WORKING PAPERS

N° 1705

JANUARY 2026

“Institutional Design For Environmental Acts”

Louise Guillouet and David Martimort

Toulouse School of Economics

Institutional Design For Environmental Acts

Louise Guillouet
International Finance Corporation, guillouetlouise@gmail.com

David Martimort
Toulouse School of Economics, david.martimort@tse-fr.eu

Abstract. Environmental decisions often involve irreversibility and uncertainty. We develop a dynamic model in which actions yield flow benefits but, once an unknown tipping point is crossed, the accumulated stock of past actions raises the risk of irreversible catastrophe. Because constitutions are incomplete contracts, contingent policies cannot be fully specified, and authority is delegated to decision-makers who observe the current stock but hold biased beliefs, decreasing welfare. Such delegation is thus costly. Imposing caps on early actions can limit discretion and improve welfare, providing foundations for the *Precautionary Principle* as a second-best institution.

Keywords. Environmental Risk, Tipping Point, Uncertainty, Irreversibility.
JEL classification. D83, Q55.

1. Introduction

Many economically valuable activities – fossil-fuel consumption, land conversion, etc. – produce flow payoffs today, while simultaneously contributing to cumulative stocks that may trigger discontinuous changes in ecological or climatic systems. Society may not know where the relevant tipping point (thereafter TP) lies, whether it has already been crossed, or how close it is. The combination of irreversibility and uncertainty has long been central to environmental economics (Arrow and Fisher, 1974). It also raises an institutional question: who should decide, when, and under what constraints? The policy response is often framed through the *Precautionary Principle*. In the absence of reliable assurances of safety, decision-makers refrain from actions that may harm future generations.

This paper offers a mechanism-design rationale for precautionary caps in a dynamic environment with uncertain TPs and institutional incompleteness. Building on Guillouet

We thank Derek Lemoine for kindly having invited us to prepare this paper for the *2026 AEA Meeting* and Ian Ball and Liam Lods for useful comments. This paper has also been funded by the Agence Nationale de la Recherche under grant ANR-17-EURE-0010 (Investissements d'Avenir program) and under grant ANR-24-CE26-7833. The findings, interpretations, views, and conclusions expressed herein are those of the authors and do not necessarily reflect the views of the Executive Directors of the*International Finance Corporation* or of the *International Bank for Reconstruction and Development* (the *World Bank*) or the governments they represent. The usual disclaimer applies.

and Martimort (2025), we analyze an economy in which (i) actions generate current benefits, (ii) the accumulated stock of past actions increases the probability of catastrophe once an unknown TP is crossed, and (iii) the location of that TP is uncertain and may already have been passed when acting.

Two economic forces organize the analysis. The optimal policy follows a stock-based feedback rule that trades off current benefits against the shadow cost of moving the system closer to the (maximal) threshold. This is an *Irreversibility Effect*: higher actions today raise the stock, thereby increasing catastrophe exposure and lowering the value of acting. With uncertainty about the TP location, a second force emerges. If the agent chooses higher actions today and the system survives, then survival is more informative that the TP lies ahead. In this sense, actions influence how future survival updates beliefs. This generates a *Pseudo-Learning Effect*: taking higher actions and surviving brings "good news" and makes DM more cautious in the future. Relative to the case where the TP is known, uncertainty strengthens precaution even though no direct signal is observed.

Institutional frictions undermine these forces. Constitutions cannot encode fully contingent plans in terms of the stock. Instead, authority is delegated to a sequence of DMs who can commit only for a short time and who discount the future using their own biased priors. Incoming DMs may underweight the inference from past survival—whether due to bounded rationality, political turnover, or systematic biases—and fail to internalize the pseudo-learning externality that current actions impose on future behavior. As a result, delegated policy is too aggressive early on.

We study simple constitutional constraints that are feasible under incompleteness, such as temporary caps (or floors) on early actions. These rules resemble delegation mechanisms in the mechanism design literature (Holmström, 1984), but here the motivation is dynamic and informational: constraints are most valuable when belief-driven distortions are largest. When delegated DMs are over-pessimistic (assigning too much probability to having crossed the TP), an early cap is optimal under broad conditions. This result provides a formal underpinning for precautionary restrictions as second-best institutions.

2. Model

A representative agent chooses an action (e.g. consumption, production) $x(\tau)$ at each date τ. Time is continuous, and payoffs are discounted at rate $r > 0$. The flow payoff is $u(x(\tau)) = \zeta x(\tau) - \frac{x^2(\tau)}{2}$ and the myopic (static) optimum is $x = \zeta$. Actions accumulate into a stock

$X(t)$. An environmental catastrophe may occur at an arrival rate $\theta(t)$ that depends on whether the stock has crossed a tipping point $\tilde{X}$, $\theta(t) = \Delta\, \mathbb{1}_{\{X(t)>\tilde{X}\}}$, where $\Delta > 0$ is the hazard rate once the TP is crossed. If the TP is crossed at $t < \tau$, the probability of surviving until τ is $e^{-\Delta(\tau-t)}$. After a catastrophe, no further payoffs are obtained.

The TP takes one of two values: $\tilde{X} = 0$ (prob. q) and $\tilde{X} = \overline{X} > 0$ (prob. $1-q$). If $\tilde{X} = 0$, catastrophe risk is present immediately; if $\tilde{X} = \overline{X}$, risk begins only after the stock exceeds $\overline{X}$. Importantly, it is never observed whether the TP has been crossed until $\overline{X}$ is passed.[1] Beyond $\overline{X}$, the TP must have been crossed, so the optimal action then equals the myopic level ζ, yielding a payoff $\mathcal{V}_\infty = \frac{u(\zeta)}{\lambda}$, where $\lambda = r + \Delta$ is the effective discount rate.

Along any trajectory that has not yet reached $\overline{X}$, survival up to date t occurs with probability $Z(q,t) = 1 - q + qe^{-\Delta t}$. Conditional on survival up to t, the posterior probability that the TP lies at 0 is $q(t) = \frac{qe^{-\Delta t}}{1-q+qe^{-\Delta t}}$. Survival is "good news" since $q(t) \le q$. The longer the system survives, the less likely it is that the hazard has been active since $t = 0$.

Two features are worth emphasizing. First, the stock $X(t)$ is a sufficient statistic for the physical state of the system, but the posterior belief depends on how much time has elapsed without catastrophe, which itself depends on past actions. Second, even though no direct signal arrives, the absence of catastrophe is informative and affects optimal behavior.

3. The Planner's Optimum

A planner chooses an action plan that depends on the full past history of actions and her beliefs. At each point in time, she maximizes continuation payoff given current beliefs, which themselves are endogenously generated by the chosen policy.[2] The planner follows a feedback rule σ^c mapping current stock into current action. This rule is sufficient to reconstruct the full history of past actions and, therefore, the belief path. Given the stock X, the time elapsed without catastrophe can be inferred as $T^c(X) = \int_0^X \frac{d\tilde{X}}{\sigma^c(\tilde{X})}$ and the associated posterior belief is $q^c(X) = q(T^c(X))$. Accordingly, the planner's value function is $\mathcal{W}^c(X) = \int_0^{+\infty} e^{-r\tau}\, Z(q^c(X),\tau)\, u(\sigma^c(X^c(\tau,X)))\, d\tau$. If instead the TP has in fact already been crossed (i.e., $\tilde{X} = 0$) but the planner behaves as if it may still lie ahead, her realized payoff is $\varphi^c(X) = \int_0^{+\infty} e^{-\lambda\tau}\, u(\sigma^c(X^c(\tau,X)))\, d\tau$. The difference $\mathcal{W}^c(X) - \varphi^c(X)$ is the value

[1] Tsur and Zemel (1995) study related problems in which crossing a threshold may be detected, altering the nature of learning and policy.

[2] Lemoine and Traeger (2014) analyze how environmental policies should optimally respond to tipping points.

of optimism: expected continuation value is higher when future payoffs are discounted at r rather than at $r+\Delta$ in the event the hazard is already active.

Proposition 1. *At the planner's optimum, the feedback rule $\sigma^c(X)$ satisfies*

$$\sigma^c(X) = \zeta + \dot{\mathcal{W}}^c(X) + \frac{\dot{q}^c(X)}{1-q^c(X)}\big(\mathcal{W}^c(X) - \varphi^c(X)\big), \qquad X \in [0, \overline{X}). \tag{3.1}$$

Proposition 1 delivers a transparent decomposition. The benchmark case $q=1$ (TP at 0 for sure) yields $\sigma^c(X)=\zeta$: once the hazard is certainly active, additional stock does not change the hazard and the optimal act is the myopic outcome. The opposite benchmark $q=0$ (TP at $\overline{X}$ for sure) reduces to a standard depletion problem. Then $\sigma^c(X) = \zeta + \dot{\mathcal{W}}^c(X) < \zeta$ for all $X < \overline{X}$. As the stock rises, future exposure becomes more imminent, and the shadow value of preserving distance increases; the planner acts more conservatively. With uncertainty ($q \in (0,1)$), the term $\dot{\mathcal{W}}^c(X)$ still reflects the irreversibility channel: a marginal increase in today's action raises the stock and reduces continuation value by bringing the system closer to $\overline{X}$. This is the standard logic in the quasi-option value and irreversible investment traditions, adapted here to a hazard that turns on discontinuously.

The additional term, $\frac{\dot{q}^c(X)}{1-q^c(X)}\big(\mathcal{W}^c(X)-\varphi^c(X)\big)$, captures pseudo-learning. Because survival lowers the posterior that $\tilde{X}=0$, $q^c(X)$ declines along the optimal path. A higher action today accelerates the accumulation of stock and therefore a given stock level is reached after less time has elapsed. With less time for survival to "prove" safety, the posterior remains more pessimistic. Conversely, a lower action slows accumulation, lengthens the survival window, and makes the planner more optimistic about not having crossed the TP—which changes her future incentives. Put differently, the policy affects how informative the absence of catastrophe is. This mechanism strengthens precaution relative to a world with a known TP. A useful implication is that the optimal policy may be more conservative early on, when little knowledge has been generated. This is exactly the region in which institutional constraints can generate the greatest welfare benefits.

4. Delegated Decision-Making

We now introduce institutional incompleteness and delegation. Decision-making is delegated over time to an infinite sequence of short-lived DMs. Each DM observes the current stock but enters office with a fixed prior and discounts the future accordingly. Specifically, every incoming DM believes that the TP lies at 0 (resp. $\overline{X}$) with probability p (resp. $1-p$), thereby ignoring the information contained in survival up to the date she takes office. Because survival is “good news" that the TP has not yet been crossed, an incoming DM is always more pessimistic than the predecessors would be at the same stock level. When $p > q$, the DM is also over-pessimistic relative to the planner in the sense that she assigns too much probability to having already crossed the TP.[3] Because all DMs begin with the same prior, they disagree about how to discount future payoffs and thus about which future actions should be taken. This generates a dynamic continuous-time game among successive office-holders (Karp 2007; Ekeland and Lazrak 2010). At date t the incumbent DM_t can commit only over an infinitesimal interval and chooses the current action $x(t)$ to maximize welfare from that point onward, given her beliefs and the current stock. At a *stock-based equilibrium* (SBE), all DMs follow a common feedback rule σ^* mapping stock into action: the incumbent finds no infinitesimal impulse deviation profitable, anticipating that all future DMs will also follow σ^*. Accordingly, the payoff for a DM observing stock level X is $\mathcal{W}^*(X) = \int_0^{+\infty} e^{-r\tau} Z(p,\tau)\, u(\sigma^*(X^*(\tau,X)))\, d\tau$, where $X^*(\tau,X)$ is the stock trajectory under σ^*. If the TP has in fact already been crossed but the current DM wrongly believes otherwise, her realized payoff is $\varphi^*(X) = \int_0^{+\infty} e^{-\lambda\tau}\, u(\sigma^*(X^*(\tau,X)))\, d\tau$.

Proposition 2. *At any SBE, the feedback rule satisfies for all $X \in [0,\overline{X})$:*

$$\sigma^*(X) = \zeta + \dot{\mathcal{W}}^*(X). \tag{4.1}$$

Comparing (4.1) with (3.1) highlights the institutional wedge: delegated DMs ignore how their actions influence future beliefs and the value of optimism. The *Pseudo-Learning Effect* disappears. Delegated policy is still affected by irreversibility (since $\dot{\mathcal{W}}^*(X) < 0$ implies $\sigma^*(X) < \zeta$), but precaution is weaker because the informational externality is not internalized. Intuitively, each DM would like successors to behave more cautiously, because aggressive future actions increase catastrophe risk when the TP lies ahead. Yet successors do not share the incumbent's more optimistic posterior formed after survival. This generates a dynamic externality: current actions affect future behavior through the evolution of the stock and through the information that survival would convey, but the latter channel is severed by institutional turnover and biased priors. The equilibrium thus tends to feature

[3] One interpretation is bounded rationality or limited institutional memory; another is political economy, where incoming DMs cater to constituencies with systematically distorted beliefs.

excessively high actions early on, when the planner would otherwise value the option of learning through survival and would shade actions down.

Proposition 3. *There exists a unique SBE, $(\mathcal{W}^*(X), \sigma^*(X))$, with boundary conditions $\mathcal{W}^*(\overline{X}) = \varphi^*(\overline{X}) = \mathcal{V}_\infty$. 1) $\mathcal{W}^*$ is always decreasing. 2) φ^* is increasing when $\overline{X}$ is sufficiently small, and first decreasing then increasing when $\overline{X}$ is sufficiently large. 3) Whether $\overline{X}$ is sufficiently small or sufficiently large, we have*

$$\dot{\mathcal{W}}^*(0) - \dot{\varphi}^*(0) < 0. \tag{4.2}$$

The inequality (4.2) concerns the marginal value of optimism at the start. When $\overline{X}$ is large, early DMs view the TP as remote and choose actions close to ζ, so both $\mathcal{W}^*$ and φ^* initially decline only slightly. As X approaches $\overline{X}$, irreversibility binds more strongly and $\mathcal{W}^*$ falls more significantly. Since φ^* is discounted more heavily, it declines less quickly, making the marginal value of optimism negative at the origin. When $\overline{X}$ is small, the trajectory hits $\overline{X}$ quickly. The *Irreversibility Effect* binds early, and $\mathcal{W}^*$ falls sharply while φ^* rises along the entire path because reaching $\overline{X}$ quickly resolves uncertainty and moves the system into the "hazard-on" regime where the continuation value is pinned by $\mathcal{V}_\infty$. In both cases, slowing stock accumulation increases the gap between how an optimistic planner and a pessimistic DM value continuation, a force that will shape the optimal constitutional constraint.

5. Constraints on Early Actions

We now evaluate simple constitutional rules that are feasible under incompleteness. An *impulse regulation* imposes an action x from the start of the trajectory for an infinitesimal duration ε. Such a rigid requirement captures the idea that constitutions cannot condition on the endogenous stock, but can impose coarse constraints at early dates. After this initial period, the economy follows the delegated trajectory, with biased DMs applying the SBE feedback rule σ^*. The planner's payoff under the impulse regulation (x, ε) is

$$\mathcal{W}^r(x,\varepsilon) = \left(\int_0^\varepsilon e^{-r\tau} Z(q,\tau)\, d\tau\right) u(x) + \int_\varepsilon^{+\infty} e^{-r\tau} Z(q,\tau)\, u(\sigma^*(X^*(\tau-\varepsilon, x\varepsilon)))\, d\tau. \tag{5.1}$$

Without any regulation, the planner receives $\mathcal{W}^r(0,0) = \mathcal{W}^*(0) + \frac{p-q}{1-p}\big(\mathcal{W}^*(0) - \varphi^*(0)\big)$. Both the planner and DM_0 discount future payoffs using survival probabilities; the only difference lies in their priors q and p. When $p > q$, the planner is more optimistic, values the future more, and therefore obtains a greater payoff than DM_0 would under her own criterion; this difference is magnified when DM_0's value of optimism $\mathcal{W}^*(0) - \varphi^*(0)$ is large.

To determine whether an impulse regulation should constrain delegated DMs—and in which direction—we expand $\mathcal{W}^r(x,\varepsilon)$ for small ε: $\mathcal{W}^r(x,\varepsilon) = \mathcal{W}^r(0,0) + \varepsilon \frac{\partial \mathcal{W}^r}{\partial \varepsilon}(x,0) + o(\varepsilon)$. Let x^r denote the action that maximizes the linear term in ε. Regulation imposes a *cap* (resp. *floor*) when $x^r < \sigma^*(0)$ (resp. $x^r > \sigma^*(0)$).

Proposition 4. *Suppose $p > q$. An impulse regulation optimally imposes a cap x^r on early actions when $\overline{X}$ is either sufficiently small or sufficiently large, where*

$$x^r = \sigma^*(0) + \frac{p-q}{1-p}\Big(\dot{\mathcal{W}}^*(0) - \dot{\varphi}^*(0)\Big) < \sigma^*(0). \tag{5.2}$$

The planner's continuation payoff following the impulse regulation is always greater than that of the biased DM_ε who inherits stock $x\varepsilon$, by $\delta(x,\varepsilon) = \frac{p-q(\varepsilon)}{1-p}\big(\mathcal{W}^*(x\varepsilon) - \varphi^*(x\varepsilon)\big)$. This difference again reflects the value of optimism. For small ε, $q(\varepsilon) \approx q < p$, so the planner remains more optimistic than DM_ε. When the delegated DM_0's marginal value of optimism at zero is negative, $\dot{\mathcal{W}}^*(0) - \dot{\varphi}^*(0) < 0$, capping early actions below $\sigma^*(0)$ increases the planner's gain. Intuitively, slowing the evolution of the stock makes survival more informative over the relevant horizon and raises the value of policy caution that future DMs would otherwise fail to adopt. Proposition 3 establishes that this situation arises robustly when $\overline{X}$ is either sufficiently large (learning is slow and valuable) or sufficiently small (irreversibility binds immediately and early mistakes are hard to undo).

These circumstances rationalize a *Precautionary Principle*: overly pessimistic delegated DMs should be forced to keep early actions low when meaningful learning and welfare-relevant information accumulation occur only over time, and when institutional turnover prevents DMs from internalizing the informational consequences of early policy choices.

Conversely, when $p < q$ (delegated DMs are over-optimistic), the same formula (5.2) implies that the planner may benefit from a floor rather than a cap: pushing early actions above $\sigma^*(0)$ is welfare-improving when delegated DMs would otherwise be too conservative relative to society's true objective. In that case, accelerating stock accumulation can reduce disagreement between the planner and delegated DMs. This scenario motivates a reverse *Precautionary Principle*, in which early minimum standards (rather than caps) correct excessive conservatism.

6. Discussion and Extensions

Several extensions preserve the core logic and can strengthen the case for early constraints. First, allowing catastrophe to impose a fixed loss at the moment it occurs (in addition to terminating payoffs) strengthens the irreversibility force and increases the welfare gain from caps that curb early excesses under delegation. Second, the result is robust to a noisy mapping between stock and information (e.g., imperfect measurement of cumulative emissions or uncertain climate response). Third, richer structures (multiple thresholds or a continuum of possible TPs) would increase the appeal of early constraints, since fully contingent plans become even less implementable. The model suggests that institutions that preserve informational continuity across political turnovers (e.g., independent agencies with long horizons, or legally mandated learning protocols) should behave more like the planner and rely less on blunt caps. Conversely, environments with high turnover and contested information should exhibit stronger demand for simple, early, rule-like constraints.

7. References

Arrow, K. and A. Fisher (1974). “Environmental Preservation, Uncertainty, and Irreversibility,” *The Quarterly Journal of Economics*, 88: 312–319.

Ekeland, I. and A. Lazrak (2010). “The Golden Rule When Preferences Are Time Inconsistent,” *Mathematical Financial Economics*, 4: 29–55.

Gardiner, S. (2006). “A Core *Precautionary Principle*,” *Journal of Political Philosophy*, 14: 33–60.

Gollier, C., B. Jullien and N. Treich (2000). “Scientific Progress and Irreversibility: An Economic Interpretation of the *Precautionary Principle*," *Journal of Public Economics*, 75: 229-253.

Guillouet, L. and D. Martimort (2025). “Acting in the Dark,” Working Paper, Toulouse School of Economics.

Holmström, B. (1984). “On the Theory of Delegation,” in M. Boyer and R. E. Kihlstrom (eds.) *Bayesian Models in Economic Theory.* Elsevier.

Karp, L. (2007). “Non-Constant Discounting in Continuous Time,” *Journal of Economic Theory*, 132: 557–568.

Lemoine, D. and C. Traeger (2014). “Watch Your Step: Optimal Policy in a Tipping Climate,” *American Economic Journal: Economic Policy*, 6: 137–166.

Tsur, Y. and A. Zemel (1995). “Uncertainty and Irreversibility in Groundwater Resource Management,” *Journal of Environmental Economics and Management*, 29: 149–161.

8. Appendix

Proof of Proposition 1. Consider an impulse deviation in which the planner chooses action x for an infinitesimal time ε, raising the stock by εx, after which σ^c is resumed. The planner’s deviation payoff is

$$\hat{\mathcal{W}}^c(x,\varepsilon,X) = \left(\int_0^{\varepsilon} Z(q^c(X),\tau)e^{-r\tau}\,d\tau\right)u(x) + \int_{\varepsilon}^{+\infty} e^{-r\tau} Z(q^c(X),\tau)u(\sigma^c(X^c(\tau-\varepsilon, X+\varepsilon x)))\,d\tau. \tag{8.1}$$

Observe that $q^c(X) = \frac{qe^{-\Delta T^c(X)}}{1-q+qe^{-\Delta T^c(X)}}$ evolves according to

$$\sigma^c(X)\dot{q}^c(X) = -\Delta q^c(X)(1-q^c(X)), \text{ with } q^c(0) = q. \tag{8.2}$$

Observe now that the survival probability satisfies the simple rule

$$Z(q^c(X),\tau+\varepsilon) = Z(q^c(X),\tau) - q^c(X)\left(1-e^{-\Delta\varepsilon}\right)e^{-\Delta\tau}.$$

Equipped with this condition and changing variables, we rewrite the benefit of a deviation as

$$\hat{\mathcal{W}}^c(x,\varepsilon,X) = \int_0^{\varepsilon} Z(q^c(X),\tau)e^{-r\tau}u(x)d\tau$$

$$+e^{-r\varepsilon}\Big(\mathcal{W}^c(X+x\varepsilon) + \int_0^{+\infty} e^{-r\tau}\left(Z(q^c(X),\tau) - Z(q^c(X+x\varepsilon),\tau)\right)u(\sigma^c(X^c(\tau,X+x\varepsilon)))d\tau$$

$$-q^c(X)\left(1-e^{-\Delta\varepsilon}\right)\varphi^c(X+x\varepsilon)\Big).$$

Taking a first-order Taylor approximation in ε yields

$$\hat{\mathcal{W}}^c(x,\varepsilon,X) = \mathcal{W}^c(X) + \varepsilon\frac{\partial\hat{\mathcal{W}}^c}{\partial\varepsilon}(x,0,X) + o(\varepsilon).$$

At an optimum $(\mathcal{W}^c(X),\sigma^c(X))$ of the planner's problem, any impulse deviation must be weakly dominated. Hence, we must have

$$0 = \max_{x\in\mathbb{R}_+}\frac{\partial\hat{\mathcal{W}}^c}{\partial\varepsilon}(x,0,X)$$

or

$$0 = \max_{x\in\mathbb{R}_+}\left(-r\mathcal{W}^c(X) - \Delta q^c(X)\varphi^c(X) + u(x) + x\left(\dot{\mathcal{W}}^c(X) + \dot{q}^c(X)\frac{\mathcal{W}^c(X)-\varphi^c(X)}{1-q^c(X)}\right)\right).$$

The necessary conditions for optimality are then

$$r\mathcal{W}^c(X) + \Delta q^c(X)\varphi^c(X) = \max_{x\in\mathbb{R}_+} u(x) + x\left(\dot{\mathcal{W}}^c(X) + \frac{\dot{q}^c(X)}{1-q^c(X)}\left(\mathcal{W}^c(X)-\varphi^c(X)\right)\right), \tag{8.3}$$

$$\sigma^c(X) \in \arg\max_{x\in\mathbb{R}_+} u(x) + x\left(\dot{\mathcal{W}}^c(X) + \frac{\dot{q}^c(X)}{1-q^c(X)}\left(\mathcal{W}^c(X)-\varphi^c(X)\right)\right). \tag{8.4}$$

Given strict concavity of the maximand above, an interior solution is given by (3.1). Inserting (3.1) into (8.3) and using (8.2) also yields

$$r\mathcal{W}^c(X) + q^c(X)\Delta\varphi^c(X) = u(\sigma^c(X)) + \sigma^c(X)\dot{\mathcal{W}}^c(X) - q^c(X)\Delta(\mathcal{W}^c(X)-\varphi^c(X)) \tag{8.5}$$

with the boundary condition $\mathcal{W}^c(\overline{X}) = \mathcal{V}_\infty$. □

Proof of Proposition 2. We define the deviation payoff $\hat{\mathcal{W}}^*(x,\varepsilon,X)$ as

$$\hat{\mathcal{W}}^*(x,\varepsilon,X)=\left(\int_0^{\varepsilon}e^{-r\tau}Z(p,\tau)d\tau\right)u(x)+\int_{\varepsilon}^{+\infty}e^{-r\tau}Z(p,\tau)u(\sigma^*(X^*(\tau-\varepsilon,X+x\varepsilon)))d\tau.$$

Changing variables, we rewrite this expression as

$$\hat{\mathcal{W}}^*(x,\varepsilon,X)=\left(\int_0^{\varepsilon}e^{-r\tau}Z(p,\tau)d\tau\right)u(x)+e^{-r\varepsilon}\int_0^{+\infty}e^{-r\tau}Z(p,\tau+\varepsilon)u(\sigma^*(X^*(\tau,X+x\varepsilon)))d\tau.$$

Using the identity $Z(p,\varepsilon+s)=Z(p,\varepsilon)Z(p(\varepsilon),s)$ we have

$$\hat{\mathcal{W}}^*(x,\varepsilon,X)=\left(\int_0^{\varepsilon}e^{-r\tau}Z(p,\tau)d\tau\right)u(x)+e^{-r\varepsilon}Z(p,\varepsilon)\Big(\int_0^{+\infty}e^{-r\tau}Z(p(\varepsilon),\tau)u(\sigma^*(X^*(\tau,X+x\varepsilon)))d\tau \tag{8.6}$$

where $p(\varepsilon)=\frac{pe^{-\Delta\varepsilon}}{1-p+pe^{-\Delta\varepsilon}}$ and, for ε small enough, $p(\varepsilon)=p-\Delta p(1-p)\varepsilon+o(\varepsilon)$. The deviation payoff can actually be expressed as

$$\hat{\mathcal{W}}^*(x,\varepsilon,X)=\left(\int_0^{\varepsilon}e^{-r\tau}Z(p,\tau)d\tau\right)u(x)$$

$$+e^{-r\varepsilon}Z(p,\varepsilon)\left(\mathcal{W}^*(X+x\varepsilon)-\frac{p(\varepsilon)-p}{1-p}\left(\mathcal{W}^*(X+x\varepsilon)-\varphi^*(X+x\varepsilon)\right)\right).$$

Assuming $\mathcal{W}^*$ differentiable, a first-order Taylor approximation in ε yields

$$\hat{\mathcal{W}}^*(x,\varepsilon,X)=\mathcal{W}^*(X)+\varepsilon\left(-(r+p\Delta)\mathcal{W}^*(X)+u(x)+x\dot{\mathcal{W}}^*(X)-\frac{\dot{p}(0)}{1-p}\left(\mathcal{W}^*(X)-\varphi^*(X)\right)\right)+o(\varepsilon)$$

or, using $\dot{p}(0)=-\Delta p(1-p)$,

$$\hat{\mathcal{W}}^*(x,\varepsilon,X)=\mathcal{W}^*(X)+\varepsilon\left(-r\mathcal{W}^*(X)-p\Delta\varphi^*(X)+u(x)+x\dot{\mathcal{W}}^*(X)\right)+o(\varepsilon)$$

We rewrite this Taylor expansion as

$$\hat{\mathcal{W}}^*(x,\varepsilon,X)=\mathcal{W}^*(X)+\varepsilon\frac{\partial\hat{\mathcal{W}}^*}{\partial\varepsilon}(x,0,X)+o(\varepsilon).$$

At a *SBE* $(\mathcal{W}^*(X),\sigma^*(X))$, any impulse deviation must be weakly dominated for the current decision-maker. We should thus have

$$0=\max_{x\in\mathbb{R}_+}\frac{\partial\hat{\mathcal{W}}^*}{\partial\varepsilon}(x,0,X)=\max_{x\in\mathbb{R}_+}\left(-r\mathcal{W}^*(X)-\Delta p\varphi^*(X)+u(x)+x\dot{\mathcal{W}}^*(X)\right). \tag{8.7}$$

Given strict concavity of the maximand above, an interior solution is given by (4.1). Inserting into the maximand of (8.7) yields

$$r\mathcal{W}^*(X)+p\Delta\varphi^*(X)=u(\sigma^*(X))+\sigma^*(X)\dot{\mathcal{W}}^*(X). \tag{8.8}$$

□

Proof of Proposition 3. Inserting (4.1) into (8.8), we obtain

$$r\mathcal{W}^*(X)+p\Delta\varphi^*(X)=\frac{(\zeta+\dot{\mathcal{W}}^*(X))^2}{2}.$$

Taking the highest root yields

$$\begin{cases} \dot{\mathcal{W}}^*(X) = -\zeta + \sqrt{2\left(r\mathcal{W}^*(X) + p\Delta\varphi^*(X)\right)} & \text{if } X \in [0, \overline{X}), \\ \mathcal{W}^*(X) = \mathcal{V}_\infty & \text{if } X \geq \overline{X}. \end{cases} \tag{8.9}$$

Finally, we obtain

$$\sigma^*(X) = \begin{cases} \sqrt{2\left(r\mathcal{W}^*(X) + p\Delta\varphi^*(X)\right)} & \text{if } X \in [0, \overline{X}), \\ \zeta & \text{if } X \geq \overline{X}. \end{cases} \tag{8.10}$$

Backwards and Forwards System of Ordinary Differential Equations. We now form a system of *ODE*s for $\varphi^*(X)$ and $\mathcal{W}^*(X)$. Studying its properties provides existence and uniqueness of a solution. Observe that, for $X \in [0, \overline{X})$, φ^* is differentiable when σ^* is, with

$$\dot{\varphi}^*(X) = \int_0^{+\infty} e^{-\lambda\tau} u'(\sigma^*(X^*(\tau, X))\dot{\sigma}^*(X^*(\tau, X))\frac{\partial X^*}{\partial X}(\tau, X)d\tau$$

where $\frac{\partial X^*}{\partial X}(\tau, X) = \frac{\sigma^*(X^*(\tau,X))}{\sigma^*(X)}$. Manipulating and integrating by parts yields the *ODE* for φ^*:

$$\sigma^*(X)\dot{\varphi}^*(X) = \lambda\varphi^*(X) - u(\sigma^*(X)). \tag{8.11}$$

Using (8.10), we rewrite this condition as

$$\dot{\varphi}^*(X) = \frac{\lambda\varphi^*(X) - u(\sigma^*(X))}{\sqrt{2\left(r\mathcal{W}^*(X) + p\Delta\varphi^*(X)\right)}}. \tag{8.12}$$

Together with the definition of σ^* in (8.10), (8.9) and (8.12) define a system of *ODE*s whose trajectories may be represented in the $(\varphi, \mathcal{W})$ plan. The terminal conditions for this system are

$$\mathcal{W}^*(\overline{X}) = \varphi^*(\overline{X}) = \mathcal{V}_\infty \tag{8.13}$$

where equalities follow from the fact that uncertainty is resolved for all decision-makers once the stock reaches $\overline{X}$.

Existence and Unicity. The locus L_1 of points such that $\dot{\varphi}^*(X) = 0$ is defined as $\lambda\varphi = u(\sigma^*)$ or

$$\sigma^* = \zeta \pm \sqrt{2\lambda(\mathcal{V}_\infty - \varphi)}.$$

Using the expression of σ^* from (8.10), we find that L_1 is made of two branches:

$$2(r\mathcal{W} + p\Delta\varphi) = \left(\zeta \pm \sqrt{2\lambda(\mathcal{V}_\infty - \varphi)}\right)^2. \tag{8.14}$$

It turns out that $\dot{\varphi}^*(X) > 0$ (resp. < 0) when the trajectory lies below (resp. above) this locus. Observe also that $\mathcal{W}^*(\overline{X}) = \varphi^*(\overline{X}) = \mathcal{V}_\infty$ lies below L_1 since $2(r + p\Delta)\mathcal{V}_\infty < \zeta^2 = 2\lambda\mathcal{V}_\infty$.

Similarly, the locus L_2 of points such that $\dot{\mathcal{W}}^*(X) = 0$ is, from (8.9),

$$2(r\mathcal{W} + p\Delta\varphi) = \zeta^2. \tag{8.15}$$

We have $\dot{\mathcal{W}}^*(X) < 0$ (resp. > 0) when the trajectory lies below (resp. above) this locus.

Observe that L_1 and L_2 intersect at

$$\mathcal{W}_\infty = \left(1 + \frac{1-p}{r}\Delta\right)\mathcal{V}_\infty \text{ and } \varphi_\infty = \mathcal{V}_\infty. \tag{8.16}$$

Consider now the backward trajectories of the system (8.9)-(8.12). Accordingly, let $\mathcal{W}_*(Y) = \mathcal{W}^*(\overline{X} - Y)$ and $\varphi_*(Y) = \varphi^*(\overline{X} - Y)$ be the solution to the backward system; thus considering values of the stock $X = \overline{X} - Y \leq \overline{X}$ for $Y \geq 0$ and viewing $Y = 0$ as the initial point of those backward trajectories. For this backward system, the initial conditions are the same as the terminal conditions (8.13) of the forward system, i.e.,

$$\mathcal{W}_*(0) = \varphi_*(0) = \mathcal{V}_\infty. \tag{8.17}$$

Using (8.9)-(8.12), we can write this backward system as

$$\dot{\mathcal{W}}_*(Y) = \zeta - \sqrt{2\left(r\mathcal{W}^*(Y) + p\Delta\varphi_*(Y)\right)} \tag{8.18}$$

$$\dot{\varphi}_*(Y) = \frac{u(\sigma_*(Y)) - \lambda\varphi_*(Y)}{\sqrt{2\left(r\mathcal{W}_*(Y) + p\Delta\varphi_*(Y)\right)}} \tag{8.19}$$

where, from (8.10), we define

$$\sigma_*(Y) = \sigma^*(\overline{X} - Y) = \sqrt{2\left(r\mathcal{W}_*(Y) + p\Delta\varphi_*(Y)\right)}. \tag{8.20}$$

Trajectories of the backward system, if they converge when Y goes to infinity (or equivalently, considering $\overline{X}$ going to infinity in the forward system), do so towards the stationary point $(\varphi_\infty, \mathcal{W}_\infty)$ where L_1 and L_2 intersect.

Observe that the action $\sigma^*(\overline{X}^-)$ is given by (8.10), i.e.,

$$\sigma^*(\overline{X}^-) = \zeta\sqrt{1 - \frac{1-p}{\lambda}\Delta} < \zeta. \tag{8.21}$$

The backward system (8.18)-(8.19) satisfies a Lipschitz condition at $Y = 0$ with

$$\dot{\mathcal{W}}_*(0) = \zeta\left(1 - \sqrt{1 - \frac{1-p}{\lambda}\Delta}\right) > 0 \text{ and } \dot{\varphi}_*(0) = \frac{u(\sigma^*(\overline{X}^-)) - \lambda\mathcal{V}_\infty}{\sigma^*(\overline{X}^-)} < 0 \tag{8.22}$$

where the last inequality follows from the fact that $\lambda\mathcal{V}_\infty = \frac{\zeta^2}{2}$ is maximizing flow payoff. From Cauchy-Lipschitz Theorem, locally in a right-neighborhood of $Y = 0$, there exists a unique solution to the backward system (8.18)-(8.19) together with the initial condition (8.17). This solution can be extended over the whole interval $\left[0, \overline{X}\right]$. Observe that the initial conditions (8.17) of the backward system are the same for all possible values of $\overline{X}$ and that the derivatives $\dot{\mathcal{W}}_*(0)$ and $\dot{\mathcal{W}}_*(0)$ in (8.22) are also independent of $\overline{X}$ since, from (8.21), $\sigma^*(\overline{X}^-) = \sigma_*(0^+)$ is itself so. For all values of $\overline{X}$, solutions of the forward system thus lie on the same one-dimensional locus $\mathcal{L}$ in the $(\varphi, \mathcal{W})$ space. Hence, fixing a particular value $\overline{X}$ amounts in fact to choosing a point along the locus $\mathcal{L}$ that corresponds to the initial values $(\mathcal{W}^*(0), \varphi^*(0))$ for the forward system.

Long-run behavior of the backward system. Because $(\varphi_*(0), \mathcal{W}_*(0))$ lies on the diagonal, below both L_1 and L_2, the trajectory of the backward system starting from $\varphi_*(0) = \mathcal{W}_*(0) = \mathcal{V}_\infty$ is such that $\mathcal{W}_*$ is increasing while φ_* is first decreasing (and always so for all values of $Y \in \left[0, \overline{X}\right]$ when $\overline{X}$ is small) before it eventually reaches L_1 and is increasing afterwards (case where $\overline{X}$ is large). Observe that φ_*, once it has already crossed L_1 at some Y_1 and lies above L_1 cannot cross it one more time at some finite $Y_2 > Y_1$ since when it crosses L_1 , it must cross it necessarily with $\dot{\varphi}_*(Y_2) = 0$ but this cannot be for a trajectory coming from above L_1 unless it is at point where L_1 admits a vertical tangent; and the only such point corresponds to $\varphi = \mathcal{V}_\infty$, which is only reached for $Y_2 = +\infty$.

For the forward system, we have the reverse pattern. First, $\mathcal{W}^*$ is always decreasing starting from $\mathcal{W}^*(0)$ and going to $\mathcal{W}^*(\overline{X}) = \mathcal{W}_*(0) = \mathcal{V}_\infty$. Hence Item 1. is proved.

Second, when $\overline{X}$ is large enough, $(\varphi^*(0), \mathcal{W}^*(0))$ lies above L_1 and thus φ^* is first decreasing starting from $\varphi^*(0)$ before it necessarily reaches L_1 and becomes increasing afterwards towards $\varphi^*(\overline{X}) = \mathcal{V}_\infty$. When $\overline{X}$ is small enough, $(\varphi^*(0), \mathcal{W}^*(0))$ lies below L_1 and thus φ^* is always increasing starting from $\varphi^*(0)$. Hence, Item 2. is proved.

That $\mathcal{W}^*$ is always decreasing in the forward system, also implies that $\sigma^*(X) < \zeta$, for $X \in [0, \overline{X})$. When $\overline{X}$ increases towards $+\infty$, $(\varphi^*(0), \mathcal{W}^*(0))$ converges towards $(\varphi_\infty, \mathcal{W}_\infty)$ which is the stationary point of the backward system. Moreover, by definition, we have $\mathcal{W}^*(X) \leq \lambda \mathcal{V}_\infty \int_0^{+\infty} e^{-r\tau} \left(1 - p + pe^{-\Delta\tau}\right) d\tau = \mathcal{W}_\infty$. Hence, a trajectory of the backward system crossing L_2 stays below the horizontal line $\mathcal{W} = \mathcal{W}_\infty$ when converging towards $(\varphi_\infty, \mathcal{W}_\infty)$. Moreover, in the neighborhood of $(\varphi_\infty, \mathcal{W}_\infty)$, the backward system can be linearized as

$$\dot{\mathcal{W}}_*(Y) = -\frac{r}{\zeta}(\mathcal{W}_*(Y) - \mathcal{W}_\infty) - \frac{p\Delta}{\zeta}(\varphi_*(Y) - \varphi_\infty) \tag{8.23}$$

$$\dot{\varphi_*}(Y) = -\frac{\lambda}{\zeta}(\varphi_*(Y) - \varphi_\infty). \tag{8.24}$$

This linear system has two negative eigenvalues $\chi_1 = -\frac{\lambda}{\zeta}$ and $\chi_2 = -\frac{r}{\zeta}$ and $(\varphi_\infty, \mathcal{W}_\infty)$ is thus a stable node. The solutions to the linearized system are of the form

$$\mathcal{W}_*(Y) - \mathcal{W}_\infty = p\phi_0 e^{-\frac{\lambda}{\zeta}Y} + w_0 e^{-\frac{r}{\zeta}Y} \tag{8.25}$$

$$\varphi_*(Y) - \varphi_\infty = \phi_0 e^{-\frac{\lambda}{\zeta}Y} \tag{8.26}$$

for (ϕ_0, w_0) arbitrary constants. We can eliminate Y from those two equations to get

$$\mathcal{W}_*(Y) - \mathcal{W}_\infty = p(\varphi_*(Y) - \varphi_\infty) + w_0 \left(\frac{\varphi_*(Y) - \varphi_\infty}{\phi_0}\right)^{\frac{r}{\lambda}} \tag{8.27}$$

When $Y \to +\infty$, $\varphi_*(Y) \to \varphi_\infty$ and $\mathcal{W}_*(Y) \to \mathcal{W}_\infty$ and the behavior of these solutions is thus given by the dominant terms corresponding to the highest eigenvalue χ_2, i.e., $\mathcal{W}_*(Y) - \mathcal{W}_\infty \sim w_0 \left(\frac{\varphi_*(Y) - \varphi_\infty}{\phi_0}\right)^{\frac{r}{\lambda}}$. In the $(\varphi, \mathcal{W})$ spaces these solutions are thus tangent to the vertical line $\varphi = \varphi_\infty = \mathcal{V}_\infty$ but remain above loci L_1 since we know that $\varphi_*(Y)$ do not cross L_1 twice.

For $\overline{X}$ large enough, $(\varphi^*(0), \mathcal{W}^*(0))$ thus lies close to the vertical line $\varphi = \mathcal{V}_\infty$ with $\mathcal{W}^*(0) < \mathcal{W}_\infty$ and $\varphi^*(0) < \mathcal{V}_\infty$ but $\varphi^*(0) - \mathcal{V}_\infty$ is of a much lower order of magnitude than $\mathcal{W}^*(0) - \mathcal{W}_\infty$. Hence, we must have $\mathcal{W}^*(0) < \left(1 + \frac{1-p}{r}\Delta\right)\varphi^*(0)$ for $\overline{X}$ large enough.

For $\overline{X}$ small enough, $(\varphi^*(0), \mathcal{W}^*(0))$ lies close to $(\mathcal{V}_\infty, \mathcal{V}_\infty)$ and again $\mathcal{W}^*(0) < \left(1 + \frac{1-p}{r}\Delta\right)\varphi^*(0)$. Using (8.8) and (8.12) at $X = 0$, we now compute

$$\sigma^*(0)\left(\dot{\mathcal{W}}^*(0) - \dot{\varphi}^*(0)\right) = r\left(\mathcal{W}^*(0) - \left(1 + \frac{1-p}{r}\Delta\right)\varphi^*(0)\right). \tag{8.28}$$

Hence, whether $\overline{X}$ is large or small enough, we have $\dot{\mathcal{W}}^*(0) - \dot{\varphi}^*(0) < 0$ which proves Item 3.

Phase Diagram Illustration. This section illustrates the phase–plane analysis of the backward system defined in equations (8.25)-(8.26).We fix parameters

$$\zeta = 1, \quad \lambda = 1.01, \quad \Delta = 1, \quad r = 0.01, \quad p = 0.5, \quad \mathcal{V}_\infty = \frac{1}{2.02}.$$

Figure 1 displays the loci

$$L_1: \quad 2(r\mathcal{W} + p\Delta\varphi) = \left(\zeta - \sqrt{2\lambda(\mathcal{V}_\infty - \varphi)}\right)^2, \qquad L_2: \quad 2(r\mathcal{W} + p\Delta\varphi) = \zeta^2,$$

together with backward trajectories $(\varphi_*(Y), \mathcal{W}_*(Y))$ for $\overline{X} \in \{1, 10, 100\}$. All trajectories start from $(\mathcal{V}_\infty, \mathcal{V}_\infty)$ and initially move to the right. When $\overline{X}$ is sufficiently large, the trajectory crosses the lower branch of L_1 once and thereafter remains strictly above it.

The horizontal axis on Figure 1 is zoomed around $\varphi = \mathcal{V}_\infty$, while the vertical axis is extended to display convergence toward the stationary point $(\varphi_\infty, \mathcal{W}_\infty)$. All trajectories start on the right of the lower branch of L_1, cross it and converge monotonically without further crossings.

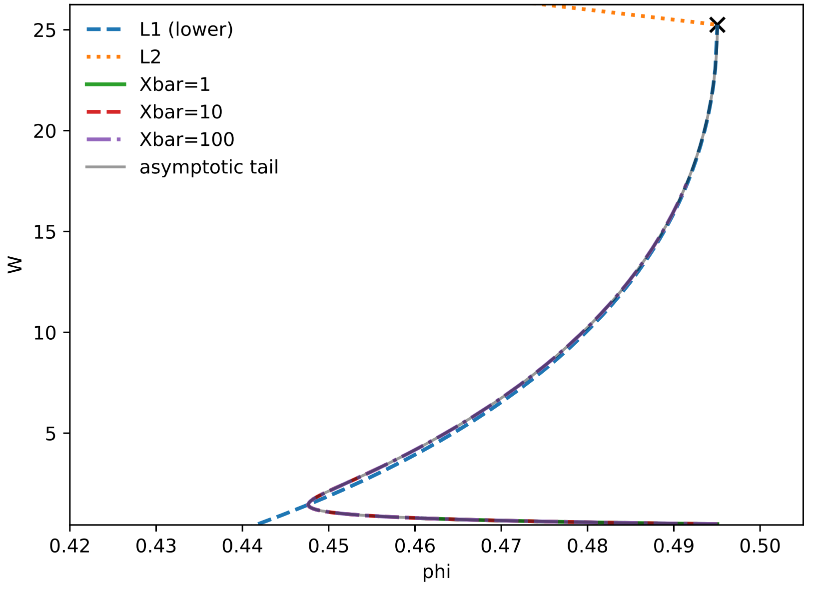

FIGURE 1. Backward trajectories $(\varphi_*(Y), \mathcal{W}_*(Y))$ for $\overline{X} \in \{1, 10, 100\}$ together with the lower branch of L_1 and locus L_2.

□

PROOF OF PROPOSITION 4. Using the identity $Z(q, \varepsilon + s) = Z(q, \varepsilon)Z(q(\varepsilon), s)$, the planner's expected payoff in (5.1) writes also as

$$\mathcal{W}^r(x,\varepsilon) = \left(\int_0^\varepsilon e^{-r\tau} Z(q,\tau)d\tau\right) u(x) + e^{-r\varepsilon} Z(q,\varepsilon) \int_0^{+\infty} e^{-r\tau} Z(q(\varepsilon),\tau) u(\sigma^*(X^*(\tau, x\varepsilon)))d\tau \tag{8.29}$$

First, we rewrite (8.29) and obtain

$$\mathcal{W}^r(x,\varepsilon) = \left(\int_0^\varepsilon e^{-r\tau} Z(q,\tau)d\tau\right) u(x) + e^{-r\varepsilon} Z(q,\varepsilon) \left(\mathcal{W}^*(x\varepsilon) + \frac{p - q(\varepsilon)}{1-p}\left(\mathcal{W}^*(x\varepsilon) - \varphi^*(x\varepsilon)\right)\right). \tag{8.30}$$

We observe that $q(\varepsilon)$ admits the following Taylor approximation when ε is close to zero $q(\varepsilon) = q - \Delta q(1-q)\varepsilon + o(\varepsilon)$. We obtain a first-order Taylor approximation of $\mathcal{W}^r(x,\varepsilon)$ for ε close to zero:

$$\mathcal{W}^r(x,\varepsilon) = \mathcal{W}^r(0,0) + \varepsilon\left(u(x) + x\left(\dot{\mathcal{W}}^*(0) + \frac{p-q}{1-p}\left(\dot{\mathcal{W}}^*(0) - \dot{\varphi}^*(0)\right)\right)\right) \tag{8.31}$$

$$-\varepsilon\left((r+q\Delta)\mathcal{W}^r(0,0) - \frac{q(1-q)\Delta}{1-p}\left(\mathcal{W}^*(0) - \varphi^*(0)\right)\right) + o(\varepsilon).$$

Using $\mathcal{W}^r(0,0) = \mathcal{W}^*(0) - \frac{q-p}{1-p}\left(\mathcal{W}^*(0) - \varphi^*(0)\right)$ and (8.8) for $X=0$, namely $r\mathcal{W}^*(0) + p\Delta\varphi^*(0) = u(\sigma^*(0)) + \sigma^*(0)\dot{\mathcal{W}}^*(0)$, and (8.28), we obtain

$$\mathcal{W}^r(x,\varepsilon) = \mathcal{W}^r(0,0) + \varepsilon\left(u(x) - u(\sigma^*(0)) + (x-\sigma^*(0))\left(\dot{\mathcal{W}}^*(0) + \frac{p-q}{1-p}\left(\dot{\mathcal{W}}^*(0) - \dot{\varphi}^*(0)\right)\right)\right) + o(\varepsilon). \tag{8.32}$$

It follows from (8.32) that the optimal impulse regulation entails

$$x^r \in \arg\max_{x\in\mathbb{R}_+} u(x) + x\left(\dot{\mathcal{W}}^*(0) + \frac{p-q}{1-p}\left(\dot{\mathcal{W}}^*(0) - \dot{\varphi}^*(0)\right)\right).$$

An interior solution thus satisfies (5.2). For $p > q$, we have

$$x^r < \sigma^*(0) \Leftrightarrow \dot{\mathcal{W}}^*(0) - \dot{\varphi}^*(0) < 0$$

Using (4.2), we conclude that, for $\overline{X}$ large (resp. small) enough and $p > q$, a cap on actions is optimal and (5.2) immediately follows. □