

An Admissibility Framework for Risk Generating Technologies

Safety Fixed Constraints and Irreversibility

Kutay Cingiz Justus Wesseler

Wageningen University

ICABR Conference
Ravello

July 8, 2026

Outline

- 1 Motivation
- 2 Maximum Incremental Social Tolerable Irreversible Cost (MISTIC)
- 3 Setup
- 4 Safety-Fixed Constraint
- 5 Integrated Admissibility
- 6 Results
- 7 References

Motivation I: Why a new framework?

- Regulatory approval is rarely a simple *approve vs. reject* problem.
- In practice, decision-makers (regulator, policy maker) choose: approval time and safeguard.
- Two distinct questions arise:
 - ① Is the technology *safe enough* under a safeguard policy?
 - ② Is approval *socially tolerable now*, given irreversibility and the value of waiting?
- The paper's contribution is to combine these two decision margins into one admissibility rule.

Motivation II: Why risk control is not enough?

- Safety regulation often focuses on controlling the probability that risk exceeds a standard.
- But meeting a probability constraint does *not* automatically imply that approval is desirable.
- Some safeguard bundles may satisfy the safety target while imposing irreversible burdens.
- Hence:

Safety feasibility $\neq$ overall admissibility.

Motivation III: Why timing matters?

- Delay reflects the option value of waiting.
- Waiting may be optimal when new information arrives in the future.
- The key policy object is therefore not just “is approval feasible?”
- *When* approval becomes justified?

Option Value of Waiting

- Approving a technology under irreversibility, uncertainty and flexibility.
- Following Wesseler et al. (2007) model, let:
 - W_T : social incremental reversible net benefits,
 - J_T : incremental irreversible costs,
 - B_T : incremental irreversible benefits.
- Welfare from approval at time T :

$$\max F(W) = \max \mathbb{E} \left[(W_T - (J_T - B_T)) e^{-rT} \right].$$

- $T = 0$: immediate approval,
 - $T > 0$: postponement to preserve flexibility and learn.
- The option value of waiting creates a hurdle factor (Demont et al. (2004))

$$H = \frac{\beta}{\beta - 1} > 1, \quad \beta > 1.$$

- The approval rule implies a threshold on tolerable irreversible costs.
- Rearranging the hurdle condition yields the Maximum Incremental Social Tolerable Irreversible Cost (MISTIC):

$$J_T^* = B_T + \frac{W_T}{H}.$$

- Interpretation:
 - J_T^* is the largest irreversible burden society should tolerate at approval time T , when irreversibility, uncertainty, and flexibility are present.
- For any given policy $x \in \mathbb{R}_+^n$, an approval at time $T \geq 0$ is tolerable only if irreversible costs $J_T(x)$ do not exceed this threshold:

$$J_T(x) \leq J_T^*(x).$$

Table: SIRBs, SIIBs, hurdle rates, and MISTICs for *Bt* maize for grain maize production on average per year for the EU-15 at 10.5% discount rate with and without CAP subsidies (in 2005 prices)

Country	SIRB		SIIB		Hurdle rate	MISTIC			
	Mio. €	€/ha	Mio. €	€/ha		Mio. €	€/ha	€/capita	€/farmhl.
<i>With CAP</i>									
France	61.90	203.80	0.24	0.81	1.14	54.31	178.81	0.90	467.12
Greece	11.76	280.34	0.04	1.03	1.79	6.60	157.34	0.60	73.75
Italy	59.90	299.29	0.19	0.98	1.23	48.87	244.16	0.84	214.27
Portugal	4.48	194.31	0.02	1.08	1.21	3.73	161.56	0.36	30.84
Spain	27.24	340.13	0.07	0.90	1.28	21.42	268.73	0.51	257.86
<i>Without CAP</i>									
France	35.89	117.89	0.24	0.81	1.16	31.09	102.11	0.52	267.40
Greece	7.11	169.32	0.04	1.03	2.50	2.89	68.75	0.26	32.25
Italy	37.25	186.71	0.19	0.98	1.31	28.55	143.13	0.49	125.20
Portugal	2.00	87.19	0.02	1.08	1.19	1.71	74.48	0.16	14.11
Spain	17.86	222.52	0.07	0.90	1.03	17.47	217.66	0.41	210.31

- The regulator chooses a safeguard policy vector

$$x = (x_I, x_F, x_S) \in \mathbb{R}_+^3.$$

- Interpretation of safeguards:
 - x_I : introduction controls,
 - x_F : dissemination / exposure controls,
 - x_S : susceptibility mitigation.
- The risk is defined as:

$$\mathcal{R}(x) = I(x_I) F(x_F) S(x_S).$$

- Following the literature (Crouch and Wilson (1981), Lichtenberg and Zilberman (1988) and Lichtenberg (2006b)):
- Let $\mathcal{I} = \{1, \dots, n\}$ denote the set of risk-generating mechanisms and let

$$x = (x_i)_{i \in \mathcal{I}} \in \mathbb{R}_+^n$$

be the vector of safeguard intensities. For each mechanism $i \in \mathcal{I}$, let $Y_i(x_i) > 0$ denote the corresponding risk-generating factor under safeguard intensity x_i . Total incremental risk is modelled as the product of these n factors:

$$R(x) = \prod_{i \in \mathcal{I}} Y_i(x_i).$$

Safety Fixed Requirement

- Let $M > 0$ be the nominal risk standard.
- Let $\alpha \in [0, 1]$ denote the regulatory margin of safety.
- The safety-fixed constraint is (Kataoka (1963)):

$$\mathbb{P}(\mathcal{R}(x) > M) \leq 1 - \alpha.$$

- Larger α means stricter precaution.

Log Transformation

- For analytical convenience, define log variables:

$$r(x) = \ln R(x), \quad y_i(x_i) = \ln Y_i(x_i) \quad \text{then} \quad r(x) = \sum_{i \in \mathcal{I}} y_i(x_i).$$

Let $M > 0$ be the nominal risk standard (maximum acceptable risk level) and let $m = \ln M$. Let $\alpha \in (0, 1)$ denote the regulatory margin of safety, interpreted as the probability that the standard is achieved. The safety-fixed constraint (Kataoka (1963)) is:

$$\Pr(R(x) > M) \leq 1 - \alpha \quad \text{equivalently} \quad \Pr(r(x) \leq m) \geq \alpha.$$

Assumption: Lognormal Risk Factors

Assumption

The vector of log-risk factors $y(x) = (y_i(x_i))_{i \in \mathcal{I}}$ is jointly normal for each x , i.e.

$$y(x) \sim \mathcal{N}(\mu(x), \Sigma(x)),$$

with finite mean vector $\mu(x) \in \mathbb{R}^n$ and covariance matrix $\Sigma(x) \in \mathbb{R}^{n \times n}$.

Constraint in Lognormal Risk Factors

- Let z_α denote the standardized cutoff exceed with probability $1 - \alpha$
- Then

$$\mathbb{P}(\mathcal{R}(x) > M) \leq 1 - \alpha \iff \mu_r(x) + z_\alpha \sigma_r(x) \leq m.$$

Define the safety-feasible set:

$$\mathcal{X}(\alpha, m) = \left\{ x \in \mathbb{R}_+^3 : \mu_r(x) + z_\alpha \sigma_r(x) \leq m \right\}.$$

- A safeguard policy is feasible if it keeps the risk below the regulatory threshold.
- z_α can be defined for any distribution as the standardized α -quantile. But only under normality it is the familiar standard-normal cutoff independent of the particular risk process. For other distributions, the cutoff must be computed from the relevant distribution.

To keep the admissibility framework independent of a specific distributional assumption, we define the safety-fixed condition in quantile form. Let

$$Q_\alpha[R(x)] = \inf \{z \in \mathbb{R} : \mathbb{P}(R(x) \leq z) \geq \alpha\}$$

denote the α -quantile of risk $R(x)$. Then the safety-feasible set can be written as

$$\mathcal{X}(\alpha, m) = \left\{x \in \mathbb{R}_+^n : Q_\alpha[R(x)] \leq M\right\}. \quad (1)$$

The safety-fixed constraint requires that the risk level remains below the nominal risk standard M with probability at least α . Equivalently, the α -quantile of risk must not exceed M . Hence, α determines the required margin of safety, while M defines the maximum acceptable risk level.




Figure: Safety condition: $Q_\alpha[R(x)] \leq M$

Integrated Admissibility Rule

For any given (α, m) , an approval at time T under safeguard policy x is admissible if both hold:

① **Safety-fixed compliance:**

$$x \in \mathcal{X}(\alpha, m).$$

② **Irreversibility tolerability (MISTIC):**

$$J_T(x) \leq J_T^*(x).$$

where $\bar{x}$ is the status quo safeguard policy. Thus the admissible set is

$$\mathcal{A}_T(\alpha, m) = \{x \in \mathcal{X}(\alpha, m) : J_T(x) \leq J_T^*(x)\}.$$

Earliest Admissible approval Time

The earliest admissible approval time is defined as

$$T^*(\alpha, m) = \inf \{ T \geq 0 : \mathcal{A}_T(\alpha, m) \neq \emptyset \}.$$

- If $\mathcal{A}_T(\alpha, m) = \emptyset$, no safeguard policy supports approval at time T .
- If $\mathcal{A}_0(\alpha, m) = \emptyset$, immediate approval is infeasible.
- Delay can occur for two distinct reasons:
 - ① no policy satisfies the safety constraint,
 - ② feasible safety policies exist, but their irreversible burden is still too high.

Definitions I

Definition (Safety-index function)

For fixed $(\alpha, m) \in (0, 1) \times \mathbb{R}$, define the safety-index function $g_{\alpha, m} : \mathbb{R}_+^n \rightarrow \mathbb{R}$ is defined as

$$g_{\alpha, m}(x) := Q_{\alpha}[R(x)] - m.$$

Definition (Asymptotic safeguard effectiveness)

The safety-index function $g_{\alpha, m}$ is said to satisfy *asymptotic safeguard effectiveness* if

$$\limsup_{\|x\| \rightarrow \infty} g_{\alpha, m}(x) \leq 0.$$

The asymptotic safeguard effectiveness condition is not tied to a particular distribution. Rather, it states that if safeguards become sufficiently strong, the safety-index eventually becomes non-positive, so the safety condition is eventually satisfied. This condition may fail if risk has a lower bound above the standard.

Results I: Non-emptiness of the Safety-Feasible Set

Theorem (1)

Fix $(\alpha, m) \in (0, 1) \times \mathbb{R}$. If the safety-index function $g_{\alpha, m}$ satisfies asymptotic safeguard effectiveness, then the safety-feasible set

$$\mathcal{X}(\alpha, m) := \{x \in \mathbb{R}_+^n : g_{\alpha, m}(x) \leq 0\}$$

is non-empty.

Proof idea.

Since $\sup g_{\alpha, m}(x) \rightarrow \leq 0$ as $\|x\| \rightarrow \infty$, there exists $R_0 > 0$ such that

$$\|x\| \geq R_0 \implies g_{\alpha, m}(x) \leq 0.$$

Choosing any $x \in \mathbb{R}_+^3$ with $\|x\| \geq R_0$ gives $x \in \mathcal{X}(\alpha, m)$. □

Irreversible costs vs MISTIC

- Safeguards reduce irreversible costs by lowering the probability and magnitude of adverse outcomes. However, safeguards also reduce the social incremental reversible benefits of the technology and increase the hurdle rate by reflecting greater uncertainty or irreversibility.
- Therefore, increasing safeguards has an ambiguous effect on admissibility: it lowers the irreversible costs but may also lower the MISTIC threshold.
- Admissibility improves only when the reduction in irreversible costs exceeds the reduction in the MISTIC threshold.

Assumption

For all $(\alpha, m) \in (0, 1) \times \mathbb{R}$, there exists $S_{\alpha, m} \subseteq \mathbb{R}_+^n$ with $S_{\alpha, m} \cap \mathcal{X}(\alpha, m) \neq \emptyset$ such that for any safeguard policy $x \in S_{\alpha, m}$, the MISTIC admissibility gap $G_T(x) = J_T^*(x) - J_T(x)$ is increasing in $x = (x_i)_{i \in \mathcal{I}}$,

$$\frac{\partial G_T(x)}{\partial x_i} > 0 \quad \text{i.e.} \quad -\frac{\partial J_T(x)}{\partial x_i} > -\frac{\partial J_T^*(x)}{\partial x_i}.$$

We call $S_{\alpha, m}$ gap monotonic set.

This assumption is consistent with the interpretation that, starting from insufficiently protected or weakly regulated states, additional safeguards primarily reduce irreversible damages.

However, the assumption need not hold globally. Once safeguards are already stringent, further safeguards may have diminishing effects on irreversible damages while continuing to reduce reversible benefits, in which case the MISTIC admissibility gap may no longer increase.

Results II: Non-emptiness of the Admissibility Set

Theorem

Fix $(\alpha, m) \in (0, 1) \times \mathbb{R}$. If the safety-index function $g_{\alpha, m}$ satisfies asymptotic safeguard effectiveness then every admissibility set contained in $S_{\alpha, m}$ is non-empty.

Proof idea.

For any $x' \in \mathbb{R}_+^n$, let $\mathcal{M}_T(x') := \{x \in \mathbb{R}_+^n : J_T(x) \leq J_T^*(x')\}$. Let $\bar{x}$ be the status quo policy. By definition of MISTIC, we know that $\exists y \neq \bar{x} \in M_T(\bar{x})$ s.t. $y \notin \mathcal{X}(\alpha, m) \cap S_{\alpha, m}$ and $d(\bar{x}, \mathcal{X}(\alpha, m) \cap S_{\alpha, m}) \leq \|y - \bar{x}\|$. Define a point in $M_T(\bar{x})$:

$$x^* := \bar{x} + \theta^*(y - \bar{x}), \quad \text{s.t.} \quad \theta^* := \frac{d(\bar{x}, \mathcal{X}(\alpha, m) \cap S_{\alpha, m})}{\|y - \bar{x}\|} \in (0, 1].$$



Results II: Non-emptiness of the Admissibility Set

Proof idea cont.

from Theorem 1, we know that $d(\bar{x}, \mathcal{X}(\alpha, m) \cap S_{\alpha, m})$ is well-defined. Then

$$\|x^* - \bar{x}\| = \|\theta^*(y - \bar{x})\| = \theta^*\|y - \bar{x}\| = d(\bar{x}, \mathcal{X}(\alpha, m) \cap S_{\alpha, m}).$$

Moreover, $x^* = (1 - \theta^*)\bar{x} + \theta^*y$ and $\bar{x} \leq x^*$. Since MISTIC admissibility gap is increasing in x , $\frac{\partial G_T(x^*)}{\partial x_i^*} > 0$ and safety index function is asymptotic safeguard effective,

$$G_T(x^*) > G_T(\bar{x}) \quad \text{and} \quad x^* \in \mathcal{X}(\alpha, m) \cap \mathcal{M}_T(\bar{x}) = \mathcal{A}_T(\alpha, m) \neq \emptyset.$$



- Approval decisions should not be evaluated by safety alone.
- A technology is admissible only when:
 - ① there exists a safeguard policy that satisfies safety fixed constraint, and
 - ② the resulting irreversible burden is socially tolerable.
- The integrated framework provides a criterion for *whether* and *when* approval becomes justified.

$$\text{Admissibility} = \text{Safety Feasibility} + \text{Irreversibility Tolerability}$$

References I

-  Crouch, E., and Richard W. (1981)
Regulation of Carcinogens.
Risk Analysis, 1 , 47-57.
-  Demont, M., Wesseler, J., Tollens, E. (2004)
Biodiversity versus transgenic sugar beet: the one euro question.
Eur. Rev. Agric. Econ. 31, 1–18.
-  Kataoka, S. (1963)
A Stochastic Programming Model.
Econometrica, XXXI 181-96.
-  Lichtenberg, E. and Zilberman, D. (1988)
Efficient Regulation of Environmental Health Risks.
Quarterly Journal of Economics 103: 167-78



Lichtenberg, E. (2006b)

Regulation of technology in the context of risk generation.

In Regulating Agricultural Biotechnology: Economics and Policy. Springer.



Wesseler, J., Scatasta, S., & Nillesen, E. (2007)

The maximum incremental social tolerable irreversible costs (MISTICs) and other benefits and costs of introducing transgenic maize in the EU-15.

Pedobiologia, 51, 261–269.

Thank you