

CAHIER D'ÉTUDES WORKING PAPER

N° 180

OPTIMAL TIMING OF ENVIRONMENTAL POLICY UNDER PARTIAL INFORMATION

PABLO GARCIA

JANUARY 2024



BANQUE CENTRALE DU LUXEMBOURG

EUROSYSTÈME

OPTIMAL TIMING OF ENVIRONMENTAL POLICY UNDER PARTIAL INFORMATION

PABLO GARCIA

Abstract. I explore the optimal timing of environmental policy when the stock of natural capital is unobserved and can only be imperfectly measured. I present two key insights. First, noisy signals about the natural capital stock blur the inference process, thereby easing the conditions under which policy adoption becomes optimal. Second, the interaction between natural capital stock volatility and the inference process gives rise to new effects that are absent under perfect information. Specifically, the impact of increased volatility on the conditions for optimal policy adoption varies depending on the information set. My work contributes to both the environmental policy timing literature and the field of resource management under incomplete information.

JEL Codes: E20, Q50.

Keywords: Natural capital, Partial information, Optimal stopping.

December, 2023. E-mail: pablo.garciasanchez@bcl.lu. Banque centrale du Luxembourg, Département Économie et Recherche, 2 boulevard Royal, L-2983 Luxembourg. For useful comments and suggestions, I thank Patrick Fève, Paolo Guarda, Olivier Pierrard, Jean-Charles Rochet and BCL colleagues. This paper should not be reported as representing the views of the BCL or the Eurosystem. The views expressed are those of the author and may not be shared by other research staff or policymakers in the BCL or the Eurosystem.

RÉSUMÉ NON TECHNIQUE

Le changement climatique est pertinent pour la politique monétaire de la zone euro, vu qu'il concerne non seulement l'objectif primaire des banques centrales, la stabilité des prix, mais aussi leurs objectifs secondaires, dont le soutien aux politiques économiques générales dans l'Union européenne. C'est ce que Mme Lagarde, présidente de la Banque centrale européenne (BCE), a rappelé à la Commission des affaires économiques et monétaires du Parlement européen lors du dialogue monétaire de novembre 2023. En effet, la BCE a adopté un programme d'action pour le climat en 2022 à l'issue de son évaluation stratégique de la politique monétaire. Entre autres, ce programme d'action prévoit d'améliorer les modèles macroéconomiques afin qu'ils tiennent compte des risques liés au changement climatique.

Dans ce contexte, cet article étudie le moment optimal pour l'adoption de la politique environnementale. Etant donné que les mesures pour protéger l'environnement peuvent comporter des coûts économiques et sociaux, à quel moment faut-il introduire une telle politique de protection ? Ce cahier analyse comment la présence d'incertitude influe sur la réponse à cette question.

En effet, une des difficultés majeures face à la politique environnementale est la mesure du capital naturel, qui englobe les ressources naturelles telles que la géologie, le sol, l'air, l'eau et toutes les formes de vie. Le capital naturel fournit les biens et les services qui rendent possible la vie sur Terre. L'eau que nous buvons, la nourriture que nous mangeons et l'air que nous respirons en sont des exemples concrets. Mesurer des éléments aussi complexes est difficile ; il n'existe pas de métrique simple et universelle. Les décideurs politiques doivent souvent se baser sur plusieurs indicateurs imparfaits pour mesurer l'évolution du capital naturel.

Cette étude propose un modèle mathématique simple qui s'affranchit de l'hypothèse traditionnelle selon laquelle les décideurs politiques auraient une information complète sur l'état du capital naturel. En fonction de ce modèle, le stock de capital naturel se détériore selon un processus stochastique avec tendance baissière et le décideur doit choisir le moment optimal pour arrêter cette détérioration, sachant que cette décision engendre des coûts socio-économiques, tels que la perte de certains emplois ou de certains actifs devenus irrécupérables

ou échoués. La nouveauté du modèle est que le décideur ne peut pas observer directement le niveau du stock de capital naturel mais doit l'inférer à partir d'informations incertaines.

Par exemple, imaginons un décideur responsable de la protection d'un écosystème fragile, sujet à des chocs imprévisibles, tels que des variations climatiques affectant l'apparition d'une maladie liée à la pollution. Il doit choisir le meilleur moment pour mettre en place une politique de conservation en considérant ses coûts socio-économiques ainsi que de l'incertitude entourant la fragilité de l'écosystème.

Les résultats de l'analyse véhiculent deux messages importants. Tout d'abord, l'incertitude rend plus probable l'adoption de la politique. L'intuition est simple : l'incertitude encourage l'implémentation de mesures afin de prévenir des dommages potentiels graves. Ensuite, l'interaction entre la volatilité du capital naturel et l'estimation de son niveau affecte la stratégie optimale à travers des canaux supplémentaires par rapport à la situation sans incertitude. Plus précisément, une hausse de la volatilité du capital naturel a des effets plus complexes sur la stratégie d'adoption, comparé à la situation où le décideur dispose de toutes les informations.

La prise de décision en matière de protection de l'environnement est cruciale ; en fin de compte, l'objectif est de préserver un environnement sain pour les générations présentes et futures. En montrant comment les décisions peuvent tenir compte de l'incertitude et de l'information imparfaite, cette étude révèle qu'une information incomplète sur l'état du capital naturel ne devrait pas servir d'excuse pour retarder l'adoption de politiques de protection. Au contraire, le modèle montre qu'une incertitude accrue facilite les conditions sous lesquelles l'adoption de telles politiques devient optimale.

If heaven had given me a few more years, I would have made Paris the capital of the Universe.

Napoleon Bonaparte

1. INTRODUCTION

When should society implement a costly policy to safeguard the environment? In his seminal works, Pindyck (2000, 2002) addresses this question using the real option theory. In his setup, environmental quality follows a stochastic process with a downward trend. A decision-maker must then choose the optimal time to enact a policy that safeguards the environment, but incurs economic costs. He analyses how structural factors, such as environmental volatility and the discount rate, affect the timing of environmental policy.

Pindyck's work relies on a strong assumption: full information. The controller perfectly observes environmental quality at each instant. In reality, however, measuring something as complex as natural capital (or its inverse, environmental degradation) is hard; there is no simple, granular, universal metric. Hence, policy makers often monitor different indicators to understand how natural capital stocks are changing around the world (Dasgupta, 2021).

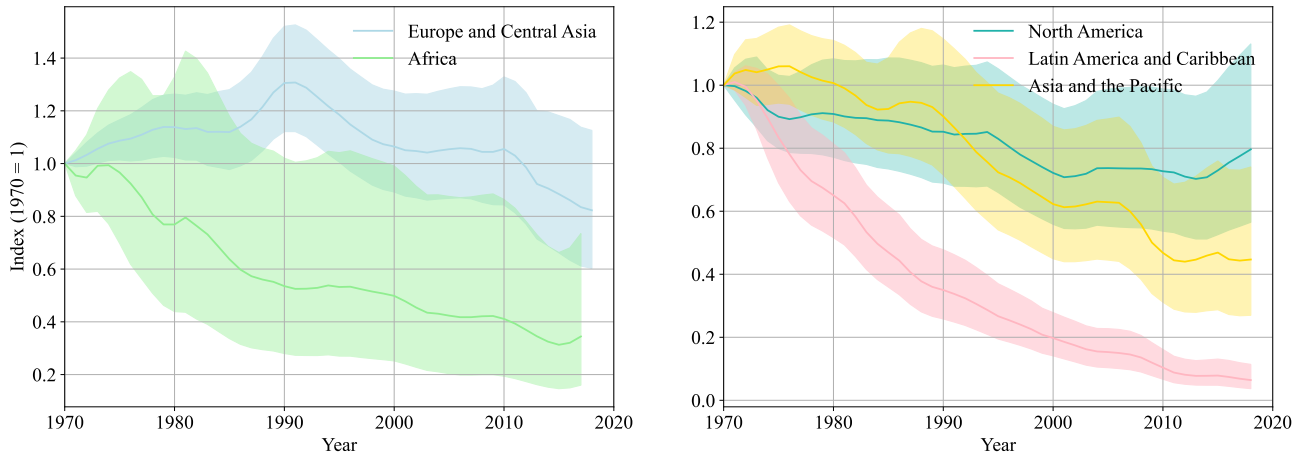
To illustrate, the Living Planet Index (LPI) measures the state of global biodiversity by tracking thousands of population trends of vertebrate species.¹ Widely used by conservationists and policymakers, the index was, for example, adopted by the Convention on Biological Diversity as an indicator of progress towards its target to 'take effective and urgent action to halt the loss of biodiversity'.² As shown in Figure 1, however, significant uncertainty surrounds the index. These wide confidence bands, a recurring feature in biodiversity indicators, highlight the challenges of measuring natural capital stocks.

Therefore, my model relaxes the full information assumption. Instead, it presents an optimal stopping problem under partial information. This still assumes a stochastic process with a negative drift for the stock of natural capital. A controller must choose the optimal time to implement a policy that permanently freezes the natural capital stock at its current level. However, implementing the policy is costly, for it has a negative impact on the reward function. Moreover, there is an exogenous terminal time T by which the policy must have been

¹For more details, please refer to <https://www.livingplanetindex.org>.

²For more details, please refer to <https://www.cbd.int/sp/elements/>

FIGURE 1. Living Planet Index



Notes. The Living Planet Index measures the state of global biological diversity based on population trends of vertebrate species from terrestrial, freshwater and marine habitats. Bold lines show index values and shaded areas represent statistical certainty surrounding the trends (95%). The Living Planet Index is maintained by the World Wildlife Fund and the Zoological Society of London.

adopted. The novel modelling feature is that the controller does not observe the stock of natural capital; that is, she faces *ignorance uncertainty* (Tsur and Zemel, 2014). She must therefore infer the level of the natural capital stock by combining a noisy signal with her own past beliefs.³

A simple example helps to illustrate the rationale of my model. A policymaker is responsible for safeguarding a fragile ecosystem susceptible to unpredictable shocks, like weather changes affecting the outburst of a pollution-induced disease. Specifically, her task is to determine the right moment to implement a conservation policy that preserves the ecosystem at its current state. In making this decision, she faces three challenges. First, adopting the policy involves economic costs, such as job losses and stranded assets. Second, the ecosystem's fragility is not known with certainty, and can only be imperfectly measured. Third, there is a deadline T by which the policy must be in place. For instance, this deadline can arise from institutional commitments to halt the degradation of natural capital by a specific date, or from the intolerable risks of irreversible harm entailed by delaying policy adoption beyond T .

³Formally, the controller maximises expected reward $G(t, x(t))$ by choosing the optimal stopping time $\tau \leq T$. She values natural capital $x(t)$, which follows a geometric Brownian motion with a negative drift coefficient; so $G_x(t, x(t)) > 0$. The controller does not observe $x(t)$; only a disturbed version of it, $s(t)$. Hence, she uses the Kalman-Bucy filter to estimate the conditional distribution of $x(t)$ given the filtration generated by $s(t)$, \mathcal{F}^s . The optimal stopping time $\tau \leq T$ must then be \mathcal{F}^s -adapted.

My model conveys two key messages. First, *a noisy signal about the current natural capital stock eases the conditions under which policy adoption becomes optimal*. As will become clear, policy adoption becomes optimal as soon as the natural capital stock exceeds an endogenous threshold. My model predicts that a noisier signal will lower the threshold. The intuition is straightforward. Noise blurs the inference process, prompting the controller to lower the policy adoption threshold in an attempt to mitigate the deterministic decline in the natural capital stock.

The second key message is that *the interaction between the volatility of the natural capital stock and the inference process gives rise to new effects that are absent in the full information baseline*. Specifically, my model suggests divergent effects of high environmental volatility on the policy adoption threshold when compared to the full information baseline. This occurs because, under partial information, increased volatility heightens the uncertainty surrounding the inference process, a channel that is not present under full information. For instance, in a full information context, my model often suggests that greater environmental volatility leads to higher adoption thresholds, as in [Pindyck \(2000, 2002\)](#). However, partial information can reverse this effect, actually leading to lower thresholds.

For simplicity, I begin by analysing an infinite horizon setup; that is, $T \rightarrow \infty$. This is not the most realistic scenario, because the controller is shielded against extreme negative realisations of the natural capital stock: she can take no action and obtain a zero reward. Nonetheless, letting $T \rightarrow \infty$ enables me to solve the full information baseline analytically and provide straightforward insights into the partial information setup. Once this simpler scenario is fully understood, I move on to the finite horizon setup where $T < \infty$. Though this new setup accepts no closed-form solution, and certain outcomes are slightly more involved, the two key messages emphasised above remain perfectly valid.

As mentioned already, this work contributes to the literature on environmental policy timing. Several papers have followed [Pindyck \(2000, 2002\)](#) in using the real option approach to deal with environmental protection. In highly cited works, [Kassar and Lasserre \(2004\)](#) study biodiversity preservation decisions, and [Saphores and Shogren \(2005\)](#) assess how to optimise the use of pesticides. Also, [Ben Abdallah and Lasserre \(2012\)](#) explores when to stop logging

when the survival of an endangered species hinges on forest habitat. Furthermore, [Nishide and Ohyama \(2009\)](#) extend Pyndick's framework by examining how the size of the underlying economy influences the optimal policy timing. Similarly, [Agliardi and Sereno \(2012\)](#) expand the framework by introducing a public finance dimension. Lastly, [Sims and Finnoff \(2016\)](#) reveal that modelling tipping points leading to irreversible environmental damages lowers the value of delaying policy adoption, for policies that were once effective become ineffective. I deviate from these papers by relaxing the full information hypothesis.

In addition, my paper relates to the literature studying resource management under uncertainty and learning. As [Sloggy et al. \(2020\)](#) note, most models of natural resource management ignore partial information regarding changes in state variables, even though, in reality, owners and managers routinely invest in information on the stock of their resources. There are, nonetheless, important exceptions. [Clark and Kirkwood \(1986\)](#), for instance, explores the optimal harvest policy of a renewable resource when the owners cannot accurately measure current stock levels. In turn, [Roughgarden and Smith \(1996\)](#) finds that the inherent problem of over-fishing is exacerbated by uncertainty in fish stock size and dynamics. More recently, several authors have presented large quantitative models -all related to partially observable Markov decision processes- studying resource management under state uncertainty. For example, [MacLachlan et al. \(2017\)](#) explore the optimal control of bovine tuberculosis in New Zealand cattle when the prevalence of the disease is imperfectly observed. Also, [Memarzadeh and Boettiger \(2018\)](#) argue that underestimating the role of uncertainty results in aggressive decision rules which might lead to the dramatic decline and possible collapse of a population, species, or ecosystem. My work has a different emphasis. It proposes a tractable, stylised optimal stopping setup to assess how the interaction between state uncertainty and volatility affects the timing of environmental policy decisions.

From a mathematical standpoint, my work relies on the optimal stopping literature. Standard references include [Øksendal \(2003\)](#) and [Peskir and Shiryaev \(2006\)](#). However, partial information brings specific challenges that diverge from the general theory. For example, the reward function is not adapted to the filtration generated by the observable process. To address these challenges, I leverage well-known techniques, my key sources being [Bertsekas \(1995\)](#), [Ludkovski \(2009\)](#) and [Zhou \(2013\)](#). Lastly, most of the resulting free boundary partial

differential equations do not accept a closed-form solution. Therefore, I use the explicit finite difference method presented in [Brandimarte \(2013\)](#).

The remainder of the paper is organised as follows. Section 2 studies the full information baseline with an infinite time horizon. Section 3 relaxes the full information assumption. Section 4 confirms that the paper’s key messages remain valid in a finite time horizon setup. Section 5 concludes.

2. OPTIMAL STOPPING UNDER FULL INFORMATION

2.1. Set up. Consider an infinite-horizon environment in continuous time. Let $x(t)$ be a state variable summarising the stock of natural capital. For example, $x(t)$ might refer to forest cover, marine fisheries or the atmosphere’s ability to absorb carbon emissions. Let $\delta x(t)$ be the rate of change of $x(t)$ in the absence of stochastic disturbances. In the real world, two competing forces determine whether δ is positive or negative. On the one hand, natural capital regenerates, e.g., forests generate new growth by sowing, animal populations reproduce, and carbon concentrations gradually leave the atmosphere. On the other hand, human activity degrades most of the planet’s ecosystems. Clearly, the strong deterioration of natural capital observed in recent decades ([Diaz et al., 2019](#)) indicates the dominance of the second force. Hence, I assume $\delta \leq 0$. Lastly, let $W(t)$ be a 1-dimension Brownian motion. Natural capital then evolves by:

$$dx(t) = \delta x(t)dt + \sqrt{\sigma_x}x(t)dW(t), \quad (1)$$

where $\sigma_x \geq 0$ governs how volatile $x(t)$ is. Formally, $x(t)$ follows a geometric Brownian motion with drift δ and volatility $\sqrt{\sigma_x}$. For future reference, let me assume that $x(0)$ is log-normally distributed.

The decision maker can stop process $x(t)$ at any instant $\tau \in (0, \infty)$, and obtain the reward $G(\tau, x(\tau))$. Stopping the process could represent, for instance, the implementation of a new environmental policy that halts the decline of natural capital. Reward $G(\tau, x(\tau))$ would then reflect the social or economic benefits derived from the policy. For mathematical tractability, I assume that the reward obtained per unit of time is zero as long as the process $x(t)$ has not

been stopped.⁴ Therefore, the optimal stopping problem is:

$$V(t, x(t)) = \sup_{\tau \geq t} \mathbb{E}^t G(\tau, x(\tau)) = \sup_{\tau \geq t} \mathbb{E}^t e^{-\rho\tau} \left[\frac{x(\tau)^\alpha}{\alpha} - C \right], \quad (2)$$

where \mathbb{E}^t denotes the expectation operator based on information available at t , $\rho \geq 0$ is the discount rate, $\alpha \in (0, 1)$ determines the coefficient of constant relative risk aversion, and $C \geq 0$ represents the cost associated with stopping the process. This cost captures the social or economic damages from adopting a new environmental policy, including job losses and stranded assets.

In sum, given the stochastic process $x(t)$, the optimal stopping problem is to compute the value function $V^*(t, x(t))$ and characterise the optimal stopping time τ^* at which the supremum is attained.⁵

2.2. Solution. Formally, solving an optimal stopping problem for a Markov process is equivalent to finding the smallest superharmonic function that dominates the reward function on the state space (see e.g. Øksendal, 2003; Peskir and Shiryaev, 2006). Therefore, the optimal stopping problem can be reduced to solving a partial differential equation with a free boundary condition dividing the state space in two regions: a continuation region and a stopping region. This boundary is not known in advanced and must be found as part of the problem's solution. I solve (1)-(2) using this approach.

Let

$$A := \{(t, x) \in \mathbb{R}^+ \times \mathbb{R}^+ : V(t, x(t)) > G(t, x(t))\}$$

be the continuation region; as long as $(t, x(t)) \in A$, stopping the process is not optimal. Similarly, let

$$B := \{(t, x) \in \mathbb{R}^+ \times \mathbb{R}^+ : V(t, x(t)) = G(t, x(t))\}$$

be the optimal stopping region; as soon as $(t, x(t)) \in B$, stopping the process is optimal. The optimal stopping time is defined as $\tau^* := \inf \{t \in \mathbb{R}^+ : (t, x) \notin A\}$. The value function in (2)

⁴Appendix A discusses this assumption.

⁵More formally, let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space hosting the Markov process $x(t)$. Let \mathcal{F}_t^x be the filtration of $x(t)$; that is, the sigma-algebra generated by $x(t)$. A random time $\tau : \Omega \rightarrow \mathbb{R}^+$ is a \mathcal{F}_t^x -stopping time if $\{\omega \in \Omega : \tau(\omega) \leq t\} \in \mathcal{F}_t^x$. Let \mathcal{T}^x be the set of \mathcal{F}_t^x -stopping times. The optimal stopping time for (2) is the supremum taken over the set of all stopping times in \mathcal{T}^x .

then solves the free-boundary problem:

$$\begin{aligned} V_t(t, x(t)) + \mathcal{L}_x V(t, x(t)) &= 0 \quad \text{in } A, \\ V(t, x(t))|_B &= G(t, x(t))|_B, \\ \nabla V(t, x(t))|_{\partial B} &= \nabla G(t, x(t))|_{\partial B}. \end{aligned} \tag{3}$$

Here \mathcal{L}_x is the characteristic operator of $x(t)$ and ∂B the boundary of set B . The third equation captures the well-known high contact (or smooth fit) principle, stating that the optimal boundary ∂B is selected so that the value function is smooth on it. Importantly, system (3) only provides a candidate for the solution of the optimal stopping problem (1)-(2). Hence, to confirm that such a candidate actually provides the optimal solution, a verification theorem is required. As discussed in Appendix B, I rely on the verification theorem 4.28 in Seierstad (2009), which is closely related to theorem 10.4.1 in Øksendal (2003). Furthermore, Appendix B shows that for the optimal stopping problem (1)-(2) to be well-defined, I must assume the following.

Assumption 1. Define $\mu = \frac{(\frac{\sigma_x}{2} - \delta) + \sqrt{(\frac{\sigma_x}{2} - \delta)^2 + 2\rho\sigma_x}}{\sigma_x}$. Then $\alpha < \mu < \frac{\alpha}{1-\alpha}$.

Briefly, this assumption guarantees that the boundary between the continuation and the stopping regions is well-defined and that the candidate solution obtained from solving system (3) is optimal.

Proposition 1. Under Assumption 1, the value function of the optimal stopping problem (1)-(2) is:

$$V(t, x(t)) = \begin{cases} e^{-\rho t} \frac{x^{*\alpha}}{\mu} \left(\frac{x(t)}{x^*} \right)^\mu & \text{for } 0 < x(t) < x^*, \\ e^{-\rho t} \left(\frac{x(t)^\alpha}{\alpha} - C \right) & \text{for } x(t) \geq x^*, \end{cases} \tag{4}$$

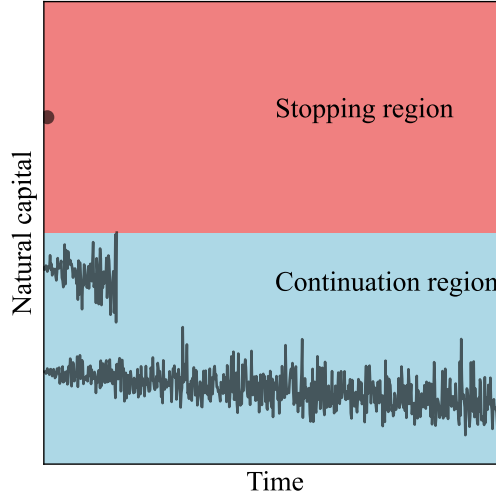
where $x^* = \left(\frac{\alpha\mu C}{\mu - \alpha} \right)^{\frac{1}{\alpha}}$. Moreover, the optimal stopping time is

$$\tau^* = \inf \{ t \in \mathbb{R}^+ : x \geq x^* \}.$$

Proof. See Appendix B. □

Thus, it is optimal to adopt the policy as soon as the natural capital stock exceeds x^* . Figure 2 illustrates the optimal stopping strategy, depicting the three possible alternatives: immediate policy adoption, policy adoption at some future date, or no policy adoption.

FIGURE 2. Graphical illustration of the optimal stopping strategy.



Notes. It is optimal to adopt the environmental policy as soon as the natural capital stock enters the stopping region.

To better understand this finding, consider the analogy of selling an asset with a fluctuating price. Selling only becomes optimal when the price of the asset is high (see e.g. Øksendal (2003)). The same logic applies here. Because the reward function depends on the natural capital stock at the time the policy is adopted, the controller waits for *favourable* realisations of $x(t)$. The structural parameters of the model give the word *favourable* its meaning.

For example, if there were no adoption costs (i.e. $C = 0$), the best policy would be to stop immediately. Given the deterministic decline in the stock of natural capital, this result is self-explanatory. In fact, we do not even need to solve the optimal stopping problem to recognise this. If $C = 0$, the set $U := \{(t, x) : G_t(t, x(t)) + \mathcal{L}_x G(t, x(t)) > 0\}$ is empty. According to Dynkin's formula, $G(t, x)$ is superharmonic with respect to (t, x) , making it optimal to stop $x(t)$ immediately (Øksendal, 2003). However, the analysis that follows focuses on the volatility parameter σ_x .

Proposition 2. *The critical value x^* triggering policy adoption features:*

$$\frac{\partial x^*}{\partial \sigma_x} > 0.$$

In addition, $\forall (t, x) \in \mathbb{R}^+ \times (0, x^)$, the value function features:*

$$\frac{\partial V(t, x(t))}{\partial \sigma_x} > 0.$$

Proof. Immediate computations from Proposition 1. □

The policy adoption threshold increases as the state variable becomes more volatile. This is reminiscent of the incentive to wait that arises with irreversible investment decisions. In problem (1)-(2), the decision maker has the *option* to stop process $x(t)$ at any future time t , whose value is $V(t, x(t))$. Said differently, exercising the option (i.e. stopping the process) is an irreversible decision, with opportunity cost equal to $V(t, x(t))$. On the contrary, inaction over any small time interval only involves a continuous change in $x(t)$. Therefore, the higher the volatility of the state variable, the larger the opportunity cost of exercising the option, and hence, the greater the incentive to wait rather than to adopt the policy now.

3. OPTIMAL STOPPING UNDER PARTIAL INFORMATION

The optimal stopping problem (1)-(2) assumed that the controller perfectly observed the natural capital stock at each instant. However, measuring something as complex as natural capital is hard. Since there is no simple, granular, universal metric, policy makers often rely on proxy signals (Dasgupta, 2021). Therefore, I now drop the full information assumption. The controller no longer observes the natural capital stock directly, but must make an inference about it by combining a noisy signal with her own past beliefs.

3.1. Set up. Suppose now the decision maker no longer directly observes $x(t)$, but only a disturbed version of the state:

$$\hat{x}(t) = \lambda \log x(t) + \sigma_s \hat{W}(t),$$

where $\lambda \in \mathbb{R}^+$, $\sigma_s \geq 0$, and $\hat{W}(t)$ is a white noise. Introducing $s(t) = \int_0^t \hat{x}(u) du$ yields the standard stochastic integral representation:

$$ds(t) = \lambda \log x(t) dt + \sigma_s dB(t), \quad s(0) = 0, \tag{5}$$

where $B(t)$ is a 1-dimension Brownian motion. As an example, if $x(t)$ represents the state of global biodiversity, then $s(t)$ could be one of the biodiversity indicators that monitor biodiversity change, such as the Living Planet Index, the Biodiversity Intactness Index or the Red List Index (Mace et al., 2018).⁶

As before, the decision maker can stop process $x(t)$ at any instant t , and obtain the reward $G(\tau, x(\tau))$. However, her choices are now based on the signal $s(t)$, as the state $x(t)$ is not revealed and can only be inferred through its effects on the drift of $s(t)$. Formally, the decision maker faces the optimal stopping problem under partial information:

$$V(t, s) = \sup_{\tau \geq t, \mathcal{F}_t^s\text{-adapted}} \mathbb{E}^t G(\tau, x(\tau)), \quad (6)$$

The optimal stopping time must be adapted to the filtration generated by $s(t)$, \mathcal{F}_t^s . That is, the decision maker must decide whether to stop process $x(t)$ based only on the history of $s(t)$.⁷ As Ludkovski (2009) explains, partial information problems like (6) feature two particular difficulties. First, the signal $s(t)$ does not reveal the eventual reward of stopping the process (i.e. the reward function is not adapted to \mathcal{F}_t^s). Second, the signal $s(t)$ is not Markovian with respect to \mathcal{F}_t^s . The next subsection tackles these difficulties, and presents the numerical scheme used to solve the partial information model.

3.2. Solution. Solving the optimal stopping problem under partial information requires a two-step inference/optimisation approach. The first filtering step transforms it into an equivalent full information optimal stopping problem. The second step solves the latter problem using standard techniques.

3.2.1. Filtering. Define $z(t) = \log x(t)$. The Ito formula yields:

$$dz(t) = \left(\delta - \frac{\sigma_x}{2}\right)dt + \sqrt{\sigma_x}dW(t). \quad (7)$$

Moreover, $z(0)$ is normally distributed, since $x(0)$ is log-normally distributed. In addition, using the definition of $z(t)$ eq.(5) becomes:

$$ds(t) = \lambda z(t)dt + \sigma_s dB(t). \quad (8)$$

⁶Please refer to chapter 2 in Dasgupta (2021) for more details on these indexes.

⁷In the previous section the optimal stopping time was adapted to the filtration of the state itself.

Crucially, (7) and (8) are a pair of linear stochastic differential equations, and hence fit into the Kalman-Bucy filter framework. Therefore, by Theorem 10.3 in [Liptser and Shiryaev \(2000\)](#), $z(t)|\mathcal{F}_t^s \sim \mathcal{N}(m(t), P(t))$, where

$$dm(t) = \left(\delta - \frac{\sigma_x}{2} \right) dt + \frac{\lambda P(t)}{\sigma_s} d\bar{W}(t), \quad (9)$$

$$ds(t) = \lambda m(t) dt + \sigma_s d\bar{W}(t), \quad (10)$$

$$\frac{dP(t)}{dt} = \sigma_x - \frac{\lambda^2 P(t)^2}{\sigma_s^2}, \quad (11)$$

and $\bar{W}(t)$ is a 1-dimension Brownian motion. The initial values $m(0) \in \mathbb{R}$ and $P(0) \geq 0$ are given. In sum, conditional on the history of the signal $s(t)$, variable $z(t)$ is normally distributed with known mean and variance.

To transform the partially observable optimal stopping problem (6) into a fully observable one, I follow [Bertsekas \(1995\)](#) (see also [Zhou \(2013\)](#) and [Ludkovski \(2009\)](#)). Because the pair $(m(t), P(t))$ provides sufficient statistics for the conditional distribution of $z(t)|\mathcal{F}_t^s$, I can compute the expected reward from stopping the process $x(t)$ conditional on \mathcal{F}_t^s

$$\begin{aligned} \mathbb{E}^t [G(t, x(t)) | \mathcal{F}_t^s] &= \mathbb{E}^t \left[e^{-\rho t} \left(\frac{x(t)^\alpha}{\alpha} - C \right) | \mathcal{F}_t^s \right], \\ &= \mathbb{E}^t \left[e^{-\rho t} \left(\frac{e^{\alpha z(t)}}{\alpha} - C \right) | \mathcal{F}_t^s \right], \\ &= \int_{-\infty}^{\infty} e^{-\rho t} \left(\frac{e^{\alpha u}}{\alpha} - C \right) \phi(u) du, \\ &\equiv g(t, m(t), P(t)), \end{aligned} \quad (12)$$


where $\phi(\cdot)$ is the normal density probability function with mean $m(t)$ and variance $P(t)$. Therefore, problem (6) reduces to the fully observable optimal stopping problem:

$$V(t, m(t), P(t)) = \sup_{\tau \geq t, \mathcal{F}_t^s\text{-adapted}} \mathbb{E}^t g(\tau, m(t), P(t)), \quad (13)$$

where $m(t)$ and $P(t)$ follow equations (9) and (11), respectively. The conditional variance $P(t)$ is deterministic and, fortunately, accepts the following closed-form solution:

$$P(t) = -P^* \frac{1 + \frac{P(0)+P^*}{P(0)-P^*} e^{\frac{2\sqrt{\sigma_x\lambda}}{\sigma_s} t}}{1 - \frac{P(0)+P^*}{P(0)-P^*} e^{\frac{2\sqrt{\sigma_x\lambda}}{\sigma_s} t}}, \quad (14)$$

FIGURE 3. Graphical analysis of the Riccati equation (11)



where $P^* = \frac{\sigma_s}{\lambda} \sqrt{\sigma_x}$. Hence, the optimal stopping problem (9)-(13) only has two dimensions: t and $m(t)$. Before attacking this problem, let us explore the dynamics of $P(t)$. Since eq.(14) is not particularly enlightening, Figure 3 offers a graphical analysis. The left panel plots $\frac{dP(t)}{dt}$ versus $P(t)$, revealing two fixed points: $\pm \frac{\sigma_s}{\lambda} \sqrt{\sigma_x}$. The positive fixed point is clearly stable, since the slope of the function is negative at that point. As a result, small disturbances away from it will be eliminated in time. The opposite is true of the negative fixed point. In this case, any disturbance, however tiny, grows in time.

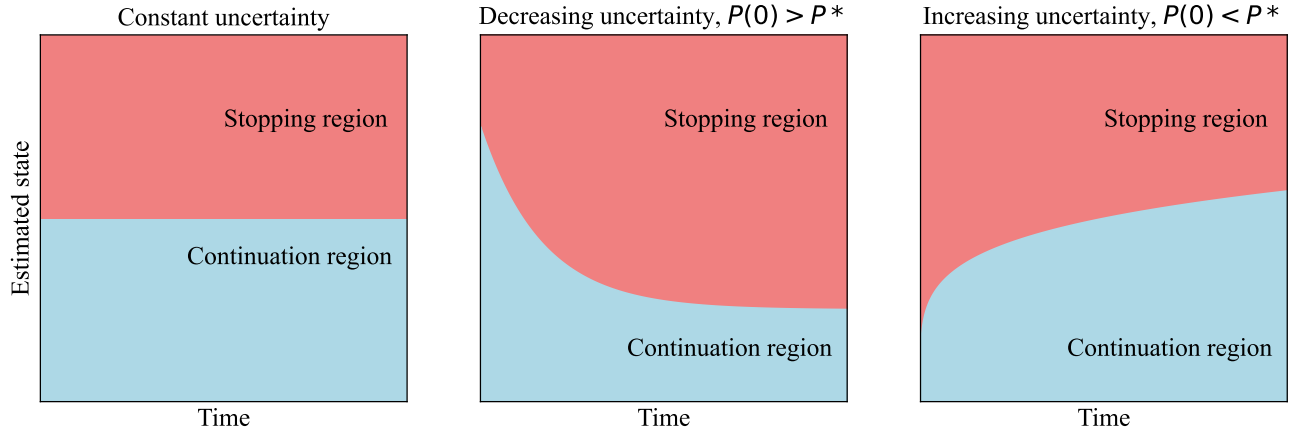
In my set-up, $P(t)$ represents a conditional variance, which is always positive. Therefore, in the relevant domain, $(0, \infty)$, $P(t)$ always approaches the stable fixed point $P^* = \frac{\sigma_s}{\lambda} \sqrt{\sigma_x}$.⁸ For example, if $P(0) < P^*$, the derivative $\frac{dP(t)}{dt}$ decreases over time, and so $P(t)$ is concave down as it asymptotes to the horizontal line $P = P^*$ (right panel in Figure 3). By the same token, if $P(0) > P^*$, $P(t)$ decreases toward P^* and is concave up.

The dynamics of $P(t)$ suggest time-varying effects of partial information on the optimal stopping strategy. Once $P(t)$ has converged to P^* , only the *level* of uncertainty surrounding the estimated state might affect the optimal adoption threshold. However, during the transition period when $P(t)$ approaches P^* , *changes* in uncertainty could play a role too. To better understand these effects, I move on to the second step of the solution procedure.

3.2.2. *Optimisation.* Like the full information baseline, problem (9)-(13) can be reduced to solving a partial differential equation with a free boundary condition dividing the state space into a continuation region and a stopping region. In the current setup, the value function is characterised by the quasi-variational inequality (see, for instance, theorem 5.2.1 in Pham,

⁸Applying L'Hôpital's rule in eq.(14) also shows that, in the relevant domain, $\lim_{t \rightarrow \infty} P(t) = P^*$.

FIGURE 4. Graphical illustration of the optimal stopping strategy.



Notes. Continuation and stopping regions obtained by solving eq.(15). To the best of my knowledge, these boundaries hold for any calibration of the model.

2009, or Proposition 2 in Ludkovski (2009)):

$$\max [V_t(t, x(t)) + \mathcal{L}_m V(t, x(t)), g(t, m(t), P(t)) - V(t, x(t))] = 0, \quad (15)$$

where $\mathcal{L}_m V(t, x(t)) = (\delta - \frac{\sigma_x}{2}) V_m(t, x(t)) + \frac{\lambda^2}{2\sigma_x^2} P(t)^2 V_{mm}(t, x(t))$, and $P(t)$ is given by eq.(14). Unfortunately, eq. (15) does not accept a closed-form solution. Therefore, I follow Brandimarte (2013) and solve it using an explicit finite difference scheme. Appendix D describes the implementation of this numerical technique and confirms its accuracy.

Before discussing some numerical examples, let me illustrate the continuation and stopping regions obtained by solving eq.(15). As shown in Figure 4, if the level of uncertainty is constant (because $P(0) = P^*$), then the adoption threshold is constant too. However, if the level of uncertainty changes over time (because $P(0) \neq P^*$), then the adoption threshold changes too. To the best of my knowledge, these insights remain true under any model calibration.

When uncertainty declines over time (because $P(0) > P^*$), waiting leads to more accurate inferences of the natural capital stock. This *learning opportunity* raises the adoption threshold at the beginning of the time horizon. As uncertainty gradually declines towards its equilibrium value, this learning opportunity weakens, resulting in a lower adoption threshold. Once uncertainty stabilises, the adoption threshold becomes constant. The same logic applies when uncertainty rises over time (because $P(0) < P^*$). Here the uncertainty linked to the inference process increases over time. Therefore, waiting at the beginning of the time horizon becomes less attractive, generating an upward path for the adoption threshold.

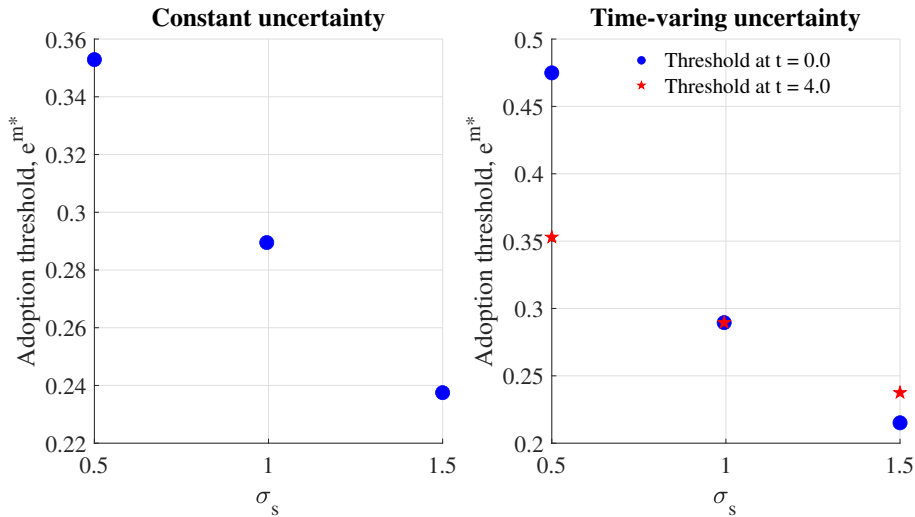
3.3. Numerical examples. Let me begin by introducing the baseline parametrisation: $T = 30$, $\alpha = 0.6$, $\rho = 0.15$, $\delta = -0.01$, $C = 0.5$, $\lambda = 1$, $\sigma_s \in [0.5, 1.5]$ and $\sigma_x \in [0.5, 1.5]$. This parametrisation, while arbitrary, facilitates the numerical analysis and yields some desirable features. For example, it considers one time unit as 3 years, implying a 5% annual interest rate and ensuring that ρT is closed to 0. Furthermore, it leads to $\frac{\mathbb{E}[x(T)]}{\mathbb{E}[x(0)]} \approx 0.75$. Additionally, it guarantees that $P(t)$ gradually converges to its long-term equilibrium, reaching it after 9 years (i.e. roughly at $t = 3$). Lastly, it satisfies Assumption 1. Of course, alternative parameterisations could be used; nevertheless, the discussion below remains valid regardless of the specific parameterisation.

3.3.1. The role of σ_s . Here is the paper's first key message: *a noisy signal lowers the adoption threshold*. Indeed, noise blurs the inference process, prompting the controller to lower the threshold to at least mitigate losses due to the deterministic decline in the natural capital stock. For simplicity, this subsection sets $\sigma_x = 1$, and hence $P^* = \sigma_s$.

First, suppose $P(0) = P^*$, so that the uncertainty associated with the estimated state is constant. The left panel in Figure 5 plots the optimal adoption threshold against the volatility parameter of the signal, σ_s , confirming the anticipated negative link. The mathematical rationale is as follows. As detailed in Appendix C, the reward function $g(t, m(t), P(t))$ increases in P^* , and hence, in σ_s . As a result, raising σ_s directly shrinks the continuation region, for the constraint $V(t, m(t), P(t)) \geq g(t, m(t), P(t))$ binds more easily.

Assume now $P(0) = 1 \neq P^*$. Then the uncertainty linked to the inference process, and hence the adoption threshold, changes over time as $P(t)$ approaches P^* . The right panel in Figure 5 plots the adoption threshold against σ_s at $t = 0$ when $P(t) \neq P^*$ and at $t = 4$ when $P(t) \approx P^*$. In either case, the key message still holds true: the noisier the signal, the lower the threshold.

In addition, $\sigma_s = 0.5$ lowers the adoption threshold over time, while $\sigma_s = 1.5$ raises it. Figure 4 anticipated the logic. When $\sigma_s < 1$, $P(0) > P^*$, so the uncertainty linked to the inference process decreases over time. Waiting provides a learning opportunity, thus raising the policy adoption threshold early on. In contrast, when $\sigma_s > 1$, $P(0) < P^*$, so uncertainty increases

FIGURE 5. Threshold x^* as a function of volatility parameter σ_s .

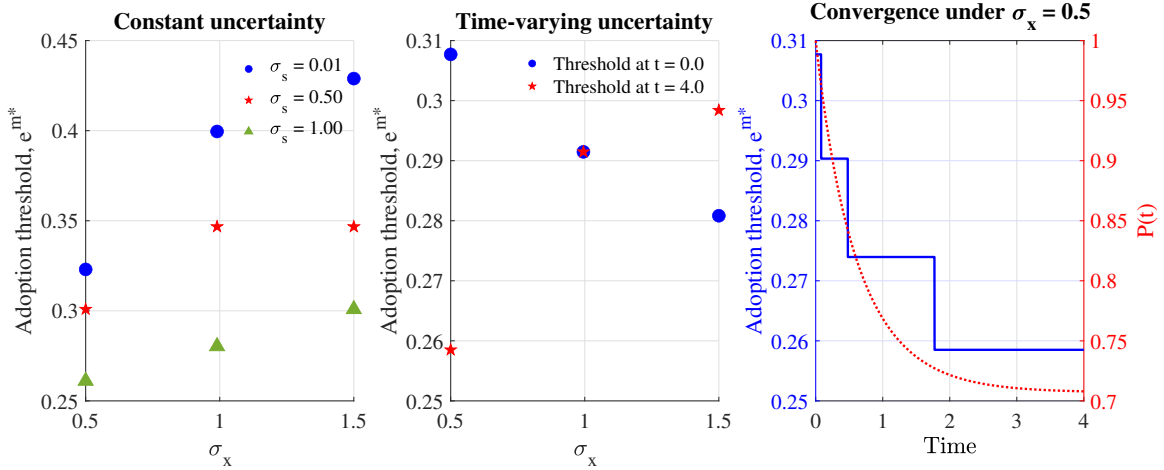
Notes. In the left panel, $P(0) = P^*$; in the right panel, $P(0) = 1 \neq P^*$. In the right panel, $P(t)$ converges to P^* at around $t = 3$.

over time. Waiting becomes less appealing, lowering the policy adoption threshold early on. Lastly, when $\sigma_s = 1$, $P^* = P(0)$, and the threshold is constant.

All told, whether the conditional variance $P(t)$ changes over time has no impact on the paper's first key message. However, this changes when I analyse the volatility parameter σ_x .

3.3.2. *The role of σ_x .* Here is the paper's second key message: *the interaction between the volatility of the natural capital stock and the inference process gives rise to new effects that are absent in the full information baseline.* Under partial information, increasing σ_x results in a less predictable natural capital stock, which in turn amplifies the uncertainty associated with the inference process. As a result, raising σ_x can produce significantly different effects on the adoption threshold compared to what is observed in the full information baseline. For simplicity, this subsection sets $\sigma_s = 1$, and hence $P^* = \sigma_x$.

As before, first assume that $P(0) = P^*$. In this case, the left panel of Figure 6, which plots the threshold against σ_x , reveals that the main finding from the full information baseline remains true. The higher the volatility of the natural capital stock, the higher the threshold. In other words, under a constant conditional variance $P(t)$, the effects of σ_x on the threshold remain monotone. However, partial information weakens the effects of σ_x . This occurs because if

FIGURE 6. Threshold x^* as a function of volatility parameter σ_x .

Notes. In the left panel, $P(0) = P^*$; in the middle panel, $P(0) = 1 \neq P^*$. In the middle panel, $P(t)$ converges to P^* at around $t = 3$. In the right panel, the adoption threshold follows a step function, because the state space is discretised (see Appendix D).

$\sigma_s \rightarrow \infty$, and hence, $P^* \rightarrow \infty$, then $g(t, m(t), P^*) \rightarrow \infty$, making it optimal to stop the process immediately regardless of the value of σ_x .

Now, suppose $P(0) = 1 \neq P^*$. The middle panel plots the adoption threshold against σ_x at two different dates: $t = 0$ and $t = 4$ when $P(t) \approx P^*$. Crucially, raising σ_x reduces the adoption threshold at the beginning of the time horizon, but increases it later on. For example, at $t = 0$, the continuation region shrinks from $(0, 0.31)$ to $(0, 0.26)$ as σ_x goes from 0.5 to 1.5. In contrast, at $t = 4$ the continuation region expands from $(0, 26)$ to $(0, 30)$.

Proposition 2 stated that increased volatility always raised the adoption threshold in the full information baseline. Hence, this middle panel shows that partial information can invert the effect on the adoption threshold from raising the volatility of the state variable.

Again, the intuition is illustrated by Figure 4. When $\sigma_x < 1$, $P^* < P(0)$, so the uncertainty surrounding the estimated state declines over time. Waiting offers a learning opportunity, which raises the adoption threshold early on. The right panel in Figure 6 illustrates these dynamics. On the other hand, when $\sigma_x > 1$, $P^* > P(0)$, so the uncertainty surrounding the estimated state increases over time. Waiting leads to noisier estimates, thus reducing the adoption threshold at $t = 0$.

TABLE 1. Summary of the effects of changes of volatility parameters on the adoption threshold

	$\frac{\Delta x^*}{\Delta \sigma_s} \Delta \sigma_s > 0$	$\frac{\Delta x^*}{\Delta \sigma_x} \Delta \sigma_x > 0$
Constant uncertainty, $P(0) = P^*$	< 0	> 0
Changing uncertainty, $P(0) \neq P^*$	< 0	Time-varying

3.3.3. *Summary.* Table 1 condenses the results from this section. In the next section, I shall prove that the two key messages arising from the table (and emphasised in italics above) also hold in a finite horizon setting.

4. OPTIMAL STOPPING IN A FINITE TIME HORIZON

Thus far, I have considered an infinite horizon model. As a result, if the state $x(t)$ evolved unfavourably, the decision maker could let it run forever to obtain a zero reward. This assumption allowed me to derive an analytical solution for the full information baseline and simplified the intuitions in the partial information setting. However, an infinite horizon might not always be a satisfactory assumption, since some environmental policies do require eventual adoption.

Therefore, this section introduces a finite time horizon $T < \infty$ at which the decision maker must halt process $x(t)$ if she has not already done so. Formally, the optimal stopping problem is:

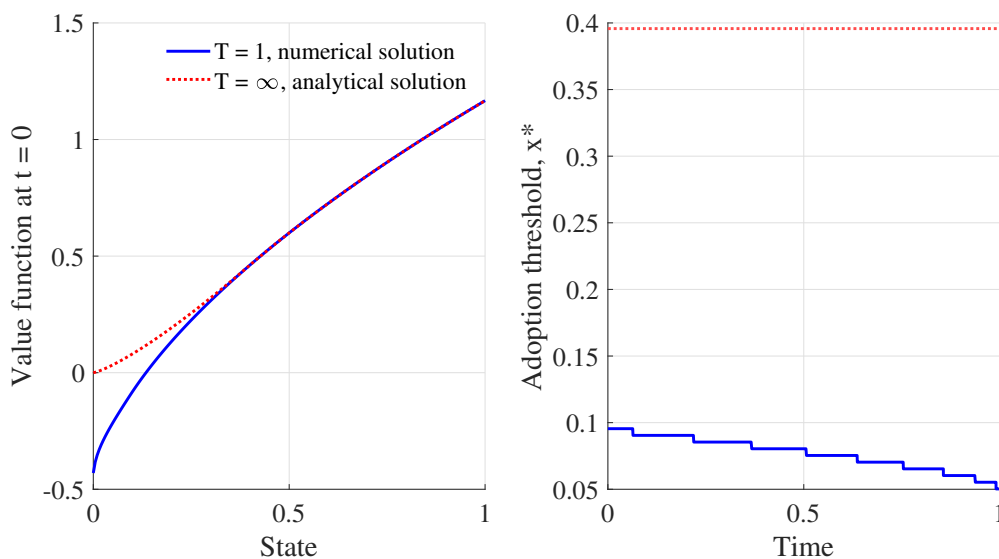
$$V(t, x(t)) = \sup_{\tau \leq T} \mathbb{E}^t G(\tau, x(t)) = \sup_{\tau \leq T} \mathbb{E}^t e^{-\rho\tau} \left[\frac{x(\tau)^\alpha}{\alpha} - C \right]. \quad (16)$$

As before, under full information the decision is based on the history of $x(t)$, and under partial information it is based on the history of $s(t)$. Neither of these problems accepts a closed-form solution. Hence, I use the finite difference method mentioned earlier after adapting the boundary conditions to the new modelling assumption. Unless otherwise stated, I maintain the baseline parametrisation, though I now set $T = 1$, implying a finite time horizon of three years.

First, I examine the shift from $T \rightarrow \infty$ to $T < \infty$ in the full information baseline. Then, I verify that the key messages from the previous section hold in the finite horizon setting.

4.1. Full information baseline. Figure 7 solves the full information setup and compares it with its infinite horizon counterpart. The left panel shows that the value function is no longer bounded by a minimum of zero; it can now dip into negative territory. I mentioned the

FIGURE 7. Value function and policy adoption threshold in full information baseline




intuition earlier: under an infinite time horizon, the decision maker's payoff cannot fall below zero, but under a finite time horizon the payoff can be negative if she reaches T with a low $x(t)$.

It follows that the adoption thresholds will be lower in the finite horizon setting. Since the reward of inaction can now be negative, the decision maker is less inclined to wait for high values of $x(t)$. As T gets closer, this effect intensifies, thus shrinking the continuation region over time. Nash bargaining provides a useful metaphor for understanding these dynamics. When a player's outside option declines, her bargaining position deteriorates, making her more willing to accept terms she might have rejected with a better outside option. The same logic applies here.

Turning to the volatility parameter σ_x , Figure 8 reveals that its effect on the adoption threshold now depends on the curvature of the payoff function. When the latter approaches linearity (e.g. $\alpha = 0.99$), raising σ_x increases the adoption threshold, leading to the same outcome as with an infinite horizon. However, as the payoff function becomes increasingly concave, the decision maker becomes increasingly risk-averse, and raising σ_x reduces the threshold. The α at which the effect of σ_x on the threshold changes sign is roughly 0.85.

Conceptually, Figure 8 recalls the literature on precautionary behaviour. Higher values of σ_x make extreme realisations of $x(t)$ more likely. A nearly risk-neutral decision maker might

FIGURE 8. The effects of the volatility parameter σ_x on the adoption threshold

Notes. Adoption threshold as a function of time in the full information baseline under a finite time horizon. Parameter α governs the curvature of the payoff function.

be inclined to try her luck and wait for a very good outcome before halting the process. Instead, a more risk-averse decision maker might lower her adoption threshold to reduce the chances of bad outcomes. That the worst possible realisation was bounded by zero under an infinite horizon could explain why an increase in σ_x always increased the adoption threshold. Specifically, the lower bound on the payoff limited the negative effects of volatility, but left the positive effects untouched.

4.2. Partial information setting. I now establish that the paper's two main messages remain valid with a finite horizon. First, going from $T \rightarrow \infty$ to $T < \infty$ does not alter the first insight: a noisy signal always reduces the continuation region. This holds true regardless of whether the uncertainty related to the inference process changes over time. The underlying logic is unaffected by the time horizon: noise blurs the inference process, prompting the controller to lower the threshold to insure against the deterministic decline in natural capital. Supporting figures can be found in Appendix E.

I now turn to the paper's second insight: the interaction between the volatility of the natural capital stock and the inference process gives rise to new effects that are absent in the full information baseline. Under $T < \infty$, the impact on the threshold from changes in σ_x is less clear-cut than under $T \rightarrow \infty$. In fact, the adoption threshold in the full information baseline no longer converges to an equilibrium value. Risk aversion also plays an important role. Numerical examples show that interactions among α , σ_x , and the filtering process can affect the

TABLE 2. Adoption threshold at $t = 0$ as a function of parameters

Full information			Partial information			
α	σ_x	Threshold (x^*)	α	σ_s	σ_x	Threshold (e^{m^*})
0.90	0.50	0.52	0.90	0.50	0.50	0.57
0.90	1.50	0.56	0.90	0.50	1.50	0.48
			0.90	1.00	0.50	0.39
			0.90	1.00	1.50	0.33
0.95	0.50	0.62	0.95	0.50	0.50	0.65
0.95	1.50	0.74	0.95	0.50	1.50	0.61
			0.95	1.00	0.50	0.46
			0.95	1.00	1.50	0.41

adoption threshold differently, contingent on the values assigned to the model parameters. Therefore, I cannot provide definitive statements, as I did for the infinite horizon setting, such as ‘raising σ_x always reduces the adoption threshold at the beginning of the time horizon, but increases it later on.’

Nevertheless, my second insight remains perfectly valid under $T < \infty$. Increasing σ_x can still reduce the adoption threshold in the partial information setting, but increases it in the full information baseline. Table 2 provides a few examples featuring this phenomenon. In all cases, $P(0) \neq P^*$. As with an infinite horizon, a time-varying conditional variance appears to be a necessary condition for the effects of σ_x on the threshold to differ in the partial and full information models.

5. CONCLUDING REMARKS

I study the optimal timing of environmental policy when the stock of natural capital is unobserved and can only be imperfectly measured. In reality, however, most environmental measures are not all-or-nothing decisions; they can be continuously adjusted post-adoption. Therefore, jointly considering the optimal adoption timing and the intensity of environmental measures under partial information is a fruitful avenue for future research.

REFERENCES

- Agliardi, E. and Sereno, L. (2012). Environmental protection, public finance requirements and the timing of emission reductions. *Environment and Development Economics*, 17(6):715–739.
- Ben Abdallah, S. and Lasserre, P. (2012). A real option approach to the protection of a habitat dependent endangered species. *Resource and Energy Economics*, 34(3):295–318.
- Bertsekas, D. P. (1995). *Dynamic Programming and Optimal Control*. Athena Scientific.
- Brandimarte, P. (2013). *Numerical Methods in Finance and Economics: A MATLAB-Based Introduction, 2nd Edition*. Wiley.
- Clark, C. W. and Kirkwood, G. P. (1986). On uncertain renewable resource stocks: Optimal harvest policies and the value of stock surveys. *Journal of Environmental Economics and Management*, 13(3):235–244.
- Dasgupta, P. (2021). The economics of biodiversity: The Dasgupta review. Technical report, HM Treasury, London.
- Diaz, S., Settele, J., Brondizio, E. S., Ngo, H. T., Agard, J., Arneth, A., Balvanera, P., Brauman, K. A., Butchart, S. H. M., Chan, K. M. A., Garibaldi, L. A., Ichii, K., Liu, J., Subramanian, S. M., Midgley, G. F., Miloslavich, P., Molnar, Z., Obura, D., Pfaff, A., Polasky, S., Purvis, A., Razzaque, J., Reyers, B., Chowdhury, R. R., Shin, Y.-J., Visseren-Hamakers, I., Willis, K. J., and Zayas, C. N. (2019). Pervasive human-driven decline of life on earth points to the need for transformative change. *Science*, 366(6471):eaax3100.
- Kassar, I. and Lasserre, P. (2004). Species preservation and biodiversity value: a real options approach. *Journal of Environmental Economics and Management*, 48(2):857–879.
- Liptser, R. S. and Shiryaev, A. N. (2000). *Statistics of Random Processes*. Stochastic Modelling and Applied Probability. Springer Berlin, Heidelberg, 2 edition. Published: 06 November 2000.
- Ludkovski, M. (2009). A simulation approach to optimal stopping under partial information. *Stochastic Processes and their Applications*, 119(12):4061–4087.
- Mace, G. M., Barrett, M., Burgess, N. D., Cornell, S. E., Freeman, R., Grooten, M., and Purvis, A. (2018). Aiming higher to bend the curve of biodiversity loss. *Nature Sustainability*, 1:448–451.
- MacLachlan, M. J., Springborn, M. R., and Fackler, P. L. (2017). Learning about a moving target in resource management: Optimal bayesian disease control. *American Journal of Agricultural Economics*, 99:140–162.

- Memarzadeh, M. and Boettiger, C. (2018). Adaptive management of ecological systems under partial observability. *Biological Conservation*, 224:9–15.
- Nishide, K. and Ohyama, A. (2009). Using real options theory to a country's environmental policy: considering the economic size and growth. *Operational Research: An International Journal*, 9(3):229–250.
- Øksendal, B. (2003). *Stochastic Differential Equations: An Introduction with Applications*. Universitext. Springer Berlin, Heidelberg.
- Peskir, G. and Shiryaev, A. (2006). *Optimal Stopping and Free-Boundary Problems*. Lectures in Mathematics. ETH Zürich. Birkhauser Basel, 1 edition.
- Pham, H. (2009). *Continuous-time Stochastic Control and Optimization with Financial Applications*. Stochastic Modelling and Applied Probability. Springer Berlin, Heidelberg, 1 edition.
- Pindyck, R. S. (2000). Irreversibilities and the timing of environmental policy. *Resource and Energy Economics*, 22(3):233–259.
- Pindyck, R. S. (2002). Optimal timing problems in environmental economics. *Journal of Economic Dynamics and Control*, 26(9):1677–1697. Special issue in honour of David Kendrick.
- Roughgarden, J. E. and Smith, F. (1996). Why fisheries collapse and what to do about it. *Proceedings of the National Academy of Sciences*, 93:5078.
- Saphores, J.-D. M. and Shogren, J. F. (2005). Managing exotic pests under uncertainty: optimal control actions and bioeconomic investigations. *Ecological Economics*, 52(3):327–339. Integrating Ecology and Economics in Control Bioinvasions.
- Seierstad, A. (2009). *Stochastic Control in Discrete and Continuous Time*. Springer, Boston, MA.
- Sims, C. and Finnoff, D. (2016). Opposing Irreversibilities and Tipping Point Uncertainty. *Journal of the Association of Environmental and Resource Economists*, 3(4):985–1022.
- Sloggy, M. R., Kling, D. M., and Plantinga, A. J. (2020). Measure twice, cut once: Optimal inventory and harvest under volume uncertainty and stochastic price dynamics. *Journal of Environmental Economics and Management*, 103(C).
- Tsur, Y. and Zemel, A. (2014). Dynamic and stochastic analysis of environmental and natural resources. In Fischer, M. M. and Nijkamp, P., editors, *Handbook of Regional Science*. Springer, Berlin, Heidelberg.
- Zhou, E. (2013). Optimal stopping under partial observation: Near-value iteration. *IEEE Transactions on Automatic Control*, 58(2):500–506.

APPENDIX A. ON THE REWARD FUNCTION $G(t, x(t))$

For mathematical tractability, my model uses a *Mayer* performance measure, where the reward function $G(t, x(t))$ only depends on the stock of natural capital at the stopping time τ . However, an alternative *Lagrange* performance measure could involve an integral process of the form $\int_0^\tau L(x(t))dt$, with $L : \mathbb{R}^+ \rightarrow \mathbb{R}^+$.

This appendix argues that, given the model's structure, these two approaches need not differ significantly. Since policy adoption freezes the natural capital stock at its current level, the dynamics of $x(t)$ can be described as

$$dx(t) = \begin{cases} \delta x(t)dt + \sqrt{\sigma_x}x(t)dW(t) & \text{if } t \leq \tau, \\ 0 & \text{if } t > \tau. \end{cases}$$

Here $\tau \in (0, T]$ is the stopping time, and T is the exogenous deadline for policy implementation. Unlike the main text, the controller selects τ to maximize an infinite sum of future utilities:

$$V = \mathbb{E} \int_0^\infty e^{-\rho t} u(x(t))dt - e^{-\rho\tau} C,$$

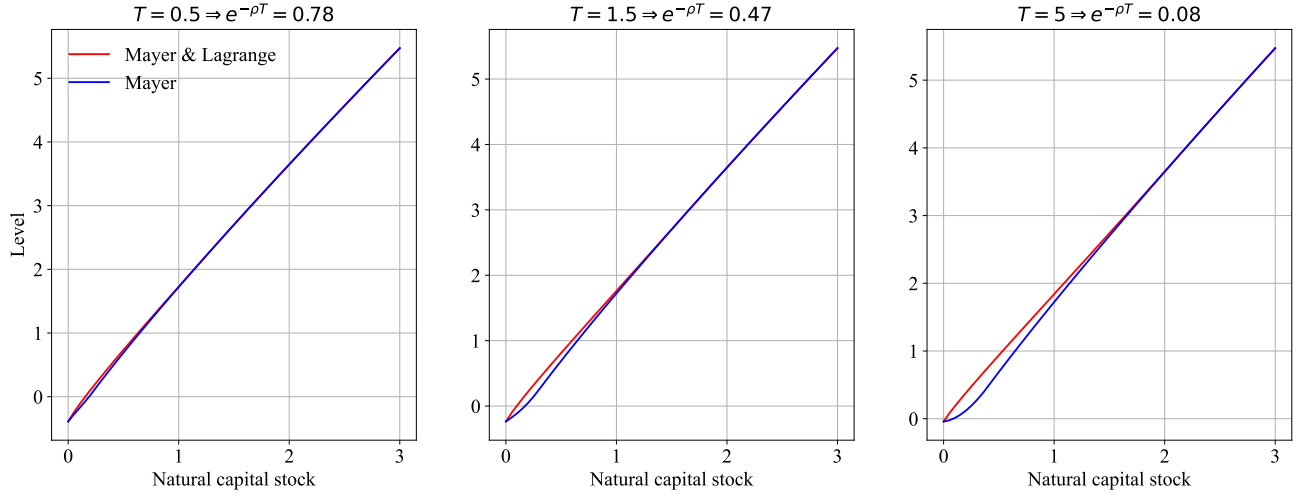
where $u(\cdot) : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ captures the utility flows derived from the stock of natural capital, $x(t)$, at each instant. Inserting the dynamics of $dx(t)$ yields

$$V = \underbrace{\mathbb{E} \int_0^\tau e^{-\rho t} u(x(t))dt}_{\text{Lagrange}} + \underbrace{\mathbb{E} e^{-\rho\tau} \left[\frac{u(x(\tau))}{\rho} - C \right]}_{\text{Mayer}}.$$

Therefore, maximising future utilities is equivalent to optimising a combination of Lagrange and Mayer performance measures.

Nonetheless, the stopping time τ , and hence the upper limit of the integral process, are bounded by T . Hence, for relatively small values of T , the Mayer performance measure should dominate the reward function V . In such cases, neglecting the Lagrange term appears a reasonable assumption.

Figure 9 confirms this intuition. It shows that as T decreases, the value function resulting from considering only a Mayer performance measure (blue line) converges towards the value

FIGURE 9. Value function V under different terminal times T 

Notes. Value function V plotted against state variable, x . The parameter values are: $\alpha = 0.90$, $\delta = -0.1$, $\sigma = 1.0$, $\rho = 0.5$ and $C = 0.5$.

function derived from considering a combination of Lagrange and Mayer performance measures (red line).⁹ When $T = 0.5$, both value functions are nearly identical. It is worth noting that with the discount rate used here ($\rho = 0.5$), the time horizon $(0, 0.5)$ can be interpreted as spanning 20 quarters with an annual interest rate of 4%, which is not unreasonably short. Even for larger terminal times, the disparities between the value functions remain modest.

All told, ignoring the Lagrange performance measure is not an overly strong assumption.

APPENDIX B. PROOF OF PROPOSITION 1

Considering the partial differential equation $V_t + \mathcal{L}_x V = 0$, let me seek a solution in the form of $v(t, x) = e^{-\rho t} x^\mu$. This transforms the equation into the following expression:

$$0 = -\rho + \mu\delta + \frac{1}{2}\sigma_x\mu(\mu - 1),$$

which has two solutions for μ . The first solution is positive:

$$\mu_1 = \frac{(\frac{\sigma_x}{2} - \delta) + \sqrt{(\delta - \frac{\sigma_x}{2})^2 + 2\rho\sigma_x}}{\sigma_x},$$

⁹To obtain these value functions, I solve the corresponding optimal stopping problems numerically using the finite difference scheme in Appendix D.

and the second solution μ_2 , which arises from using a minus sign in front of the square root, is negative. The general solution for v can then be expressed as: $v(t, x) = e^{-\rho t} (c_1 x_1^\mu + c_2 x_2^\mu)$, where c_1 and c_2 are two arbitrary constants.

Next, I guess that the continuation region A has the form

$$A = \{(t, x) \in \mathbb{R}^+ \times \mathbb{R}^+ : 0 < x < x^*\},$$

for some $x^* > 0$. Now, for $v(t, x)$ to be bounded as $x \rightarrow 0$, c_2 must be equal to zero. Then I have two constants to determine $\{c_1, x^*\}$ and two boundary conditions:

$$\begin{aligned} c_1 x^{*\mu_1} &= \frac{x^{*\alpha}}{\alpha} - C, \\ c_1 \mu_1 x^{*\mu_1 - 1} &= x^{*\alpha - 1}. \end{aligned}$$

Solving this system of equations yields $c_1 = \frac{x^{*\alpha - \mu}}{\mu}$ and $x^* = \left(\frac{\alpha \mu C}{\mu - \alpha}\right)^{\frac{1}{\alpha}}$. By substituting c_1 back into $v(t, x)$, I obtain the final expression: $v(t, x) = e^{-\rho t} \frac{x^{*\alpha}}{\mu_1} \left(\frac{x}{x^*}\right)^{\mu_1}$ for $x \in A$. In addition, if $x \notin A$, the controller stops immediately, and the value function equals the payoff function.

This is my candidate for the optimal solution. As mentioned in the main text, however, a verification theorem is required. I rely on theorem 4.28 in [Seierstad \(2009\)](#), which is closely related to theorem 10.4.1 in [Øksendal \(2003\)](#). My candidate solution satisfies the following properties:

- $v(t, x)$ is twice continuously differentiable everywhere and Lipschitz continuous in $[0, \infty)$.
- The equation $V_t + \mathcal{L}_x V = 0$ holds in A . Moreover, $v(t, x^*) = g(t, x^*)$.
- Since $v(t, 0) > G(t, 0)$ and $\frac{v_x(t, x)}{G_x(t, x)} < 1 \forall x \in (0, x^*)$, $v(t, x) > G(t, x)$ in A . By construction, furthermore, $v(t, x) = G(t, x)$ in the stopping region B .
- $v_t(t, x) + \mathcal{L}_x v(t, x) < 0$ in B ; that is, $G_t(t, x) + \mathcal{L}_x G(t, x) < 0$ for $x > x^*$. To see this, note that $G_t(t, x) + \mathcal{L}_x G(t, x) = e^{-\rho t} [\rho C - (\rho - \delta + (1 - \alpha)\frac{\sigma_x^2}{2})] := \Omega(t, x)$. Under the parameter restriction, $\mu < \frac{\alpha}{1 - \alpha}$, $\Omega(t, x^*) < 0$. Given that $\Omega_x(t, x) < 0 \forall x$, I have $G_t(t, x) + \mathcal{L}_x G(t, x) < 0$ in B , as initially claimed.
- Both $\lim_{t \rightarrow \infty} G(t, x(0))$ and $\lim_{t \rightarrow \infty} v(t, x(0))$ equal 0. For some positive constants β , $c_v, c_G, d_1, d_2, d_3, d_4$, for all $(t, x) \in \mathbb{R}^+ \times (0, x^*)$, $|v_x(t, x)| \leq c_v e^{-\beta t}$, and for all $(t, x) \in \mathbb{R}^+ \times \mathbb{R}^+$, $|G_x(t, x)| \leq c_G e^{-\beta t}$, $|\sqrt{\sigma_x} x| \leq d_1 + d_2 |x|$, $|\delta x| \leq d_3 + d_4 |x|$, $\beta > d_4$.

Therefore, my candidate solution satisfies all the conditions required by the verification theorem 4.28 in [Seierstad \(2009\)](#). As a result, $v(t, x) = V(t, x)$ and $\tau^* = \inf \{t \in \mathbb{R}^+ : x \geq x^*\}$ is optimal.

APPENDIX C. ON THE LINK BETWEEN P AND $g(t, m, P)$

I now show that the payoff function $g(t, m, P)$ increases in P . As stated in the main text, $g(t, m, P)$ is given by:

$$g(t, m, P) = \int_{-\infty}^{\infty} e^{-\rho t} f(u) \phi(u) du,$$

where $f(u) = \frac{e^{\alpha u}}{\alpha} - C$ and $\phi(u)$ is the normal probability density function with mean m and variance P .

Consider the following thought experiment. Suppose $\sigma_s \rightarrow 0$, and hence, $P \rightarrow 0$. Then, $\phi(u)$ approaches a Dirac delta function centered at m . That is, the probability density function becomes infinitely peaked at m , so that $g(t, m, P) \rightarrow f(m)$. Denote that value by g^* . Now, suppose σ_s increases slightly. Since $\phi(u)$ is symmetric, $g(t, m, P)$ now results from equally weighting the value of $f(u)$ immediately at the left and right of m . Denote that value g^{**} . Since $f(u)$ is convex, $g^{**} > g^*$.

This reasoning holds true for any two values of σ_s , thus proving that $g(t, m, P)$ increases in P .

APPENDIX D. SOLUTION METHOD: A EXPLICIT FINITE DIFFERENCE SCHEME

Finite difference methods approximate each partial derivative with a quotient; thereby transforming the functional equation into a set of algebraic equations. I begin by setting up a discrete grid with respect of t and m . Though the domain for the partial differential equation reaches $+\infty$ in both dimensions and $-\infty$ in the m -dimension, I must bound it in some way for computational purposes. Therefore, let T and M_h be large enough numbers playing the role of $+\infty$ and M_l be a small enough number playing the role of $-\infty$. The grid then consists of points (t, m) such that:

$$m = M_l, \Delta m, 2\Delta m, \dots, M\Delta m \equiv M_h,$$

$$t = 0, \Delta t, 2\Delta t, \dots, N\Delta t \equiv T,$$

where M and N are positive integers. Approximating the first derivative with respect to m by a central difference and the first derivative with respect to t by a backward difference

transforms the first argument of eq. (15) into:

$$\begin{aligned} \frac{V(t, m) - V(t - \Delta t, m)}{\Delta t} + \left(\delta - \frac{\sigma_x}{2} \right) \frac{V(t, m + \Delta m) - V(t, m - \Delta m)}{2\Delta m} \\ + \frac{\lambda^2}{2\sigma_s^2} P(t)^2 \frac{V(t, m + \Delta m) - 2V(t, m) + V(t, m - \Delta m)}{\Delta m^2} = 0. \end{aligned} \quad (17)$$

Setting the right boundary conditions is essential. Based on the insights discussed in the full information baseline, when m is very large, I expect the policy threshold to have been crossed. Hence, I have:

$$V(t, M_h) = g(t, M_h, P).$$

Similarly, $m \rightarrow -\infty$ would indicate that the state, $x(t)$, had reached 0. Were that to happen, $x(t)$ would remain there forever, as it follows a geometric Brownian motion. So the policy threshold would never be adopted; hence I have

$$V(t, M_l) = 0.$$

The last boundary condition follows naturally from the exponential discounting term:

$$V(T, m) = 0.$$

In grid notation:

$$V(T, i\Delta m) = 0, \quad i = 0, 1, \dots, M,$$

$$V(j\Delta t, M_h) = g(j\Delta t, M_h, P(j\Delta t)), \quad j = 0, 1, \dots, N,$$

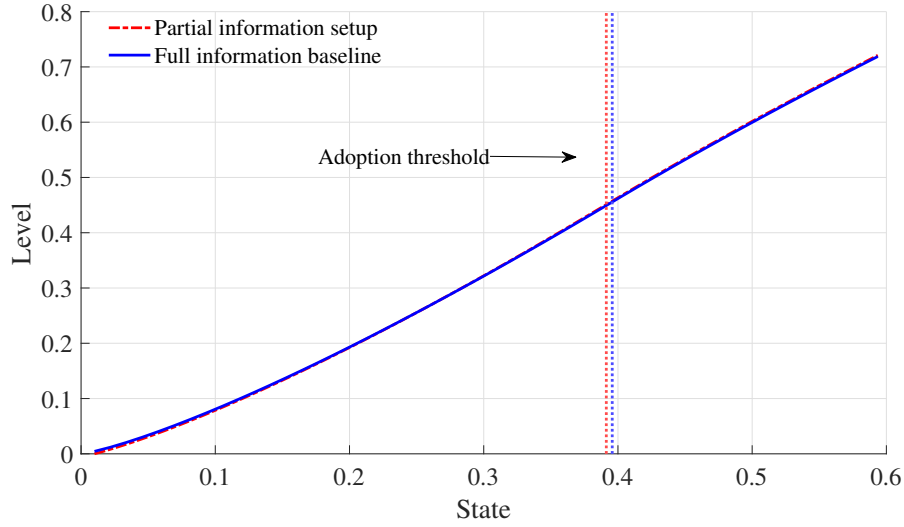
$$V(j\Delta t, M_l) = 0, \quad j = 0, 1, \dots, N.$$

Lastly, I need to consider the max operator in eq. (15) related to the free boundary condition arising from the possibility of early policy adoption. Hence, the value function at each grid point cannot be lower than the immediate reward if the policy is adopted:

$$V(j\Delta t, i\Delta m) \geq g(j\Delta t, i\Delta m, P(j\Delta t)).$$

If I evaluate eq. (17) at $t = N\Delta t$, $V(N\Delta t, m)$, $V(N\Delta t, m - \Delta m)$ and $V(N\Delta t, m + \Delta m)$ are known from the boundary conditions; the only unknown value is $V((N - 1)\Delta t, m)$, which can be obtained as an explicit function of known values. Therefore, rewriting eq. (17), I get

FIGURE 10. Comparison of value functions and adoption thresholds.



Notes. Value functions are plotted against the state variable (x in the full information baseline and e^m in the partial information setup) when $t = 0$. The parameter values used in the exercise are: $dt = 6e - 3$, $dm = 0.08$, $M_h = -0.5$, $M_l = -4.6$, $T = 30$, $\alpha = 0.60$, $\rho = 0.15$, $\delta = -0.01$, $C = 0.5$, $\lambda = 1$, $\sigma_s = 0.01$, $\sigma_x = 1$ and $P(0) = P^*$.

an explicit scheme:

$$V(t - \Delta t, m) = \alpha_1(t)V(t, m) + \alpha_2(t)V(t, m + \Delta m) + \alpha_3(t)V(t, m - \Delta m), \quad (18)$$

for $t = N\Delta t, (N - 1)\Delta t, \dots, \Delta t$ and $m = \Delta m, 2\Delta m, \dots, (M - 1)\Delta m$, where


$$\begin{aligned} \alpha_1(t) &= 1 - \frac{\lambda^2 P(t)^2}{2\sigma_s^2} \frac{dt}{dm^2}, \\ \alpha_2(t) &= \frac{\Delta t}{\Delta y} \left[\frac{(\delta - \frac{\sigma_x}{2})}{2} + \frac{\lambda^2 P(t)^2}{2\sigma_s^2 \Delta m} \right] \\ \alpha_3(t) &= \frac{\Delta t}{\Delta y} \left[\frac{\lambda^2 P(t)^2}{2\sigma_s^2 \Delta m} - \frac{(\delta - \frac{\sigma_x}{2})}{2} \right]. \end{aligned}$$

Imposing the max operator is the final step. After computing $V(t - \Delta t, m)$ using eq. (18), I check for the possibility of policy adoption, and set

$$V(t - \Delta t, m) = \max [V(t - \Delta t, m), g(t - \Delta t, m, P)].$$

One technical point deserves further comment. The payoff function $g(t, m, P)$, given by eq.(12), cannot be written in terms of elementary functions; hence, I approximate it using a standard quadrature method.

FIGURE 11. Policy adoption threshold as a function of volatility parameter σ_s when $P(0) = P^*$




Notes. Adoption threshold as a function of σ_s in the partial information setup under a finite time horizon. The uncertainty linked to the inference process is constant, since $P(0) = P^*$.

I conclude by testing the accuracy of the numerical scheme. If $\sigma_s \rightarrow 0$ and $P(0) = P^*$, the uncertainty surrounding the estimated state is nil, making the partial information setup equivalent to the full information baseline. Figure 10 compares the closed-form solution for the value function and the adoption threshold in the full information baseline with their numerical counterparts in the partial information setup under $\sigma_s = 0.01$ and $P(0) = P^*$. Specifically, the value functions are plotted against the state variable (x in the full information baseline and e^m in the partial information setup) at $t = 0$. The numerical solution does a great job at both approximating the value function and finding the boundary between the continuation and the stopping regions.

APPENDIX E. IMPACT OF σ_s ON THE ADOPTION THRESHOLD WITH A FINITE TIME HORIZON

As stated in the main text, going from $T \rightarrow \infty$ to $T < \infty$ does not alter the first insight: a noisy signal always reduces the continuation region. Figures 11 and 12 confirms that this holds true regardless of whether the uncertainty related to the inference process changes over time.

FIGURE 12. Policy adoption threshold as a function of volatility parameter σ_s when $P(0) \neq P^*$



Notes. Adoption threshold as a function of σ_s in the partial information setup under a finite time horizon. The uncertainty linked to the inference process changes over time, since $P(0) \neq P^*$.



BANQUE CENTRALE DU LUXEMBOURG

EUROSYSTEME

2, boulevard Royal
L-2983 Luxembourg

Tél.: +352 4774-1
Fax: +352 4774 4910

www.bcl.lu • info@bcl.lu