The Precautionary Principle when Project Implementation Capacity is Congestible*

Anthony Heyes
University of Ottawa
University of Birmingham

Sandeep Kapur†
Birkbeck College
University of London

November 2022

Abstract

The precautionary principle justifies postponing the implementation of development projects to await better information about their environmental impacts. But if implementation capacity is congestible, as is often the case in practical settings, a postponed project may have to vie for implementation priority with projects that arrive later. Limitations of implementation capacity create two risks. First, it may never make sense to go back to a postponed project, even if it is later revealed to be a good one. Second, the planner may find it worthwhile to go back to it, but at the expense of undesirable delay of future projects. We update Arrow and Fisher (1974) with a planner facing a sequence of projects varying stochastically in their (1) importance and (2) improvability, but knowing that implementation capacity is congestible. The scope for congestion implies a 'bonus' for earlier-than-otherwise decisions, in common parlance "keeping the desk clear" that works against the well-understood option value that encourages delay. The optimal decision rule depends upon the stochastic environment whereby future projects are generated, in ways that are not obvious. The value of the bonus is increasing in the expected importance of future projects but decreasing in their expected improvability. Higher variability of the importance of projects, in the sense of mean-preserving spread, increases the size of the bonus, but variability in their improvability has a generally ambiguous impact. We characterize the adjusted decision rule and note its implications for the conduct of cost-benefit informed policy.

Keywords: Dynamic decision-making, the Precautionary Principle, project appraisal, bottlenecks

JEL: D61, D81, H43

---

*We thank two anonymous reviewers and the coordinating editor for their comments and suggestions.

†Corresponding author: s.kapur@bbk.ac.uk, Department of Economics, Birkbeck, University of London, Malet St, London WC1E 7HX.

1 Introduction

Suppose that a project, once implemented, is hard to reverse, and the benefits of non-implementation are uncertain. Then a rational planner who anticipates the arrival of improved information on those benefits may prefer to postpone making a decision on the project, even if its expected net present value is known to be positive. That insight was formalized in a seminal paper by Arrow and Fisher (1974), the objective of which was to characterize how irreversibility should be taken account of in cost-benefit analysis of projects that entail damage to an environmental asset of uncertain value, and provides the intellectual basis for the so-called 'precautionary principle' frequently invoked in policy discussion.[1] The bias that a rational decision-maker should exhibit against project implementation in such circumstances is operationalized by the metrics of option and quasi-option value (Crabbe (1987)). More informally, Arrow and Fisher rationalize the logic commonly expressed by decision-makers, committees and other organizations in the wake of *in*-action, that they will "look to return to this later, when we know more".

Missing from this analysis, however, are the implementation constraints, bottlenecks and limitations on the scaling-up and down of activity that are important features of many real-world settings. If the capacity to implement is congestible – in other words there is some limit on the number of projects that an organization can execute (or execute well) at one time – then postponement has a cost unaccounted for in models involving a *single* potential project. The existence of an implementation constraint means that the decision-maker faces a more complex problem as the constraint makes inter-dependent decisions on proposals that could otherwise be treated as separate. Backlogs in implementation can prove costly, as new projects have to 'compete' for priority with un-executed projects carried over from earlier. This provides two paths to inefficiency. Depending on what comes along, (a) the planner might never find a time when it makes sense to go back to a previous project, even if its net benefits are later revealed to be positive; or (b) she may opt to come back to it in a future period, but at the expense of diverting attention from some subsequent project which, consid-

[1]See Atkinson et al (2006). In Arrow and Fisher (1974) and most of the subsequent literature choice-relevant information is assumed to arrive with passage of time. An alternative strand of research treats decision-makers as active gatherers of information. See, for example, Che and Mierendorff (2018), and references therein.

ered in isolation, would have merited prompt execution. A forward-looking decision-maker, in deciding to postpone acting on a proposal, will recognize that such vacillation may impact what gets done later.

Real contexts (policy, corporate, or organizational) are rarely characterized by a decision being required on a single, stand-alone project. More typically the planner can expect to face a series of proposals that arrive over time, and the scope for congestion in decision-making is usually well-understood in these settings. Activity levels often cannot be scaled up and down from period to period to accommodate 'lumpy' decision flows without loss of performance. Presbitero (2016), for example, studies a large set of World Bank projects between 1970 and 2017 and finds capacity limitations to be a significant hindrance to project success when multiple projects are executed simultaneously.[2] The could reflect the organizations' own capabilities, or the absorptive capacity of the environment. The intent of this paper is to investigate how the precautionary principle and concept of quasi-option value need to be adjusted for such settings. The analysis is primarily motivated by the stylized realities of project-focused organizations such as international development agencies, municipal development corporations and project-based NGOs. But the logic might equally apply to private firms or other entities including universities, families or individuals. Any setting where there is a flow of potential projects to engage in (some new, some ideas carried over from earlier), but where the entity in question can "only do so much at one time".[3]

More concretely, we develop a stylized model of a capacity-constrained planner facing a sequence of project proposals. Projects vary in their **importance** and their **improvability**. We will define these precisely, but in essence these relate to the size of the project and to the extent to which the planner might be able to make a better decision by postponing it. The decision-maker knows the characteristics of the proposal currently in front of her while the characteristics of future proposals are uncertain, though

[2]Readers may recognise a parallel phenomenon at individual level – 'mental bandwidth' is limited (Mullainathan and Shafir, 2013) such that an individual can only do so many things effectively at one time. Among management scholars the notion of organizational bandwidth being congestible is widely acknowledged (see for, examples, Nunamaker et al. (2001) and the associated journal special issue *Enhancing Organizations' Intellectual Bandwidth.*). The congestion may be driven by a number of factors, but by way of caricature, "... the chief executive can only do so many things at once" (Geanakoplos and Milgrom, 1991).

[3]The motive here is distinct from other models of dynamic allocation of project effort under constraints, for examples Gifford and Wilson (1995), Grossman and Shapiro (1986).

the distributions from which they are drawn is known. For tractability we assume that proposals arrive one per period and have a 'shelf life' – if not acted upon within two periods of arrival they expire. We characterize the solution to the planner's problem in this setting to compare outcomes when the planner faces the implementation constraint against the unconstrained benchmark case embedded in Arrow and Fisher's original formulation.

The central tension at the heart of the decision problem is the trade-off between wanting to delay the decision on improvable projects (especially important ones), analogous to waiting for more information in the model of Arrow and Fisher, but equally the desire to prevent backlog of projects – to "keep a clear desk". It turns out that making that trade-off optimally generates some nuanced comparative statics. We explore in particular the comparative statics of optimal decision-making with respect to the decision environment in which the planner finds himself, as parameterized by the mean and variation of the distribution from which the improvability and importance characteristics of future projects are drawn. The premium attached to keeping a clear desk is increasing in the expected importance of future projects, but decreasing in the variability of their importance. Equally, it is decreasing in the mean of how improvable future projects are expected to be, but may be increasing or decreasing in the dispersion of that improvability.

1.1 A Single Project Model

The essence of the precautionary principle can be understood in a simple two-period example.

There is a proposal for an intervention that would deliver a flow of benefits in each of two periods, an 'initial period' and a 'future period'. Following the original framing provided by Arrow and Fisher we will focus on the decision to develop or not develop a parcel of land, but could equally well refer to the implementation of any sort of policy intervention whose future environmental costs will be better understood with the passage of time. If implemented at the start of the initial period the benefits that accrue are $\tau > 0$ in the initial period and $\eta > 0$ in the future period, making a total of $(\tau + \eta)$.[4] If implementation is delayed until the start of the future period,

[4] We ignore discounting everywhere in the model. As in Arrow and Fisher (1974), none of the arguments made here rely on inter-temporal discounting.

the only benefit received is η, so that τ is a measure of the cost of delay.

Implementation is not reversible. Development of the parcel of land implies the irreversible destruction of some natural resource. The value of that resource, Ω, is unknown at the start of the initial period, though it is common knowledge that it takes the value $\omega > 0$ with probability $\pi \in (0, 1)$ and the value $\omega_0 < \omega$ with probability $(1 - \pi)$. To save notation, we set $\omega_0 = 0$.[5] In other words, with probability π the destroyed resource will turn out to be high value, and worthless otherwise. The uncertainty is resolved at the start of the future period.[6]

To rule out uninteresting cases in which there are no circumstances in which implementation makes sense, or no case in which non-intervention makes sense, we assume the following:

Assumption 1 $\omega > \tau + \eta$.

In the absence of uncertainty (i.e., if the realization of Ω was known in advance) the planner implements the project if and only if the resource is of low (zero) value. There is no gain to delaying implementation.

When Ω is uncertain we can think of two scenarios. In one scenario the decision maker is compelled to make the implement/not implement choice at the start of initial period. In the other in which he can defer until the start of the future period.

If compelled to decide at the start of the initial period a risk-neutral decision-maker compares the expected net benefits of implementation, $(\tau + \eta - \pi\omega)$, with the net benefits of not implementing, which are 0. The optimal decision rule is simple: taking values of other parameters as fixed, it is optimal to implement immediately as long as there is sufficiently small probability that the destroyed resource will turn out to be valuable. More precisely,

Result 1 *If compelled to implement or not implement at the start of the initial period, a risk-neutral decision maker implements if and only if* $\pi \leq \pi^*$, *where*

$$\pi^* = \left(\frac{\tau + \eta}{\omega}\right).$$

[5] Our findings continue to hold if we allow ω_0 to be positive, as long as $\omega_0 < \eta$.

[6] We do not require that uncertainty is totally resolved by any date, though that is what we will assume. It would be sufficient to regard ω and ω_0 as the *conditional* expected values of the natural resource, contingent on the arrival of some binary signal.

Contrast this with circumstances in which the decision-maker is able to defer the decision to the future. The attraction of deferral is that is enables a better-informed choice. If the environmental asset is revealed to be high value ω, he can preserve it by eschewing development. On the other hand if the asset is revealed to be low value, deferred development produces benefit η. Deferral has expected payoff $\pi\omega + (1-\pi)\eta$. Immediate development has benefits $\tau + \eta$. Comparing these two, we can define the expected net benefit to early implementation as

$$\Delta_0(\pi) = \tau + \eta - [\pi\omega + (1-\pi)\eta]. \tag{1}$$

For a risk-neutral planner, early implementation is warranted in this circumstance if $\Delta_0(\pi) \geq 0$. Again, this optimal decision criterion can be expressed in terms of a critical value of π:

Result 2 *If allowed to defer a decision until the start of the future period, a risk-neutral decision-maker implements at the start of the initial period if and only if $\pi \leq \hat{\pi}$ where*

$$\hat{\pi} = \left(\frac{\tau}{\omega - \eta}\right). \tag{2}$$

Observe that Assumption 1 ensures that $\hat{\pi} < \pi^*$, so that these two distinct hurdle rates partition the unit interval into three sub-intervals. For values of $\pi \in (0, \hat{\pi}]$, the likelihood of the vulnerable resource turning out to valuable is small enough that early implementation is warranted regardless. Likewise, for values of $\pi \in (\pi^*, 1)$, the high likelihood of destroying a valuable resource warrants no development in the initial period.

The intermediate range is more interesting. For $\pi \in (\hat{\pi}, \pi^*]$, if compelled to choose at the start of the initial period the planner opts to implement (development rather than conservation). However if given the chance he would prefer to postpone the decision, forgoing the short-run economic gain to wait to learn about the true value of the natural resource at risk of destruction. In policy parlance the planner invokes the precautionary principle – eschews a project (at least in the short-run) that has positive expected net present value out of caution for the potential 'worst case' environmental damage that might follow. The precautionary principle only biases the decision against irreversible development, it does not preclude it in every scenario.

A decision environment that permits deferral can be re-cast in slightly different terms, by defining the **improvability** of the decision through delay in implementation. Define $\alpha(\pi)$ as the ratio of the expected payoff to a deferred decision relative to the payoff from immediate implementation: we have

$$\alpha(\pi) = \frac{\pi\omega + (1-\pi)\eta}{(\tau+\eta)} = \frac{\eta + \pi(\omega-\eta)}{(\tau+\eta)}.$$

Clearly it is preferable, under current assumptions, to defer any decision to implement whenever $\alpha(\pi) > 1$. The improvability depends on the parameters in a natural way: other things equal, improvability is larger when τ is small relative to η, or when ω is large. In what follows, we take the payoff parameters (τ, η and ω) as given, with π as the variable of interest. The improvability of the decision, and hence the case for its postponement, is increasing in π.

2 A Sequence of Projects with Limited Implementation Capacity

We extend the simple setting to allow for two plausible features: (1) The planner does not face a single, once-and-for-all decision, but rather a flow of project proposals on which he must make decisions over time. (2) There are limits on implementation capacity within the organization or setting in which he is operating.

On (1), while sometimes a decision-maker may be appointed to examine one and only one decision in isolation, much more typical is the situation in which we have an individual (like a manager) or other decision-making entity, such as a committee, tasked with arbitrating on a flow of decisions that arise sequentially over time. Feature (2) simply recognizes organizational capacity (bandwidth) as finite.[7]

To operationalize these features we first extend the model above to a three period setting, with $t = 1, 2, 3$. Specifically, a candidate project arises in period 1 which can be implemented at the start of that period, or the decision postponed to the start of period 2. A second, independent project

[7] A softer version of (2) would be to make implementation not subject to an absolute constraint but rather congestible. In other words, an increase in the number of projects 'on the go' at any one time would reduce the efficacy of implementation – less than perfect scalability in implementation activities.

arises at the start of period 2, which can be implemented at the start of that period or postponed to the start of period 3. However implementation capacity is limited: more concretely only one project can be implemented in any period.[8]

Projects vary in two characteristics.

The first, as above, is improvability $\alpha(\pi_t)$ of the project through deferral, where the subscript t has been added to denote the period t in which candidate project is 'born'. At the start of any period t, the value of π_t and, thereby, the improvability $\alpha(\pi_t)$ of project t is revealed. Thus, while the improvability of Project 1 is known at start of period 1, π_2 (and the associated improvability of Project 2) is not known till the start of period 2. Ex ante only the probability distribution from which random $\tilde{\pi}_2$ is drawn is known.

The second dimension along which projects vary is their scale or importance. Recall that the project described in the previous section delivers payoff of $\tau + \eta$ if implemented immediately, and the greater of ω or η if postponed. To reduce the model to bare essentials we assume that projects that arrive over time are linearly-scaled versions of this base project: a project of scale or importance s entails payoffs $(\tau + \eta)s$ if implemented immediately, the greater of ωs and ηs if postponed.[9] The 'importance' of project t is given by s_t. Without loss of generality we normalize $s_1 = 1$. To capture, in the simplest manner, the possibility that the second project in the sequence may turn out to be larger or smaller (that is, more or less important) than the first, we assume $\tilde{s}_2 \in \{s_h, s_\ell\}$ where $s_h > 1 > s_\ell$.

The sequence of decision problems is completely specified by these variables: $(\tau, \eta, \omega, 1, \tilde{s}_2, \pi_1, \tilde{\pi}_2)$, along with probability distributions on $\tilde{s}_2$ and $\tilde{\pi}_2$. The two random variables $\tilde{s}_2$ and $\tilde{\pi}_2$ have known distributions but the decision-maker observes their realized values only at the start of period 2, at the birth of Project 2. Here, then, our decision-maker knows that a further project proposal will appear next period but does not know ex ante how important and how improvable it will be. In order to obtain closed-form

[8]The attentive reader will also note that this formulation implies that project opportunities expire or have a 'shelf-life' of two periods – the first project cannot be implemented in period 3. This is for tractability, as it ensures no more than one unimplemented project from a past round can be carried over.

[9]A richer setting could allow the structure of returns to sequential projects to vary more generally. Our simplified structure allows us to focus on the impact of the likely size of future projects without distorting qualitative insights.

solutions, we restrict the probability distributions as follows:

Assumption 2 *The stochastic characteristics of Project 2 are given by the following probability distributions.*

(a) $\tilde{\pi}_2 \sim U[0,1]$*: the realization of* $\tilde{\pi}_2$ *is distributed uniformly in the unit interval;*

(b) $\mathrm{pr}(s_h) = \mathrm{pr}(s_\ell) = 0.5$.

Absent constraints on implementation capacity, the sequential projects are completely separable and can be progressed (or not) independently of each other. Our interest is in the inter-dependence that is generated by limited implementation capacity, and how that affects how projects should optimally be evaluated, and the application of the precautionary principle.

To assess the impact of future projects on the optimal timing of current projects in this settings, we evaluate the consequences of early versus late implementation of the initial project that arrives in period 1.

2.1 Early implementation of Project 1

We begin by characterizing the expected payoff if the decision-maker chooses to implement Project 1 immediately.

Project 1 then delivers payoff $\tau + \eta$, and the decision on Project 2 is unencumbered by limitations of implementation capacity. The decision on Project 2 is then identical to one described in the previous section, with the threshold for its early implementation given by $\hat{\pi}$ in equation (2).[10] The optimal choice for Project 2, and its payoff, will depend on the realization of $\tilde{\pi}_2$ at the start of period 2. For realizations $\tilde{\pi}_2 \leq \hat{\pi}$, the Project 2 turns out to be not sufficiently improvable to merit its postponement: if so, its immediate implementation delivers $(\tau + \eta)\tilde{s}_2$. For $\tilde{\pi}_2 > \hat{\pi}$ it is optimal to postpone the decision to period 3, with expected payoff $[\tilde{\pi}_2\omega + (1 - \tilde{\pi}_2)\eta]\tilde{s}_2$. Summarizing, the payoff to the Project 2 conditional on $\tilde{s}_2$ and $\tilde{\pi}_2$ is

$$v_2(\tilde{s}_2, \tilde{\pi}_2) = \begin{cases} (\tau + \eta)\tilde{s}_2 & \text{if } \tilde{\pi}_2 \leq \hat{\pi} \\ [\tilde{\pi}_2\omega + (1 - \tilde{\pi}_2)\eta]\tilde{s}_2 & \text{otherwise.} \end{cases}$$

[10] As all payoffs are scaled by a common multiple s, the critical threshold $\hat{\pi} \equiv \tau/(\omega - \eta)$ is invariant to s.

By Assumption 2, $\tilde{\pi}_2$ is distributed uniformly over the unit interval $[0, 1]$. Taking expectation over $\tilde{\pi}_2$, the ex ante expected payoff to Project 2 of size $\tilde{s}_2$ is

$$\begin{aligned} V_2(\tilde{s}_2) &= \int_0^{\hat{\pi}} (\tau + \eta)\tilde{s}_2 d\tilde{\pi}_2 + \int_{\hat{\pi}}^1 [\eta + (\omega - \eta)\tilde{\pi}]\tilde{s}_2 d\tilde{\pi}_2 \\ &= \frac{1}{2}\left[(\omega + \eta) + \frac{\tau^2}{\omega - \eta}\right]\tilde{s}_2. \end{aligned}$$

As $\tilde{s}_2$ is assumed to take values s_h and s_ℓ with equal probability, we define $\overline{s} = \frac{1}{2}(s_h + s_\ell)$. Taking expectation over $\tilde{s}_2$, the payoff to unencumbered optimal choice on the Project 2 is

$$V_2 = \frac{1}{2}\left[(\omega + \eta) + \frac{\tau^2}{\omega - \eta}\right]\overline{s}. \tag{3}$$

We can now aggregate the total expected payoff to a policy that entails early implementation of Project 1, followed by the optimally-timed decision on Project 2:

$$EV^{\text{early}} = (\tau + \eta) + \frac{1}{2}\left[(\omega + \eta) + \frac{\tau^2}{\omega - \eta}\right]\overline{s}. \tag{4}$$

2.2 Deferring Project 1

Next we evaluate the expected payoff for a policy in which the implementation of Project 1 is deferred. While such a deferral allows for updating with respect to new information with on Project 1, the implementation constraint may have an unhelpful knock on effect on Project 2. The lingering Project 1 might induce a welfare-reducing delay in Project 2. Alternatively, the revealed characteristics of Project 2 may be such that it never makes sense to go back to Project 1.

To evaluate the outcomes, note that if Project 1 is deferred, the decision-maker faces one of two possibilities.

Case 1: New information reveals Project 1 to be high damage and therefore unattractive even if considered in isolation.

In period 2 the decision maker learns, with probability π_1, that the resource damaged by implementation of Project 1 is high value, ω. In this case Project 1 is welfare-reducing and is discarded. With no previously-deferred candidate project that merits implementation, the decision-maker is unrestricted with respect to Project 2.[11] Given that the optimal implementation of Project 2 is unencumbered, its expected payoff is given, as before, by equation (3).

To summarize, in the event that Project 1 is abandoned due to its revealed high environmental cost, the decision-maker retains the value ω of the preserved resource, and the expected payoff to unencumbered choice for Project 2. Aggregating those payoffs for this scenario:

$$EV^{\text{defer}}(\omega) = \omega + \frac{1}{2}\left[(\omega + \eta) + \frac{\tau^2}{\omega - \eta}\right]\bar{s}. \tag{5}$$

Case 2: New information reveals Project 1 to be low damage and therefore attractive if considered in isolation.

With probability $(1 - \pi_1)$ the decision-maker learns, at the start of period 2, that the environmental resource at risk from the implementation of Project 1 is low value (in fact, worthless, by our assumption). Considered in isolation the planner would implement Project 1 at this stage. However, the capacity constraint with regard to implementation – operationalized here by the assumption that only one project can be implemented at any one time – means that going ahead with Project 1 necessarily means not proceeding with Project 2, at least for now.

It is convenient to partition this case into two sub-cases based on the realized scale or importance, s_ℓ or s_h, of Project 2. Without loss of generality assume that $s_\ell < \frac{\eta}{\tau} < s_h$.[12]

First consider $s_2 = s_\ell$. In this case, Project 2 turns out to be small enough that the optimal policy will involve 'serial postponement' in the implementation of projects: to implement the legacy Project 1 in period 2, and implementing Project 2 in period 3 only if environmental costs are revealed to be low.[13] The ex-ante expected payoff to the optimal decision

[11] Our modeling assumption is that abandoning a legacy project immediately releases implementation capacity for current projects. In real settings, even abandonment could demand decision-making resources.

[12] To see that this does not imply loss of generality observed that we could allow $\tilde{s}_2$ to take one of any multiple values, and then partition the set of projects into two sub-sets: those that are small vs those that are large, with $\frac{\eta}{\tau}$ being the dividing line. In effect, s_ℓ and s_h can be regarded as expected values conditional on that partition.

[13] To see why 'serial postponement' is optimal in this case, note that it delivers a payoff of η from implementing Project 1 in period 2 and *at least* ηs_ℓ from an optimal decision in period 3 for Project 2. Implementing Project 2 immediately, with payoff $(\tau + \eta)s_\ell$ would require abandoning Project 1 altogether. Given $s_\ell < (\eta/\tau)$ the total payoff from sequential postponement is higher.

on Project 2 in this case is

$$V_2^{\text{defer}}(\omega_0; s_\ell) = \int_0^1 [\pi_2\omega + (1-\pi_2)\eta] d\pi_2\, s_\ell = \frac{1}{2}(\omega + \eta)s_\ell. \tag{6}$$

To summarize, in this scenario Project 1 is initially deferred but then implemented in period 2 after it is revealed it to viable, with payoff η. Project 2, of size s_ℓ, is deferred to period 3 and implemented only if viable. Aggregating the expected payoffs across the two projects in this scenario, we have

$$EV^{\text{defer}}(\omega_0; s_\ell) = \eta + \frac{1}{2}(\omega + \eta)s_\ell. \tag{7}$$

Finally, consider $s_2 = s_h$. The planner now has on his desk the deferred Project 1, which the passage of time has revealed to be an attractive one, at least when considered on its own merits but also a newly arrived Project 2 that is relatively important (larger in scale).

This final scenario poses a more complex problem and requires comparison of the period 2 return to implementation of Project 1, η, with any downside from the forced deferral of Project 2. The magnitude of that downside (and indeed whether there is any such downside at all) depends not just on the scale of the Project 2 but also its improvability.

To evaluate this trade-off, note that if the decision maker opts to implement Project 1 in period 2, then goes on to behave optimally with respect to Project 2 in the subsequent period, he anticipates payoff

$$\eta + [\pi_2\omega + (1-\pi_2)\eta]s_h.$$

If instead he abandons Project 1 in order to implement Project 2 immediately he obtains $(\tau + \eta)s_h$. The optimal selection between these two courses of action naturally depends on π_2. The value of π_2, recall, determines the improvability of the decision on Project 2 – how much 'better' that decision can be made by waiting for the environmental impacts of the project to be revealed. As π_2 becomes larger the deferral of a decision on Project 2 becomes less costly and, beyond some point even desirable in its own right.

Manipulation shows that, contingent on arriving at the start of period 2 with a viable legacy Project 1, discarding Project 1 is the optimal decision for and only for realizations of π_2 that are low enough. For higher realizations of π_2 the optimal strategy involves serial postponement. Summarizing, the total payoffs across projects in these two scenario gives:

$$v(\pi_2|\omega_0, s_h) = \begin{cases} (\tau+\eta)s_h & \text{if } \pi_2 \le \frac{\tau s_h - \eta}{(\omega-\eta)s_h} \\ \eta + [\pi_2\omega + (1-\pi_2)\eta]s_h & \text{otherwise.} \end{cases} \quad (8)$$

Taking expectation over possible realizations of π_2 gives

$$EV^{\text{defer}}(\omega_0, s_h) = \eta + \frac{1}{2}\left[(\omega+\eta)s_h + \frac{(\tau s_h - \eta)^2}{(\omega-\eta)s_h}\right]. \quad (9)$$

We can now evaluate the ex ante return to a policy that defers a decision on Project 1. Weighting expressions (5), (7), and (9) by their probabilities, the expected value of deferral of Project 1 in the initial period is

$$EV^{\text{defer}} = [\pi_1\omega + (1-\pi_1)\eta] + \frac{1}{2}(\omega+\eta)\bar{s} + \frac{1}{2}\frac{\bar{s}}{\omega-\eta}\left[\pi_1\tau^2 + (1-\pi_1)\frac{(\tau s_h - \eta)^2}{2\bar{s}s_h}\right] \quad (10)$$

2.3 The incentive for early implementation of Project 1

Finally we assess the incentive to defer a decision on Project 1 in the initial period by comparing the payoff to its early implementation, as obtained in equation (4), with that to its deferral, as in equation (10). The difference between these two is usefully denoted as

$$\Delta_c(\pi_1) = EV^{\text{early}} - EV^{\text{defer}}. \quad (11)$$

The value $\Delta_c(\pi_1)$ captures the net advantage to prompt implementation of Project 1, in a setting with a sequence of project and in which the decision-maker recognizes that his organization's implementation capacity is limited. In this constrained setting early implementation of Project 1 is warranted if and only if $\Delta_c(\pi_1) > 0$.

3 The impact of congestible implementation capacity

How does congestible implementation capacity affect the optimal implementation of projects?

The simple setting in Section 1 analyzed the case without any constraints in implementation capacity. Following equation (1), with no constraints, early implementation of Project 1 is warranted if and only if $\Delta_o(\pi_1) \geq 0$. The analysis in Section 2 arrived at an analogous criterion in the presence of congestible implementation capacity: namely that Project 1 should be implement promptly if and only if $\Delta_c(\pi_1) \geq 0$.

A comparison of $\Delta_o(\pi_1)$ and $\Delta_c(\pi_1)$ allows us to judge the impact of the constraint in implementation capacity. Substituting from equations (4) and (10) in equation (11), and comparing with equation (1), we can write

$$\Delta_c(\pi_1) = \Delta_o(\pi_1) + \delta_c(\pi_1), \tag{12}$$

where, the second term

$$\delta_c(\pi_1) \equiv \frac{1}{2}\frac{(1-\pi_1)}{\omega - \eta}\left[\tau^2\overline{s} - \frac{(\tau s_h - \eta)^2}{2s_h}\right] \tag{13}$$

quantifies the adjustment, due to limited implementation capacity, on the net benefit to early implementation.

Equation (12) decomposes the net advantage to early implementation of Project 1 into two components: $\Delta_0(\pi_1)$ captures the net benefit to early implementation of the immediate project at hand, while $\delta_c(\pi_1)$ is a measure of the premium attached to 'keeping the desk clear' to tackle future projects that might call for immediate implementation.

Proposition 1 *Congestible implementation capacity creates a bias to earlier-than-otherwise-optimal implementation of projects. The size of the bias is captured by the term $\delta_c(\pi_1)$, which is positive, so works against the precautionary principle.*

Proof. It is sufficient to check that $\delta_c > 0$. With straightforward manipulation,

$$\begin{aligned}
\delta_c(\pi_1) &= \frac{1}{4}\frac{(1-\pi_1)}{\omega-\eta}\left[(s_h + s_\ell)\tau^2 - s_h\left(\tau - \tfrac{\eta}{s_h}\right)^2\right] \\
&= \frac{1}{4}\frac{(1-\pi_1)}{\omega-\eta}\left[s_h\left(\tau^2 - \left(\tau - \tfrac{\eta}{s_h}\right)^2\right) + s_\ell\tau^2\right].
\end{aligned}$$

Recall that $\pi_1 \in (0,1)$, so $\delta_c(\pi_1)$ is strictly positive, which implies $\Delta_c(\pi_1) > \Delta_o(\pi_1)$. If so, for any π_1 there is bias towards early decisions. $\square$

The result is intuitive and central to the paper. Limitations in implementation capacity create the possibility of congestion in future decisions. The congestion may manifest itself in the potentially costly deferral of future decisions. In other circumstances, where future projects turn out to be large and not worth deferring, they might trigger the abandonment of legacy projects that have not been yet implemented. Both considerations make a case for earlier implementation of the project in hand, with the purpose of releasing implementation capacity for future.

Importantly, the consideration of limited implementation capacity only biases decision making towards pre-emptive implementation. For values of π_1 large enough, $\Delta_0(\pi_1)$ is sufficiently negative and $\delta_c(\pi_1)$ is small in magnitude, so even $\Delta_c(\pi_1)$ is negative: here the optimal policy would be to defer implementation regardless. In words, when there is a high probability that future information will reveal that implementing a project will result the loss of a valuable environmental resource, the precautionary principle will trump any apprehension about congestible implementation capacity. Our insight here is that concerns about accumulating costly backlogs of unimplemented projects push against the bias towards delay implied by the textbook precautionary principle.

3.1 Comparative statics

While it is intuitive that the possibility of congestion of implementation might create a premium (or what we will call 'bonus') for early execution – an onus on 'keeping a clear desk' – our model allows us to be more specific in quantifying that bias, and identifying what characteristics of a decision environment determine its size.

To this second end we turn to a number of comparative static exercises. First with respect to the parametric characteristics of Project 1, then with respect to the stochastic processes that generate the characteristics of future projects.

Result 3 *With congestible implementation capacity the bonus for immediate implementation of Project 1 is decreasing in π_1. That is, $\Delta_c(\pi_1)$ is decreasing in π_1.*

Proof. This follows from inspection of equations (1), (12) and (13), which establish that both $\Delta_o(\pi_1)$ and $\delta_c(\pi_1)$ are decreasing in π_1, and hence so is their sum $\Delta_c(\pi_1)$. $\square$

The two channels underlying this result are worth spelling out. First, $\Delta_o(\pi_1)$ is decreasing in π_1 because higher values for π_1 imply greater improvability of the decision on Project 1, as postponement allows better adaptation to valuable information. The second channel, which operates through $\delta_c(\pi_1)$, is more subtle: a higher value of π_1 implies a greater probability that a high revealed value of the environmental asset at risk from Project 1 will lead to the abandonment of that project in the future. In that scenario, there is no effective limitation on implementation capacity for future projects. Indeed, as π_1 approaches 1, $\delta_c(\pi_1)$ tends to zero, so the decision rules with and without the constraints come to coincide.

Next we turn to the question of how the desire to keep a clear desk depends on the payoff parameters, namely τ, η that capture the payoff structure of the projects, and on ω, the value of the underlying environmental asset at risk. For the next and subsequent results, it is helpful to re-write (13) slightly,

$$\delta_c = \frac{1}{4}\frac{(1-\pi_1)}{\omega-\eta}\left[\tau^2 s_\ell - \eta^2\frac{1}{s_h} + 2\tau\eta\right]. \tag{14}$$

Result 4 *With congestible implementation capacity the bonus for immediate implementation of early projects is increasing in τ and η, but decreasing in ω.*

Proof. We have $\Delta_c = \Delta_0 + \delta_c$, where $\Delta_0 = \tau - \pi_1(\omega - \eta)$ and, δ_c is as in equation (14). It is straightforward to verify that Δ_c is increasing in τ and η and decreasing in ω. $\square$

To see how τ affects Δ_c, consider its impact on its two components, Δ_0 and δ_c. Clearly Δ_0 is increasing in τ: early implementation is more advantageous even for an isolated project if it has higher 'front-loaded' returns. But for sequential projects in an environment of limited implementation capacity, prompt implementation of early projects relieves capacity for timely implementation of future projects that are similarly front-loaded.

Likewise it is easy to check that both Δ_0 and δ_c are increasing in η. Higher values of η indicate projects with higher expected returns, boosting the case for implementation, other things being equal.

In contrast, both Δ_0 and δ_c are decreasing in ω. Recall that ω is the value of the environmental asset in the high-value state. Even for an isolated project, the precautionary principle weakens the case for early implementation. The indirect effect in sequential decisions is more subtle: higher ω for *future* projects implies that those later projects will themselves be ones that the decision-maker will find attractive to delay for precautionary reasons. As such the reduced incentive to execute those later proposals as soon as they arise softens the imperative to preserve future implementation capacity – it *relaxes* further the pressure for rushed execution of early projects.

Next, we turn to the impact of the stochastic characteristics of future projects on the optimal implementation profile. Recall that when choosing between early and delayed implementation of Project 1, the importance, or scale ($\tilde{s}_2$), of Project 2 is not yet known. Neither is its future improvability, $\tilde{\pi}_2$, which determines how attractive or unattractive its subsequent postponement might be. Only the probability distributions from which those parameters are drawn are known at the outset. We examine how the stochastic characteristics of these distributions affect the optimal timing of projects.

Result 5 *With congestible implementation capacity the bonus for immediate implementation of Project 1 is larger when future projects are expected to be more important (i.e. larger in scale), in the sense of first-order stochastic dominance.*

Formally, the proof follows directly from inspection of equation (14): δ_c is increasing in both s_h and s_ℓ, so the claim follows. Intuitively, if future decision opportunities are likely to be bigger in scale (more important), their subsequent postponement due to congestion would be costlier.

How does the greater *variability* of the scale of future projects affect the case for early implementation of projects at hand? The specifics of model assume that the scale of future projects is equally likely to low or high, that is $\tilde{s}_2 \in \{s_\ell, s_h\}$. We consider a mean-preserving spread of this point distribution, in which $\tilde{s}_2 \in \{(s_\ell - \epsilon), (s_h + \epsilon)\}$, with $\epsilon > 0$.

Result 6 *With congestible implementation capacity the bonus for immediate implementation of Project 1 is smaller if the importance (scale) of future projects is more dispersed in the sense a mean-preserving spread of $\tilde{s}_2$.*

This result follows from (14).[14] Intuitively, the deferred implementation of Project 1 can impact Project 2 in two possible ways. One, it may simply

[14]Replacing s_ℓ with $(s_\ell - \epsilon)$ and s_h with $(s_h + \epsilon)$, it is easy to verify δ_c is decreasing in

lead to a postponement of Project 2 to period 3, effectively creating a pattern of 'serial postponement': this will be the case when Project 2 turns out to be sufficiently unimportant and/or improvable. But as equation (8) makes clear, when s_h is large enough, the decision-maker does better by abandoning Project 1. The higher payoff in that case is increasing in s_h (and unaffected by s_ℓ), so that a mean-preserving spread of $\tilde{s}$ increases the expected return to deferral of Project 1, reducing the overall gain to its prompt implementation.

Next, consider the impact of variations in the stochastic improvability of future decisions, given by the distribution of random variable $\tilde{\pi}_2$. Recall that Assumption 2 had restricted this to be uniformly distributed in the unit interval. While the choice of a precise distribution delivered closed-form solutions to highlight our central argument, that sharp assumption did not leave any scope for assessing the impact of variations in that distribution. Hence in what follows, we relax Assumption 2.

To study the impact of variations in the distribution of $\tilde{\pi}_2$ on the magnitude of Δ_c, we explore how the net advantage to early implementation of Project 1 varies with particular realizations of $\tilde{\pi}_2$.

- For $\tilde{\pi}_2 > \frac{\tau}{\omega-\eta}$, the best course for Project 2 involves postponement, *regardless of whether or not Project 1 had already been implemented.* If so, the net benefit, summed across all projects, from early implementation of Project 1 equals $\Delta_0 = \tau - (\omega - \eta)\pi_1$.
- For $\tilde{\pi}_2 \leq \frac{\tau}{\omega-\eta}$, if considered in isolation, the decision-maker would implement early. However in a sequence of projects, Project 2 has to compete with legacy projects carried over from earlier. This case admits two sub-possibilities.

  First, for $\tilde{\pi}_2 \in [0, \frac{\tau-(\eta/\tilde{s})}{\omega-\eta})$, Project 2 is important enough to merit immediate execution, even though its implementation implies abandoning any legacy Project 1. In this sub-case the net overall benefit from a strategy of early implementation of Project 1 is simply its return $\tau + \eta$.

  Second, for intermediate values, $\tilde{\pi}_2 \in [\frac{\tau-(\eta/\tilde{s})}{\omega-\eta}, \frac{\tau}{\omega-\eta})$, early implementation of Project 2 would be justified in isolation, but in a sequence

ϵ whenever, consistent with our setting, $s_h > \eta/\tau$, because

$$\left[\tau^2(s_\ell - \epsilon) - \eta^2 \frac{1}{s_h + \epsilon}\right] < \left[\tau^2 s_\ell - \eta^2 \frac{1}{s_h}\right].$$

of projects, it is optimally delayed to enable a return to implementing the legacy project. The presence of a legacy project then results in serial deferral – Project 1 is implemented in period 2, and Project 2 in period 3. The net overall benefit from a strategy that has early implementation of Project 1 is now $\tau + [\tau - (\omega - \eta)\pi_2]\tilde{s}_2$.

We can summarize these cases as follows. Let $e_c(\tilde{\pi}_2)$ denote the value of difference in total payoffs (across all projects) between early and late implementation of Project 1.

$$e_c(\pi_2) = \begin{cases} \tau + \eta & \text{if } 0 \le \pi_2 \le \frac{\tau - (\eta/\tilde{s})}{\omega - \eta} \\ \tau + [\tau - (\omega - \eta)\pi_2]\tilde{s}_2 & \text{if } \frac{\tau - (\eta/\tilde{s})}{\omega - \eta} < \pi_2 < \frac{\tau}{\omega - \eta} \\ \tau & \text{otherwise} \end{cases} \tag{15}$$

By construction the previously-defined Δ_c is the expectation of $e_c(\tilde{\pi}_2)$ across all realizations of $\tilde{\pi}_2$. Given that $e_c(\pi_2)$ is decreasing in π_2 for some realizations, and invariant in others, we have the following result.

Result 7 *With congestible implementation capacity the bonus for immediate implementation of Project 1 is lower if future projects are likely to be more improvable, in the sense of first-order stochastic dominance of the distribution of* π_2.

The more likely future proposals are to describe projects that are improvable, and therefore the more attractive/less unattractive is likely to be their postponement, the lower is the premium associated with keeping a clear desk.

How does an increase in the dispersion (say, in the sense of a mean-preserving spread) of the improvability parameter affect the premium for early implementation? Note that $e_c(\pi_2)$ is neither convex nor concave in π_2, so the impact of that variation is ambiguous.

To illustrate this consider the case in which $\tilde{\pi}_2$ is distributed uniformly in a subset of the unit interval $[a, b]$, where $0 \le a < b \le 1$;

Assumption 3 *The stochastic characteristics of Project 2 are given by the following distributions.*

(a) $\tilde{\pi}_2 \sim U[a,b]$*: the realization of* $\tilde{\pi}_2$ *is distributed uniformly in the interval* $[a,b]$*, where* $0 \leq a < b \leq 1$*. We have* $a < \hat{\pi} < b$*, where* $\hat{\pi} = \tau/(\omega - \eta)$*.*[15]

(b) $\mathrm{pr}(s_h) = \mathrm{pr}(s_\ell) = 0.5$

This assumption allows a slightly more general treatment than Assumption 2, which amounts to the special case in which $a = 0$ and $b = 1$. Replicating the analysis of the previous Section under Assumption 3, we can obtain a modified version of (14), to measure the additional premium for early implementation of Project 1:

$$\delta_c = \frac{1}{4}\frac{(1-\pi_1)}{b-a}\left[\frac{1}{\omega-\eta}\left(\tau^2 s_\ell - \eta^2\frac{1}{s_h} + 2\tau\eta\right) + a\left([(\omega-\eta)a - 2\tau]s_\ell - 2\eta\right)\right]$$

So doing allows us to study the affect of variations in the distribution of the improvability characteristics of projects. For instance, for $\epsilon > 0$, a change from $U[a,b]$ to $U[a+\epsilon, b+\epsilon]$ raises its mean while leaving its variance unchanged. It is straight-forward to verify that an increase in ϵ lowers δ_c. On the other hand, the impact of a mean-preserving spread, from $U[a,b]$ to $U[a-\epsilon, b+\epsilon]$, is ambiguous.

4 Conclusions

The precautionary principle is, rightly, an influential concept in policy analysis (Atkinson et al (2006), Steele (2006), Foster et al (2000)). The principle rationalizes a bias against early action on development projects with irreversible environmental impacts, even when the expected benefits from action exceed expected costs, making the case for deferral to a time when more information is available.

While the desire to keep options open is an enticing one, as much in a policy setting as in our personal lives, such postponement implies risk. There might never be a time when it makes sense to go back to a deferred project, even if it later transpires that it is a good one, because of competing

[15] This additional restriction avoids trivialities. If $a > \hat{\pi}$, then the second period opportunity is always sufficiently improvable to merit postponement of a decision to period 3, eliminating any potential congestion in decision. If $b < \hat{\pi}$, then the second period opportunity is never improvable enough to merit postponement of a decision, reducing the sequential decision problem to a single period choice.

opportunities that arise later. And even if it does make sense to go back, that may come at the expense of having to displace or postpone a later opportunity which, considered in isolation, would have demanded prompt attention.

As such the decision-maker faces a conflict between wanting to wait for more information to allow a really informed decision on each particular projects, but at the same time not wanting to congest implementation capacity more than necessary.

Here we provide a framework that embeds such logic in a simple way, complementing the influential work of Arrow and Fisher by recasting it in a more realistic setting in which a decision-maker faces a stream of proposals but faces implementation constraints – we can only do so many things at once. In such a setting there is a bonus in favor of early execution of projects, even those with uncertain net benefits, that acts as counter-weight to the option value associated with postponement and dilutes the logic of the precautionary principle. We develop an interpretable expression for that premium and characterize scenarios in which it fully versus only partially offsets the option value.

The bonus to keeping a 'clear desk' depends crucially on the decision environment in which the planner finds himself, as described by the distributions from which the characteristics of future projects will be drawn. It is larger the more important future projects are expected to be (in the sense of first-order stochastic dominance) and the less variable that importance (in the sense of mean-preserving spread). In other words a more 'choppy' decision environment – a stream of proposals of very variable quality – diminishes the onus for prompt action. Other things equal it is smaller if future projects are typically expected to be more improvable, though the affect of variability in that improvability is in general ambiguous.

Bibliography

1. Arrow, K. and A. Fisher (1974). "Environmental Preservation, Uncertainty and Irreversibility." *Quarterly Journal of Economics* 88(2): 312-19.

2. Atkinson, G., S. Mourato and D. Pearce (2006). *Cost-Benefit Analysis and the Environment: Recent Developments.* OECD: Paris, France.

3. Che, Y. K. and K. Mierendorff (2018). "Optimal Dynamic Allocation of Attention." *American Economic Review* 109(8): 2993-3029.

4. Crabbe, P. (1987). "The Quasi-option Value of Irreversible Investment: A Comment." *Journal of Environmental Economics and Management* 14(4): 384-5.

5. Foster, K. R., P. Vecchia (2000). "Science and the Precautionary Principle." *Science* 288(5468): 979-81.

6. Geanakaplos J. and P. Milgrom (1991). "A Theory of Hierarchies based on Limited Managerial Attention." *Journal of the Japanese and International Economies* 5(3): 205–25.

7. Gifford, S. and C. Wilson (1995). "A Model of Project Evaluation with Limited Attention." *Economic Theory* 5: 67-78.

8. Grossman, G. M. and C. Shapiro (1986). "Optimal Dynamic R&D Programs." *Rand Journal of Economics* 17(5): 581-93.

9. Mullainathan, S. and E. Shafir (2013). *Scarcity: Why Having Too Little Means So Much* Time Books, New York.

10. Nunamaker J., R. Briggs, G. De Vreede and R. Sprague (2001) 'Introduction to Special Issue: Enhancing Organizations' Intellectual Bandwidth: the Quest for Fast and Effective Value Creation.' *Journal of Management Information Systems* 17(3): 3–8.

11. Presbitero, A. (2016). "Too Much and Too Fast? Public Investment Scaling-up and Absorbtive Capacity." *Journal of Development Economics* 120(1): 17–31.

12. Steele, K. "The Precautionary Principle: A New Approach to Public Decision-making?" *Law, Probability and Risk* 15(1): 19–31.

Declarations

- Funding: None
- Conflicts of interest/Competing interests: none
- Availability of data and material (data transparency): Not applicable
- Code availability (software application or custom code): Not applicable