$-rV + \alpha yV' + \frac{1}{2}\sigma^2 y^2 V'' + q = \begin{cases} \delta yV' & \text{before intervention} \\ 0 & \text{after intervention.} \end{cases}$$

The rest is routine: solve, find the strategy as a trigger y^* by a C^1 fit, and if one wishes a fully rigorous proof, do the limit transition to an infinite horizon.

Remark 2.2: It is easy to see that (at least under the appropriate integrability conditions) the optimisation will not depend of m even under the following generalisations:

(a) The argument on M uses, essentially, only an integrating factor approach, i.e. linearity of the differential equation.

(b) Under linearity, it does not matter whether M models the total stock of the pollutant, or if it models the project's contribution: suppose the latter, and denote everyone else's contributions by L also driven by a linear differential equation with the same decay rate, then we can either formulate in terms of the total $(L + M)$ (also following a linear differential equation with the same decay, though likely with different emission levels) or split up. In either case, the optimisation will separate into a term not depending on the decision and one not depending on L or M (which enter only through the sum). In the problem of Proposition 2.1, it means that the optimisation amounts to minimising $k(\eta)\cdot\left[\gamma\min\{1,\theta/\widehat{\theta}\} - \min\{1,\theta/\widehat{\theta}\}^\gamma\right]/(\gamma - 1)$ – which is solved by choosing $\widehat{\theta} = \theta^*$ according to (4).

(c) We need not restrict Θ to be geometric Brownian. It can be any exogenously given stochastic process which does not depend on M, and the optimal rule will still only depend on the history and future law of Θ (that is, its state if it is autonomous Markovian).
(d) k can be allowed to depend on Θ as long as it does not depend on M.
(e) The discount rate need not be constant nor deterministic, as long as it does not depend on M nor our control. In the above case, $e^{-rt}\Theta$ is geometric Brownian too, and we could therefore merge the discount factor into the Θ process – but then the intervention cost would have to be represented as a process too. That is, however, not an issue.
(f) The optimisation problem need not be restricted to optimal stopping – we can replace the cost $ke^{-r\tau}$ by a cost *process* associated with $\eta(t)$ and Θ (possibly history-dependent), as long as this does not depend on M.

As an example of the latter, Framstad and Strand (2013) extend the Pindyck model by endogenising the initial infrastructure, then subject to a utility and an investment cost. In particular, the initial investment will, as shown above, not depend on m. The paper also discusses extensions like endogenising *timing* of the initial investment, and availability of emissions-reducing technology to be retrofitted to the installation. As long as these quantities do not depend on m nor the subsequent development of M, then neither will those decisions of initial technology and (if controllable) its timing. Also extending to a model with *gradual* build-up of infrastructure and subsequent reduction of emissions will have decision rules not depending of M.

Let us remark that although the question of dependence upon the stock does naturally not apply to pollutants for which only the flow causes damage, it may for actually solving the optimisation problems be useful to compare fast-decaying stocks with the limiting case of flows. The equivalence of flow and stock in terms of non-dependence upon the stock level is also exploited by, e.g. Harstad and Battaglini (2012) in a model for environmental agreements.

Section 3 will formalise the property discussed, as well as more general linear models for which it carries over.

3. Linear evolutionary equations in Banach spaces

We shall consider a more general setup than Section 2. We replace the initial state m by a given function X unaffected by the control (hence the letter X for eXogenous). This will cover, e.g. delay equations for which the evolution depends on the past, in which case X is given initially as the history, the path $t \mapsto \{M(t)\}_{t\leq 0}$. Analogously, we replace the initial state θ for Θ by an arbitrarily dimensional parameter denoted by G – for example, this could be the history $t \mapsto \{\Theta(t)\}_{t\leq 0}$.

Let us first fix some terminology.

Definition 3.1: We shall use the term «*does not depend on*» to mean invariance under partial shift, e.g. functional independence (contrasted to stochastic independence). Terms like «*might depend on*», «*will usually depend on*» and «*does depend on*» should be self-explanatory.

The property is now given slightly informally.

Definition 3.2: Consider a minimisation problem indexed by $(G,X) \in \mathbb{G} \times \mathbb{X}$. We shall say that the problem *does not depend on* X if, for each $G \in \mathbb{G}$, the ordering of the controls according to performance does not depend on $X \in \mathbb{X}$.

We shall in practice look for decompositions of the form

$$\begin{aligned} &[\text{something which does not depend on the control}] \\ &+[\text{non-negative functional of } X]\cdot[\text{functional which does not depend on } M] \end{aligned} \tag{8}$$

(where again, «does not depend on M» means functional independence, even when certain values for X may be deduced from M without knowing the control). The first line will be analogous to the «$\theta m/(r+\delta-\alpha)$» damage which incurs even without the pollutant source in question. Usually, the «non-negative functional of X» will be a constant; however, an example where it is not will be given in (13)–(15). It should be noted that Definition 3.2 is weaker than the property that the *optimal strategy* does not depend on X; consider the example from the previous section – if one replaces $ke^{-r\tau}$ by some function non-decreasing in τ, then the optimal choice will be $\tau = 0$ for all $m \geq 0$, even when the M under the integral is replaced by any positive non-decreasing function of M.

3.1. Sufficient conditions for the optimisation problem not to depend on the exogenous X

The following simple application of the Banach fixed-point theorem essentially sums up why the linear cases behave as they do. We could push the result further by making the assumption of a left-inverse $(\mathbf{I}-\mathbf{\Xi})^{-1}$ ad hoc, but the following framework is sufficiently general for the applications.

Lemma 3.3: *Consider a Banach space* $\mathbb{M}$*, with a bounded linear operator* $\mathbf{\Xi} : \mathbb{M} \to \mathbb{M}$ *such that some power is a contraction, and a linear functional* $\mathbf{\Phi} : \mathbb{M} \to \mathbb{R}$*. Then, for* $Q \in \mathbb{M}$*, the unique solution* M *of the functional equation*

$$M = Q + \mathbf{\Xi} M \tag{9}$$

is

$$M = \mathbf{\Psi} Q, \qquad \text{where} \qquad \mathbf{\Psi} := \mathbf{I} + \sum_{j=1}^{\infty} \mathbf{\Xi}^j = (\mathbf{I}-\mathbf{\Xi})^{-1} \tag{10}$$

is a well-defined bounded linear operator from $\mathbb{M}$ *onto* $\mathbb{M}$*. Furthermore,* $\mathbf{\Phi} M = (\mathbf{\Phi}\mathbf{\Psi})Q$*, a linear functional acting on* $Q \in \mathbb{M}$*.*

Assume a linear structure $Q = A + X$, where A is a (fully or partially) controllable component (in dynamic systems, interpretable as controlling the future evolution) and X does not depend on the control chosen, and the following consequence is immediate.

Proposition 3.4: *Given a set* $\mathbb{A} \subseteq \mathbb{M}$*, a (possibly nonlinear) functional* $\mathbf{\Gamma} : \mathbb{A} \to \mathbb{R}$*, a linear operator* $\mathbf{\Xi} : \mathbb{M} \to \mathbb{M}$ *with some power being a contraction, and some* $X \in \mathbb{M}$*. Then, the problem*

$$\inf_{A \in \mathbb{A}} \{\mathbf{\Phi} M + \mathbf{\Gamma}(A)\} \qquad \text{subject to} \quad (9) \text{ and } Q = A + X \tag{11}$$

can be rewritten as

$$(\mathbf{\Phi\Psi})X + \inf_{A\in\mathbb{A}} \{(\mathbf{\Phi\Psi})A + \mathbf{\Gamma}(A)\} \tag{12}$$

where the latter optimisation problem does not depend on X.

Notice again that it does not matter whether M models the project's emissions or the total emissions, as long as the evolution is modelled by the action of a linear operator. Under linearity, it does not matter whether this «evolution» is actually in (univariate) time:

Remark 3.5: The motivating simple model of Section 2 concerned aggregate damage over time. However, there is nothing in Proposition 3.4 that precludes time–space aggregates; the vector M could be of arbitrary dimension, including space indexing dimensions, and the canonical model for the dissemination of a pollutant in physical space – the heat/diffusion equation – is of course linear.

It is, however, crucial that the $\mathbf{\Xi}$ operator is not controlled.

Remark 3.6: Attempting to «fix» linearity by augmenting with more terms will violate the crucial exogeneity of the $\mathbf{\Xi}$ operator, which we need to keep the first term of (12) outside the optimisation. Let, for example, the model be $\dot{M} = \overline{\eta}1_{[0,\tau]} - \delta_2 M^2 - \delta_1 M$, so that M is trapped in the unit interval. Then, we attempt to introduce an infinite-dimensional linear model with coordinates $M_i(t) =$ the ith power of $M(t)$; it easily follows by induction that each $\dot{M}_i$ can be written as a polynomial in M, hence as a finite linear combination of the coordinates. However, then $\dot{M}_2 = 2M\dot{M} = 2M_1\overline{\eta}1_{[0,\tau]} - \delta_2 M_3 - \delta_1 M_2$, and the first term makes the new infinite-dimensional $\mathbf{\Xi}$ dependent on control – and that ruins the non-dependence argument even if the said dependence occurs only in coordinates which do not enter the running damage!

Let us work out how to fit the model of Proposition 2.1 into the applicability of Proposition 3.4. The key is the contraction property established in the usual Picard–Lindelöf iteration to hold locally, and just as in that argument, we can apply the following piecewise:

- Our control A is now the *cumulative* emissions, the function $t \mapsto \int_0^t \eta(s)\,\mathrm{d}s$. For the problem of Proposition 2.1, $\mathbb{A}$ is the set of functions of the form $\overline{\eta}\min\{t,\tau\}$ for some stopping time $\tau \geq 0$.
- The X function is the constant m.
- $\mathbf{\Xi}$ takes as input the function $t \mapsto M(t)$ and returns the function $t \mapsto -\delta \int_0^t M(s)\,\mathrm{d}s$.
- The aggregate damage is $\mathbf{\Phi}$, whose dependence on the initial condition $G = \theta \in [0,\infty)$ is notationally suppressed.

The «piecewise» version of this is carried out on the partition $0 = t_0 < t_1 < \cdots$, with t_i defined as the first time contraction failed in the previous step, by shrinking $\mathbb{M}$ to the subspace/quotient space of functions with the known past $t \mapsto M(\min\{t, t_i\})$.

Remark 3.7: As mentioned in Section 2, a strategy for flow pollutants will not depend on the state. In the case where the flow incurs a cost – ecological damage or Pigouvian tax – then this is covered by the $\mathbf{\Gamma}$ functional, as it takes as input the entire emissions path and hence can depend on the time-derivative. Nevertheless, it could be of interest to consider a flow as a limit of a fast-decaying stock. One can then let $\boldsymbol{\Xi}$ represent a fast decay, and renormalise Q. Letting $Q = X + \frac{1}{\epsilon}\tilde{A}$ and $\boldsymbol{\Xi} = \frac{1}{\epsilon}\tilde{\boldsymbol{\Xi}}$, M is solved by

$$M = -(\tilde{\boldsymbol{\Xi}} - \epsilon\mathbf{I})^{-1}(\epsilon X + \tilde{A})$$

in terms of the *resolvent* $(\tilde{\boldsymbol{\Xi}} - \epsilon\mathbf{I})^{-1}$of $\tilde{\boldsymbol{\Xi}}$. If 0 is in the closure of the resolvent set, we can then let $\epsilon \to 0$ through an appropriate sequence. Using again the problem of Proposition 2.1, as an example, we can take $\delta = 1/\epsilon$ and $\tilde{\eta} = \overline{\eta}/\epsilon$; then $\tilde{A} = -\tilde{\eta}\,\tilde{\boldsymbol{\Xi}}1_{[0,\tau]} = \tilde{\eta}\cdot\big((\epsilon\mathbf{I} - \tilde{\boldsymbol{\Xi}})1_{[0,\tau]} - \epsilon 1_{[0,\tau]}\big)$, so that

$$\begin{aligned} M &= \epsilon(\epsilon\mathbf{I} - \tilde{\boldsymbol{\Xi}})^{-1}(X + 1_{[0,\tau]}) + (\epsilon\mathbf{I} - \tilde{\boldsymbol{\Xi}})^{-1}(\epsilon\mathbf{I} - \tilde{\boldsymbol{\Xi}})\tilde{\eta}1_{[0,\tau]} \\ &\to \tilde{\eta}1_{[0,\tau]} \qquad \text{as}\,\epsilon \searrow 0,\ \text{i.e. as}\,\delta \to +\infty \end{aligned}$$

Thus, in the limit, M is precisely the flow expressed as a limit of normalised stocks.

4. Some nonlinear cases

Linearity turns out not to be a necessary condition for the property of Definition 3.2; we have merely used that X splits out additively upon application of the $(\boldsymbol{\Phi\Psi})$ functional. Furthermore, there are examples where not even this additivity holds, and where still the optimisation does not depend on X. The question is rather whether these are to be considered merely degeneracies. Of course, that is a matter of definition and opinion – e.g. one would likely consider it a degeneracy if one ad hoc, for the purpose of creating an example, restricts the set of controls in just in order to satisfy the requirement of Definition 3.2. Also, we have not ruled out $\boldsymbol{\Phi}$ functionals which do depend on X explicitly; for example, one can construct a cancelling of X by $\boldsymbol{\Phi} = \tilde{\boldsymbol{\Phi}}(\mathbf{I} - \boldsymbol{\Xi})$, i.e. $\boldsymbol{\Phi}Q = \tilde{\boldsymbol{\Phi}}M$, and $M \mapsto \tilde{\boldsymbol{\Phi}}M$ need not depend on X. For example, in the language of Proposition 2.1, a functional that takes as input the path of $e^{\delta t}M(t) - M(0)$ will yield an expression which does not depend on $M(0) = m$. The integral criteria of Sections 2 and 5 would not be prone to these kinds of constructed degeneracy, though.

The following will consider some cases which are nonlinear, but where the optimisation still does not depend on the initial stock. The first quadratic case is arguably the more «degenerate».

4.1. A quadratic case

There turn out to be quadratic cases where $\boldsymbol{\Gamma}$ cancels out the part which does *not* depend on X, leaving one which does, but in a way that might leave the *optimisation* not depending. Suppose that the objective to be minimised is no longer linear, but involving a quadratic: $\boldsymbol{\Phi}M + \langle M, M\rangle + \boldsymbol{\Gamma}(A)$, for some suitable bilinear form $\langle\,\cdot\,,\,\cdot\,\rangle$. Suppose now the particular form where $\boldsymbol{\Gamma} = \Gamma_0 - \boldsymbol{\Phi\Psi}A - \langle A, A\rangle$, where Γ_0 does not depend on A. We still assume the linear $M = \boldsymbol{\Psi}(A + X)$. Then, the minimisation problem becomes

$$\boldsymbol{\Phi\Psi}X + \langle\boldsymbol{\Psi}X, \boldsymbol{\Psi}X\rangle + 2\inf_{A\in\mathbb{A}}\langle\boldsymbol{\Psi}A, \boldsymbol{\Psi}X\rangle \tag{13}$$

which is linear in X. If now the dependence on X separates out as a single univariate multiplicative factor, the minimisation does not depend on this as long as we restrict to the non-negative range. Let us again take the example from Proposition 2.1 as a starting point. We modify the objective function (before applying the expectation) by replacing the linear integrand by the quadratic $e^{-rt}\,\Theta(t)\,(M(t))^2$ – as long as we assume $m \geq 0$, this is increasing in M – and then replacing $ke^{-r\tau}$ by the functional

$$\mathbf{\Gamma} = \overline{\eta}^2 \int_0^\infty e^{-(r+2\delta)t}\,\Theta(t)\left[\left(\int_0^t e^{\delta s}\,\mathrm{d}s\right)^2 - \left(\int_0^{\min\{t,\tau\}} e^{\delta s}\,\mathrm{d}s\right)^2\right]\mathrm{d}t \tag{14}$$

assuming r big enough to keep everything finite. It has some properties in common with the problem of Proposition 2.1; it is decreasing in τ, but positive whenever $\tau < \infty$ and $\theta > 0$. We can simplify the minimand to

$$\int_0^\infty e^{-rt}\,\Theta(t)\Big(M(t)\Big)^2\,\mathrm{d}t + \mathbf{\Gamma}$$
$$= \overline{\eta}^2 \int_0^\infty e^{-(r+2\delta)t}\,\Theta(t)\left(\int_0^t e^{\delta s}\,\mathrm{d}s\right)^2\,\mathrm{d}t + 2m\overline{\eta}\int_0^\infty e^{-(r+2\delta)t}\,\Theta(t)\left(\int_0^{\min\{t,\tau\}} e^{\delta s}\,\mathrm{d}s\right)\,\mathrm{d}t \tag{15}$$

and the minimisation does not depend upon *non-negative* m. Neither does the minimum, trivially obtained by $\tau = 0$. Again, we need to emphasise that establishing the «does not depend» property of Definition 3.2 is more modest than the corresponding property of the optimal strategy: $\mathbf{\Gamma}$ and $\mathbf{\Phi}$ could have a common minimand without applying Definition 3.2, and there could be other cases where there is a «corner solution» (i.e. for one-shot timing problems: such that the optimal τ is a.s. 0 or a.s. $+\infty$).

It seems tempting to guess that nonlinear $\mathbf{\Phi}$ will lead to an optimisation problem depending on X except degenerate cases – for a suitable opinion on «degenerate». The next subsection will consider a class where the linearity condition is weakened.

4.2. Additivity over the exogenous or over the controlled part, and worst-scenario optimisation

Suppose that we are not optimising a problem like (12), but, rather than with one fixed linear functional ($\mathbf{\Phi\Psi}$), we are given a family of functionals with the criterion being to optimise over the «worst case». In the following, a range of functionals $\mathbf{\Lambda}$ will replace the single $\mathbf{\Phi\Psi}$, with you playing against a worst-case opponent $\mathbf{\Lambda} \in \mathbb{L}$. Modify the setup of Proposition 3.4 such that the objective is

$$\inf_{A\in\mathbb{A}}\ \sup_{\mathbf{\Lambda}\in\mathbb{L}}\{\mathbf{\Lambda}(X+A) + \mathbf{\Gamma}_{\mathbf{\Lambda}}(A)\}. \tag{16}$$

Before giving examples, the reader should note that this kind of criteria will be better suited to model Knightian uncertainty than expected utility with a universally agreed upon probability measure. In this sense, a «scenario» is not an event, but should rather be represented by a probability law, and worst-scenario optimisation would correspond to picking the «most pessimistic» model. Now, which scenario is worst does (usually) depend on our control. Assuming an argmax $\mathbf{\Lambda}^*$ in (16) exists, then it is in general a function of A, and the subsequent minimisation with respect to A will depend on X except in special cases. To illustrate, consider the following very simple example.

Example 4.1: Consider the one-shot problem of Proposition 2.1, modified as follows: the intervention cost is uncertain, with disagreement over its distribution, but it is agreed upon that it is stochastically independent of everything else, and revealed only once the intervention is made. Everyone agrees on the other elements. We can then identify the various probability scenarios with their respective expectation for this cost, and use that for the «k» variable. Obviously, a higher k is bad for every strategy – strictly so except for the «never intervene» strategy. The worst-case scenario does thus not depend on the strategy – it is the highest k. With (16) as a criterion, we can thus simply plug the highest k in, and we are back in the one-shot problem.

This model is also applicable for a problem where implementation is carried out only if and when two (or more!) decision-makers unanimously agree upon doing so. Suppose that they agree upon everything but the distribution over implementation cost as above; then, the decision is made by the one who sees the highest trigger value θ^*, i.e. the highest k, i.e. the worst scenario; thus, we are within the framework of (16).

This example simplifies because it does not depend on the strategy which scenario is the worst. That need of course not be the case. Suppose, for example, that Θ is not geometric Brownian, and that it is uncertain at high levels which it has not yet hit or spent enough time for a reliable estimate. (Although Brownian volatility can be measured accurately by quadratic variation, rare-jump uncertainty is hard to assess without long time series, and if the optimal rule for a given model is to stop at first hitting time, then heuristically, we will at the critical range where we consider stopping not have data to estimate from.) If now the scenarios involve beliefs for the model for Θ and jointly the distribution for the intervention cost, it need not be so that pessimism for one of the entities corresponds to pessimism on the other. What is the worst scenario might then depend on the candidate for strategy.

In such situations, where there is no universal «worst» scenario, it is of course tempting to try to apply a minimax theorem to justify commuting the inf and the sup. Then, the minimisation would be done $\mathbf{\Lambda}$-wise, where the X splits out. This is the last part of the following proposition, which gives conditions for the minimisation not to depend on X (though not for the maximisation for the worst case, which could still depend on X!).

Proposition 4.2: *Consider the problem*

$$\inf_{A\in\mathbb{A}} \sup_{\mathbf{\Lambda}\in\mathbb{L}} \{\mathbf{\Lambda} X + \mathbf{\Lambda} A + \mathbf{\Gamma}_{\mathbf{\Lambda}}(A)\} \tag{17}$$

where $\mathbb{L}$ is a given family of linear functionals $\mathbf{\Lambda} : \mathbb{M} \to \mathbb{R}$ and A is our control. The potentially nonlinear functional $\mathbf{\Gamma} : \mathbb{A} \to \mathbb{R}$ can depend on $\mathbf{\Lambda}$, although in part (i) below we will assume it does not:

(i) If $\mathbf{\Lambda} \mapsto \mathbf{\Lambda} A$ is constant on $\mathbb{L}$ for each A, and furthermore $\mathbf{\Gamma}_{\mathbf{\Lambda}}$ is constant with respect to $\mathbf{\Lambda}$ for each A, the optimisation problem (17) *reduces to*

$$\inf_{A\in\mathbb{A}} \{\mathbf{\Lambda}_0 A + \mathbf{\Gamma}_{\mathbf{\Lambda}_0}(A)\} + \sup_{\mathbf{\Lambda}\in\mathbb{L}} \mathbf{\Lambda} X \tag{18}$$

for an arbitrarily chosen $\mathbf{\Lambda}_0 \in \mathbb{L}$.

(ii) If instead $\mathbf{\Lambda} \mapsto \mathbf{\Lambda} X$ *is constant on* $\mathbb{L}$ *for each* X*, then* (17) *reduces to*

$$\mathbf{\Lambda}_0 X + \inf_{A \in \mathbb{A}} \sup_{\mathbf{\Lambda} \in \mathbb{L}} \{\mathbf{\Lambda} A + \mathbf{\Gamma}_{\mathbf{\Lambda}}(A)\} \tag{19}$$

again, for an arbitrarily chosen $\mathbf{\Lambda}_0 \in \mathbb{L}$.

(iii) Alternatively, suppose that $\mathbb{L}$ *is convex and* $\mathbb{A}$ *is convex and compact, and that* $\mathbf{\Lambda} \mapsto \mathbf{\Gamma}_{\mathbf{\Lambda}}(A)$ *is upper semicontinuous and quasiconcave and* $A \mapsto \mathbf{\Gamma}_{\mathbf{\Lambda}}(A)$*is lower semicontinuous and quasiconvex. Then, (17) can be written as*

$$\sup_{\mathbf{\Lambda} \in \mathbb{L}} \left\{ \mathbf{\Lambda} X + \min_{A \in \mathbb{A}} \{ \mathbf{\Lambda} A + \mathbf{\Gamma}_{\mathbf{\Lambda}}(A)\} \right\}, \tag{20}$$

provided the inner minimum is attained.

Neither of these three minimisation problems with respect to A *depends on* X.

Proof: The first two parts are self-evident. The assumptions for part (iii) are those of Sion's generalisation of the celebrated von Neumann minimax theorem (cf. e.g. Komiya 1988). □

The interpretation for probability scenarios is as follows: in case (i), everyone agrees upon the consequences of our actions, *given* the exogenous X component (e.g. the history), which splits out linearly. In case (ii), everyone agrees upon the damage incurred *from* X, e.g. if it is known history with known consequences. For the last case (iii), we are in a commuting minimax situation; it should then be remarked that the convexity of $\mathbb{A}$ could be a significant restriction; for the one-shot case, the set of functions of the form $A(t) = \overline{\eta} \min\{t, \tau\}$ is not convex, and one will have to extend the problem for a hope to apply part (iii).

It is worth mentioning that worst-scenario optimisation has a connection with the so-called *risk measures* of mathematical finance, where convex preferences over Knightian uncertainties (often called «model risks») can be represented through worst-case Knightian risks with a «penalty function» adjusting for the credibility of the scenarios. See, e.g. Föllmer and Schied (2002a, 2002b) or Frittelli and Gianin (2002).

5. Linear (stochastic) differential equations

This section concerns cover the case where M is governed by a stochastic differential equation. The standard existence and uniqueness result is a Picard–Lindelöf argument, which of course applies under linearity, so it is a special case of Proposition 3.4, but if we want to *solve* the problem (11) from the representation (12), we would want to write out the latter. Notice though that it is not necessarily desirable to work with (12) – especially as Proposition 3.4 provides us with the information that we can look for an optimal rule not depending on the state of M. The next subsection will cover the semimartingale Itô differential case, and then other integrals will be sketched in Section 5.2.

5.1. Itô stochastic differential equations

Let us fix the setup and notation for this subsection:

- We shall work on a (notationally suppressed) usual stochastic basis; namely, a probability space equipped with a right-continuous filtration complete at time zero.

- Vectors are column vectors, unless indicated by the transposition superscript†. The symbol $\cdot$ denotes the Euclidean inner product on $\mathbb{R}^d$, but will be used for products between scalars as well.
- For stochastic processes, denote by superscript $^{\mathsf{c}}$ the continuous part, and for discontinuities: $\Delta^+ Y = Y(t^+) - Y(t)$, $\Delta^- Y = Y(t) - Y(t^-)$. We will use Δ when the interpretation is unambiguous. Furthermore, we use accents $\grave{Y}$ $(\acute{Y})$ for the left-continuous version (right-continuous version) of Y.
- The reader should be aware that as matrix products do not commute, notations like $\mathrm{d}\Pi(t)M(t)$ may be necessary even when M will be part of the integration.
- Differentials denote Itô-type integration.

We now specify the objective function; we assume that the optimisation problem is to minimise the expected value of the functional

$$\int_0^{\infty} \big(M(t^-)^{\dagger}\mathrm{d}D(t) + \mathrm{d}C(t)\big), \tag{21}$$

where we shall, ad hoc, assume integrability. Since we are in a multidimensional setting, the discount factor e^{-rt} has been incorporated into the processes D and C (mnemonics: *D*amage from the pollutant, *C*ost of control).

The modelling building blocks are the following entities:

Measurability/continuity assumptions on the processes. We assume that all processes are adapted, and their sample paths possess both left and right limits. In addition, the following are standing assumptions:

(a) The $\mathbb{R}^d$-valued process D, assumed right continuous, is an exogenously given (uncontrolled) process which aggregates the environmental damage of the pollutant stock. $\mathrm{d}D$ specialises $\mathbf{\Phi}$ but generalises $e^{-rt}\Theta(t)\ \mathrm{d}t$.

(b) M will be the pollutant stock, which we can affect through a predictable, hence assumed left-continuous, control denoted by S; we introduce this for the sake of interpretation, although we will not write down the explicit way it enters. $t \mapsto M$ need not be left or right continuous. We shall introduce a driving process Q for M, and the following measurability/continuity conditions will apply for Q and for M:
   - On intervals where S is constant, M will be assumed right continuous.
   - At discontinuity times T for S, henceforth *interventions*, $S(T)$ does not affect $M(T)$, only $M(T^+)$. In other words, we have that $M(T)$ does not depend on $\Delta^+ S = S(T^+) - S(T)$, which may in turn depend on the past up to and including T (the T-measurability of this difference is the assumed predictability), and which affects $\Delta^+ M$. Due to the assumed right continuity of the filtration, $M(T^+)$ is T-measurable; at time T, we know our intervention $\Delta^+ S$, and there is no randomness drawing the right limit.

(c) The process C – specialising the $\mathbf{\Gamma}$ functional but generalising $ke^{-r\tau}$ – is the incurred cost of the control, allowed to depend on the past (subject to assumptions specified below), in particular, the entire past path of S, but we shall below assume it does not depend on X. As C is only an integrator for the continuous discount factor, we do not need to worry about left-hand or right-hand jumps, as discontinuities contribute only through $C(T^+) - C(T^-)$. We can therefore work with any version.

The dynamics for M will in this subsection be assumed to obey the following form (in terms of transposes, to get the differential post-multiplied):

$$\mathrm{d}M^\dagger(t) = M^\dagger(t^-)\ \mathrm{d}\Xi^\dagger(t) + \ \mathrm{d}Q^\dagger(t), \qquad M(0) = m, \tag{22}$$

where X is fully represented through m, and the Ξ functional is specified as integration with respect to the given $\mathbb{R}^{d\times d}$-valued right-continuous process Ξ. In order to fit to the setup, put $Q(0) = m$ and $A = Q - m$; then, the $\mathbb{R}^d$-valued process $Q - Q(0) = Q - m$ can be influenced by the control (but shall not depend on M or m).

We then have the following.

Proposition 5.1: *Suppose that for each given control S, the following holds: Ξ, D, Q and C are given semimartingales, the two first right-continuous, and that the jumps satisfy, with probability 1,*

$$\Delta\Xi\Delta\acute{Q} \in \textit{column space } (I + \Delta\Xi), \textit{ all jumps}. \tag{23}$$

Suppose furthermore that M uniquely (up to version) solves the (Itô) stochastic differential equation (22)*. Then, there exists some $\mathbb{R}^{d\times d}$-valued semimartingale Π, given by*

$$\Pi(0) = I, \qquad \mathrm{d}\Pi(t) = \big(\mathrm{d}\Xi(t)\big)\ \Pi(t^-) \tag{24}$$

– form with time-differentials post-multiplied: $(\Pi^\dagger\ \mathrm{d}\Xi^\dagger)^\dagger$ – such that (21) *equals, if at least one integral converges,*

$$m^\dagger \int_0^\infty \Pi(t)^\dagger\ dD$$
$$+ \int_0^\infty \left(\left[\Delta^+ Q(0) + \int_{(0,t]} \Pi(s^-)^{-1}\ \big(\acute{Q}(s) - Y(s)\big) \right]^\dagger \Pi(t)^\dagger\ dD + \ dC(t) \right), \tag{25}$$

where Y is a right-continuous process such that, in terms of Itô differentials,

$$\mathrm{d}Y^{\mathsf{c}} = \ \mathrm{d}\Xi^{\mathsf{c}}\ \mathrm{d}Q^{\mathsf{c}}, \qquad (I + \Delta\Xi)\Delta Y = \Delta\Xi\Delta\acute{Q}. \tag{26}$$

In particular, if Ξ and D do not depend on S, then the minimisation over S does not depend on M.

Proof: Notice first that (23) ensures that (26) can be satisfied even when $\Delta\Xi$ has an eigenvalue of -1. Now M enters directly the objective only through the left-continuous version. We can therefore first do the differential calculus on the right-continuous version $\acute{M}$, which satisfies (22) except at intervention times. Notice that $\acute{M}(0) = m + \Delta^+ M(0) = m + \Delta^+ Q(0)$, and $\Delta^+ Q(0)$ depends solely on $\Delta^+ S(0)$. We claim that

$$\acute{M}(t) = \Pi(t)m + \Pi(t)\left[\Delta^{+}Q(0) + \int_{(0,t]} \Pi(s^{-})^{-1}\ \mathrm{d}\big(\acute{Q}(s) - Y(s)\big)\right]. \tag{27}$$

To see this, differentiate using the Itô formula. Suppressing time arguments,

$$\mathrm{d}\acute{M} = (\mathrm{d}\Xi)\grave{M} + \Pi\Pi^{-1}\ \mathrm{d}(\acute{Q} - Y) + (\mathrm{d}\Xi)\big(\mathrm{d}(\acute{Q} - Y)\big) \tag{28}$$

where the latter term is the cross-variation expressed as Itô differentials. Now, cancel terms using (26). □

Remark 5.2: A few comments are appropriate.

(a) The restriction (23) applies in the case where the jump amplitude could possibly have -1 as an eigenvalue (in which case, at the eigenvector, Ξ would cause a jump to null). It limits the possible jumps Q could make at the same time. The condition will be satisfied whenever a coordinate of M jumps to a state not depending on the pre-jump state $M(T^{-})$; then, the new post-jump state is zero. It thus covers cases where the pollutant should vanish at a jump (say, if we are modelling a case where some exogenous agent could at some point T choose to clean up).

(b) The model and result admit cases where we actually intervene in the pollution stock, but only through Q – that is, in absolute numbers, not in percentages. The intervention does not depend on the level; if damages and costs are so that it pays off to remove 1 unit of the pollutant, then that decision does not depend on the stock level. This may be objectionable when M models a cardinal level, as it could bring M outside the first orthant – and capping the cleansing operation to keep stocks non-negative would mean that the strategy takes the state of M into account. This objection does, however, not apply to a discrete emission which instantly *increases* one or more coordinates of M.

(c) These discrete interventions in Q may seem to be not captured in the above argument, but no information is lost. Even if the proof only uses the left- and right-continuous versions, then $\Delta^{+}S(t)$ could in principle be based on the observation of $M(t)$ as well (as that is measurable), and not merely the left limit. Without welfare loss, it *will* actually not depend on $M(t)$ – that property is now proven, not merely assumed.

(d) The dynamics of M can depend on D, but not the other way around; if D depends on our control, then the m-dependent part may of course also do so. M and D may, however, be driven by common given processes – just augment D with these, and augment M with zero-valued coordinates to match the dimension for the dot product.

(e) Even though D is assumed a semimartingale, it is still a generalisation of the «no assumptions needed» Θ of Remark 2.2(c); recall that $D(t)$ does not correspond to $\Theta(t)$, but to $\int_0^t e^{-rs}\Theta(s)\ \mathrm{d}s$.

(f) Proposition 5.1 covers linear stochastic difference equations. To those who have only familiarised themselves with the stochastic integral with respect to Brownian motion and then maybe with respect to Lévy motions, Markov chains as Itô diffusion-type processes may look as a bit of an odd approach. However, the semimartingale concept does not require jump times to be random, and processes could very well be constant between integer times – all such processes are in fact semimartingales, as long as they are adapted.

The linearity of the evolutionary operator covers higher order differential equations for M. Assuming ad hoc stability and finite expectation, the optimal strategy still does not depend on M if the model of Proposition 2.1 is modified to allow for an nth-order linear differential equation

$$\sum_{i=0}^{n} \delta_i \cdot \left(\frac{\mathrm{d}}{\mathrm{d}t}\right)^i M(t) = \overline{\eta} \cdot 1_{t\in[0,\tau]} \qquad \text{with} \quad M \in \mathrm{C}^{n-1} \cap \mathrm{C}^n([0,\infty)\backslash\{\tau\}). \tag{29}$$

In the following example, we shall cover the stable non-oscillating case with $n = 2$.

Example 5.3: Consider the model of Proposition 2.1, except that M obeys Equation (29) with $n = 2$, $\delta_2 = 1$ and $\delta_1 > 2\sqrt{\delta_0}$, $\delta_0 \geq 0$. With characteristic roots

$$\lambda_1 = -\frac{\delta_1}{2} - \sqrt{\frac{\delta_1^2}{4} - \delta_0} \quad < \quad \lambda_2 = -\frac{\delta_1}{2} + \sqrt{\frac{\delta_1^2}{4} - \delta_0} \tag{30}$$

and initial data $M(0) = m$, $\dot{M}(0) = \mu$, we get

$$\begin{aligned} M(t) = \frac{\overline{\eta}}{\delta_0} + \frac{1}{\lambda_2 - \lambda_1}\Bigg(\left[\lambda_2\left(m - \frac{\overline{\eta}}{\delta_0}\right) - \mu\right] e^{\lambda_1 t} - \left[\lambda_1\left(m - \frac{\overline{\eta}}{\delta_0}\right) - \mu\right] e^{\lambda_2 t} \\ - \frac{\overline{\eta}}{\delta_0}\left\{\lambda_2\left(1 - e^{\lambda_1 \min\{0,t-\tau\}}\right) - \lambda_1\left(1 - e^{\lambda_2 \min\{0,t-\tau\}}\right)\right\}\Bigg). \end{aligned} \tag{31}$$

We see that the m and μ terms split out linearly in a way that does not depend on emissions. However, with $\overline{\eta}$ given, we can just as well split out the entire first line, which does not depend on τ. The first line of (31) yields the damage

$$\begin{aligned} &\frac{\theta\overline{\eta}}{\delta_0(r-\alpha)} + \frac{\theta}{\lambda_2 - \lambda_1}\left(\frac{\lambda_2\left(m - \frac{\overline{\eta}}{\delta_0}\right) - \mu}{r - \lambda_1 - \alpha} - \frac{\lambda_1\left(m - \frac{\overline{\eta}}{\delta_0}\right) - \mu}{r - \lambda_2 - \alpha}\right) \\ &= \left(\frac{\overline{\eta}}{(r-\alpha)} + (r - \lambda_1 - \lambda_2 - \alpha)m + \mu\right) \cdot \frac{\theta}{(r - \lambda_2 - \alpha)(r - \lambda_1 - \alpha)}, \end{aligned} \tag{32}$$

while the contribution from the rest, including the intervention cost, is

$$ke^{-r\tau} - \frac{\overline{\eta}}{\delta_0(\lambda_2 - \lambda_1)}\int_0^\infty e^{-rt}\Theta(t)\left\{\lambda_2\left(1 - e^{\lambda_1 \min\{0,t-\tau\}}\right) - \lambda_1\left(1 - e^{\lambda_2 \min\{0,t-\tau\}}\right)\right\}\mathrm{d}t$$

which has expected value

$$\mathbf{E}\left[ke^{-r\tau} - \frac{\overline{\eta}}{\delta_0}e^{-r\tau}\Theta(\tau)\int_0^\infty e^{-(r-\alpha)t}\left\{1 + \frac{\lambda_1 e^{\lambda_2 t} - \lambda_2 e^{\lambda_1 t}}{\lambda_2 - \lambda_1}\right\}\mathrm{d}t\right], \tag{33}$$

where we have used the strong Markov property to perform the time change, and the multiplicative form of the gBm. Again, by the strong Markov property and the continuity of the gBm, it suffices to consider stopping times of the form $\widehat{\tau} =$ first hitting time for $[\widehat{\theta}, \infty)$, and then one can optimise over $\widehat{\theta}$:

$$\begin{aligned}&\left(k - \widehat{\theta}\cdot\frac{\overline{\eta}}{\delta_0}\cdot\left\{\frac{1}{r-\alpha} - \frac{r-\lambda_1-\lambda_2-\alpha}{(r-\lambda_2-\alpha)(r-\lambda_1-\alpha)}\right\}\right)\cdot\mathbf{E}[e^{-r\widehat{\tau}}]\\ &\quad = \left(k - \frac{\overline{\eta}\widehat{\theta}}{(r-\alpha)(r-\lambda_1-\alpha)(r-\lambda_2-\alpha)}\right)\cdot\min\,\{1, (\theta/\widehat{\theta})^\gamma\}\end{aligned} \tag{34}$$

with γ given by (4). The minimiser is

$$\theta^* = \frac{\gamma}{\gamma-1}\cdot\frac{k(\overline{\eta})}{\overline{\eta}}\cdot(r-\alpha)(r-\lambda_1-\alpha)(r-\lambda_2-\alpha) \tag{35}$$

– compare to (4) again – so that (34) becomes $-\min\{1, (\theta/\theta^*)^\gamma\}\cdot k/(\gamma-1)$. To get the value function, and on a form comparable with (5), we add the (32) contribution:

$$\frac{(r-\lambda_1-\lambda_2-\alpha)m+\mu}{(r-\lambda_2-\alpha)(r-\lambda_1-\alpha)}\theta + k(\overline{\eta})\cdot\begin{cases}\left[\gamma\dfrac{\theta}{\theta^*} - \left(\dfrac{\theta}{\theta^*}\right)^\gamma\right]\Big/(\gamma-1) & \text{if} \quad \theta < \theta^* \\ 1 & \text{if} \quad \theta \geq \theta^*.\end{cases} \tag{36}$$

The first term is the damage which incurs with or without the project in question, and we see that the contribution from the optimised project has precisely the same form, except with a modified formula for the optimal trigger level θ^*. Observe that λ_1 and λ_2 are both negative, so the condition $r > \alpha$ ensures that everything converges.

5.2. *Some considerations beyond the Itô integral*

The previous subsection employed the standard stochastic calculus setup: the Itô integral with respect to semimartingales. The approach does, however, apply to other integral concepts as well.

Let us first point out that *fractional integrals* – whether they are of Erdélyi–Kober type (unifying and generalising both the Weyl and Riemann–Liouville types, see Pagnini 2012) or Hadamard type – are linear and can be covered by the form (9). Fractional differential equations have been proposed to model anomalous diffusion («diffusion» here meaning the physical phenomenon, e.g. particle flows in hydrology) (see, e.g. Chen et al. 2010). Another use is to allow for non-semimartingales, e.g. the well-known fractional Brownian motion, as a driving noise in differential equations. The brief exposition on fractional calculus in what follows is intended to facilitate the latter – the Stratonovich type and Hitsuda–Skorohod/Wick–Itô type integrals are valid also for the semimartingale framework as a special case.

An example: fractional Brownian motion. The *fractional Brownian motion (fBm)* $Z^{(h)}$ *of Hurst parameter* $h \in (0, 1)$ is a Gaussian process with zero mean, and covariance structure (for the univariate case) $\mathbf{E}[(Z^{(h)}(T) - Z^{(h)}(t))^2] = |T - t|^{2h}$. We shall work with the continuous-path version (ensured by the Kolmogorov continuity theorem). The term was coined by the seminal paper Mandelbrot and Van Ness (1968), defining it as the $(h - 1/2)$-order Weyl fractional integral (/derivative) of the ordinary Brownian motion Z, as, for $h \neq 1/2$,

$$\text{(constant)} \times \int_{-\infty}^{t} \left[(t-s)^{h-1/2} - (\max\{0, -s\})^{h-1/2}\right] \mathrm{d}Z(t), \tag{37}$$

but fBm also admits finite-memory Erdélyi–Kober representations (e.g. Dzhaparidze and van Zanten 2004). fBm has negatively correlated increments for $h < 1/2$. For $h > 1/2$ it has positively correlated increments, and the *long-memory* property that the covariance of the increments $Z^{(h)}(1) - Z^{(h)}(0)$ and $Z^{(h)}(T) - Z^{(h)}(1)$ diverges to $+\infty$ with T. The long memory has been a rationale to consider it as a model for various phenomena, including finance; however, not being a semimartingale, it leads to arbitrages (i.e. riskless free lunches) in frictionless markets with continuous trading. Among the vast literature on the topic, Rogers (1997) is an example where he not only establishes an arbitrage strategy, but also shows how to fit the same long-memory property into an arbitrage-free semimartingale model. fBm has also been used in the modelling of pollution (see, e.g. Guo et al. 2009).

The non-semimartingale property means that the fBm as an integrator behaves somewhat different from the ordinary Brownian motion. We mention a few cases suited to allow these kinds of processes as integrators.

The integrals of Young and Stratonovich and beyond. For stochastic analysis with respect to Brownian motion, there is the well-known Stratonovich integral, formalised by choosing the midpoint time for the integrands, taking limits of sums $Y\big(\frac{1}{2}(t_{i+1} + t_i)\big)\big(Z(t_{i+1}) - Z(t_i)\big)$. The Stratonovich integral admits an ordinary chain rule, without second-order terms. It turns out that if the driving process is continuous (discontinuities may be handled jump by jump) and with paths of zero quadratic variation, the Itô and Stratonovich integrals coincide, and equal the Young integral, which is, in some sense, the only continuous pathwise integral under this regularity. The Stratonovich integral thus extends the Young integral while keeping its (ordinary) chain rule. If we assume continuous sample paths and the ordinary chain rule in Proposition 5.1, we put $Y = 0$ and delete ΔZ and cross terms. Even though the optimisation problem could be cumbersome, we know that the state and history of M need not be taken into account, reducing the dimensionality of the optimisation problem.

Further generalisations can be given through the theory of *rough paths* (see, e.g. Lejay 2003), and the linear differential equations that arise with those integrals admit existence/uniqueness by Picard-type iteration.

The Wick–Itô and Hitsuda–Skorohod type integrals. Originating from white noise theory, these integrals are defined on distribution spaces, wherein Brownian motion is actually differentiable. The Wick–Itô formulation does in fact use the time-derivative of Brownian motion, allowing a Riemann-sum-based integral (technically defined in the

Bochner or even Pettis sense), written as

$$\int_0^T Y(t) \diamond \dot{Z}(t) \ \mathrm{d}t, \tag{38}$$

where the $\diamond$ denotes the (associative, commutative) so-called *Wick product*. This is not a pathwise («ω-wise») integral, as the Wick product is not a product between the *realisations* of random variables, but a product of probability *distributions* (somehow in the sense that the convolution product is); it has the property that the expectation of a Wick product is the product of expectations. Furthermore, the Wick–Itô integral admits a Wick product version of the ordinary chain rule, on a certain closure of the set of Wick polynomials $\sum c_i U^{\diamond i}$ – for example, the Wick exponential $\exp^{\diamond}\big(Z(t)\big) = 1 + \sum_{i\in\mathbb{N}} Z(t)^{\diamond i}/i!$ will have the time-derivative $\dot{Z}(t) \diamond \exp^{\diamond}\big(Z(t)\big)$. References for white noise theory with applications to Wick–Itô differential equations include the book of Holden et al. (1996), and for fBm, Elliott and van der Hoek (2003).

Let us simply perform the formal algebraic manipulation using the Wick-type integral, to see the consequences for the model (22) – assuming for simplicity continuous sample paths. First, we need a Wick-integrating factor: $\Pi = \exp^{\diamond}(\Xi)$. Then, M becomes

$$M(t) = \Pi(t) \diamond m + \Pi(t) \diamond \int_0^t \Pi(s)^{\diamond(-1)} \diamond \ \mathrm{d}Q(s) \tag{39}$$

$$= \exp^{\diamond}\big(\Xi(t)\big) \diamond m + \exp^{\diamond}\big(\Xi(t)\big) \diamond \int_0^t \exp^{\diamond}\big(-\Xi(s)\big) \diamond \dot{Q}(s) \ \mathrm{d}t. \tag{40}$$

Now the first term goes outside the minimisation. Let us assume that we are again in a one-shot model where $\dot{Q}(t)$ can be written as $\overline{\eta} 1_{[0,\tau]}$. If the damage functional is on Wick form – the differential being $\diamond \ \mathrm{d}D = \diamond\, \Theta \ \mathrm{d}t$ – then we are in a sense lucky, as we then have a pure Wick formulation, and one might apply expectations first and optimisation afterwards. However, mixing the ω-wise product and the Wick product will usually lead to intractabilities, and converting back and forth is certainly not trivial. For example, if we have pathwise differential $\mathrm{d}D$, we would want to calculate the probability distribution $\exp^{\diamond}\big(\Xi(t) - \Xi(s)\big) \diamond 1_{[0,\tau]}$ and evaluate at ω; the author is not aware of any tractable way to do this for general stopping times τ, and without doing this evaluation, we only have a distribution, not a response to path; without evaluation, the Wick product does not state how to respond to observations. Thus, the *modelling choice* at each «product» occurring in the model – Wick-type vs. ω-wise type – has non-trivial consequences to model behaviour. The linearity is still key to the property of Proposition 3.4 though, and the knowledge that the optimisation can be carried out without regard to X could potentially help making the problem tractable.

The Wick–Itô integral is often employed in *anticipative* stochastic calculus – for example, in (40), the $\diamond$ in front of m allows it to be random, and the theory even admits it to depend on future states of M. Indeed, there are other integral concepts designed for anticipative stochastic calculus (see, e.g. Kuo, Sae-Tang, and Szozda 2012 for a fairly recent one).

Stochastic partial differential equations. As pointed out in Remark 3.5, the theory of Proposition 3.4 covers (linear) time–space evolution modelled by the heat equation. The dissemination of the pollutant in space could also be subject to randomness. Such models could be hard to accommodate under ordinary stochastic calculus, as one could easily encounter models where one would want multiparameter Brownian motion and its second-order derivative. However, there is a well-developed theory based on white noise analysis, using the Wick–Itô approach – potentially leading to the same difficulties as for optimal stopping, in converting the model to one for response to actual observations. Again, see the book Holden et al. (1996).

6. Closing remarks

For linear models for the decay of pollution, as considered in this paper, the optimisation can be carried out without regard to the stock. This is a property often assumed as a valid approximation for idealised infinitesimal agents, but under linearity it holds exact regardless of size and also yields non-dependence with respect to others' exogenously given future emissions. The result admits general damage functionals as long as they do not depend on M explicitly and so that the functionals themselves are not controllable, and are linear or possess the appropriate additivity property.

In Itô stochastic differential equation models, there are possible model features we have not explicitly mentioned: the impact cost factor can even covariate with M in terms of Itô differentials, and the dynamics for M could depend on D or Θ; M still separates out as long as $\mathbf{\Xi}$ and $\mathbf{\Phi}$ are given functionals unaffected by our actions. Covariating Itô differentials seems natural from a small agent point of view, when the impact could in reality depend on the *aggregate* stock; then, upward fluctuations (or jumps) in the aggregate stock could cause Θ to increase, while it would still be a reasonable approximation to disregard a small agent's contribution to this effect. However, the reverse causality should definitely be allowed in a multiagent model, where there would be feedback from the Θ level to the agents' behaviour, and their respective «θ^*» triggers will vary over the agents' cost structures represented by k or more generally the $\mathbf{\Gamma}$ functional. Allowing for this, the «size no issue» feature must be expected to break down: if the other agents' behaviour can affect Θ, then our behaviour might as well, and reasonably the effect can only be disregarded if we are small.

For future research, the non-dependence result could make it easier to guess a solution form for problems with linear models, and then fit and verify by the tool of choice, e.g. dynamic programming. Furthermore, the reduction of dimensionality might be helpful for numerical solutions; for example, is the kind of transformations employed in this paper useful for more tractable optimisation in models which include stochastic (physical) diffusion of the pollutant? Finally, to what extent are the generalisations from linearity to the additivity conditions used in Section 4 useful? Could the non-dependence property, e.g. carry over to reasonable models for heterogeneous beliefs, as remarked at the end of that section? The limitations beyond the linear models will be a topic of interest, and only future research can tell whether this is of relevance for applications.

Acknowledgements

While carrying out this research, the author has been associated with the Oslo Centre for Research on Environmentally friendly Energy (CREE) and the Centre for the Study of Equality, Social Organization (ESOP) at the University of Oslo. CREE and ESOP are supported by The Research Council of Norway. This work was initiated, and part of the results were obtained, under contract with the Development Research Group, Environment and Energy Team at the World Bank, Washington, DC. The content solely reflects the author's view, and not those of any of his affiliations.

The author is indebted to an anonymous referee for suggesting improvements to readability, and to Bård Harstad for a fruitful discussion. Any errors are, of course, the author's sole responsibility.

References

Balikcioglu, M., P.L. Fackler, and R.S. Pindyck. 2011. "Solving Optimal Timing Problems in Environmental Economics." *Resource and Energy Economics* 33: 761–768.

Borodin, A.N., and P. Salminen. 1996. *Handbook of Brownian Motion – Facts and Formulae*. Basel: Birkhäuser Verlag.

Chen, W., H. Sun, X. Zhang, and D. Korošak. 2010. "Anomalous Diffusion Modeling by Fractal and Fractional Derivatives." *Computers and Mathematics with Applications* 59: 1754–1758.

Dzhaparidze, K., and H. van Zanten. 2004. "A Series Expansion of Fractional Brownian Motion." *Probability Theory and Related Fields* 130: 39–55.

Elliott, R.J., and J. van der Hoek. 2003. "A General Fractional White Noise Theory and Applications to Finance." *Mathematical Finance* 13: 301–330.

Föllmer, H., and A. Schied. 2002a. "Convex Measures of Risk and Trading Constraints." *Finance and Stochastics* 6: 429–447.

Föllmer, H., and A. Schied. 2002b. "Robust Preferences and Convex Measures of Risk." In *Advances in Finance and Stochastics*, edited by K. Sandmann, 39–56. Berlin: Springer.

Framstad, N.C. 2011. "A Remark on R.S. Pindyck: 'Irreversibilities and the Timing of Environmental Policy'" *Resource and Energy Economics* 33: 756–760.

Framstad, N.C., and J. Strand. 2013. *Energy Intensive Infrastructure Investments with Retrofits in Continuous Time: Effects of Uncertainty on Energy Use and Carbon Emissions* (Memorandum No. 11/2013). Oslo: Department of Economics, University of Oslo.

Frittelli, M., and E.R. Gianin. 2002. "Putting Order in Risk Measures." *Journal of Banking & Finance* 26: 1473–1486.

Guo, W., Y. Wang, M. Xie, and Y. Cui. 2009. "Modeling Oil Spill Trajectory in Coastal Waters Based on Fractional Brownian Motion." *Marine Pollution Bulletin* 58: 1339–1346.

Harstad, B., and M. Battaglini. 2012. *Participation and Duration of Environmental Agreements*. Cambridge, MA: The National Bureau of Economic Research.

Holden, H., B. Øksendal, J. Ubøe, and T. Zhang. 1996. *Stochastic Partial Differential Equations*. Boston, MA: Birkhäuser Boston.

Komiya, H. 1988. "Elementary Proof for Sion's Minimax Theorem." *Kodai Mathematical Journal* 11: 5–7.

Kuo, H.-H., A. Sae-Tang, and B. Szozda. 2012. "A Stochastic Integral for Adapted and Instantly Independent Stochastic Processes." In *Stochastic Processes, Finance and Control*, edited by S.N. Cohen, D. Madan, T.K. Siu, and H. Yang, 53–71. Hackensack, NJ: World Scientific.

Lejay, A. 2003. "An Introduction to Rough Paths." In *Séminaire de Probabilités XXXVII*, 1–59. Berlin: Springer.

Mandelbrot, B.B., and J.W. Van Ness. 1968. "Fractional Brownian Motions, Fractional Noises and Applications." *SIAM Review* 10: 422–437.

Pagnini, G. 2012. "Erdélyi-Kober Fractional Diffusion." *Fractional Calculus and Applied Analysis* 15: 117–127.

Pindyck, R.S. 2000. "Irreversibilities and the Timing of Environmental Policy." *Resource and Energy Economics* 22: 233–259.

Pindyck, R.S. 2002. "Optimal Timing Problems in Environmental Economics." *Journal of Economic Dynamics and Control* 26: 1677–1697 (in honour of David Kendrick).

Rogers, L.C.G. 1997. "Arbitrage with Fractional Brownian Motion." *Mathematical Finance* 7: 95–105.