Dynamic Safety Margins in Environmental Risk Regulation: A Stochastic Extension of the Safety Rule Framework

Kutay Cingiz, Justus Wesseler

August 11, 2025

1. Introduction

In the formulation of environmental health risk of the paper Lichtenberg and Zilberman (1988), the threshold level of health risk log-risk Y is controlled by ensuring that the probability of exceeding a given level Y_0 remains below $1-P$. This is often implemented using a quantile from the standard normal distribution, denoted $F(P)$, so that the regulated threshold becomes:

$$Y(P) = \mathbb{E}[Y] + F(P)S_Y,$$

where S_Y is the standard deviation of the log-risk. The parameter $F(P)$ is treated as fixed and known, derived from the inverse cumulative distribution function (CDF) of the standard normal distribution.

2. Motivation for Extension

The assumption of a static, known $F(P)$ masks a deeper issue: in real-world regulatory settings, the distribution of risk is not perfectly known and may evolve over time due to new scientific findings, monitoring, or public perception. To reflect this dynamic uncertainty, we propose modeling $F(P_t)$ as a stochastic process, specifically a geometric Brownian motion (GBM). This introduces a time-evolving safety margin that better captures regulatory uncertainty.

3. Theoretical Extension: $F(P_t)$ as Geometric Brownian Motion

We assume that the quantile parameter $F(P_t)$ evolves according to a GBM:

$$dF(P_t) = \mu F(P_t)\, dt + \sigma F(P_t)\, dW_t, \tag{1}$$

where:

- μ is the drift (expected growth rate in perceived risk),
- σ is the volatility (degree of uncertainty in risk quantile),
- W_t is a standard Wiener process,
- $F(P_0)$ is the initial value corresponding to the original quantile.

The resulting threshold evolves over time:

$$Y_t(P_t) = \mathbb{E}[Y] + F(P_t)S_Y. \tag{2}$$

3.1. Connection to Regulatory Decision-Making

In the original model, the margin of safety parameter P — and thus the corresponding $F(P)$ — serves as a proxy for the regulator's aversion to uncertainty. By modeling $F(P_t)$ as a geometric Brownian motion, we extend the framework to account for dynamic, time-evolving uncertainty about the risk exceedance threshold. This allows regulatory decision-making to adapt over time, capturing features such as learning, scientific updates, and political shifts in risk tolerance. It also enables the study of optimal timing and intensity of interventions, akin to real options in investment decisions under uncertainty.

4. Implications

4.1. Stochastic Safety Margins

The use of a GBM implies that the safety threshold $Y_t(P_t)$ is now stochastic. As $F(P_t)$ evolves, the regulator must account not only for expected risk but also for uncertainty in how that risk is characterized. From the solution of the GBM, we know that:

$$F(P_t) = F(P_0)e^{\left(\mu - \frac{1}{2}\sigma^2\right)t + \sigma W_t},$$

which implies:

$$\mathbb{E}[F(P_t)] = F(P_0)e^{\mu t}, \quad \text{Var}(F(P_t)) = F(P_0)^2 e^{2\mu t}(e^{\sigma^2 t} - 1).$$

Therefore, the variance of the threshold Y_t becomes:

$$\mathrm{Var}(Y_t) = S_Y^2 \cdot \mathrm{Var}(F(P_t)) = S_Y^2 \cdot F(P_0)^2 e^{2\mu t}(e^{\sigma^2 t} - 1).$$

This shows that the uncertainty in the threshold level increases with both the drift μ and volatility σ, highlighting the importance of dynamic risk management and potential for more conservative policy responses over time.

4.2. Real Options Framework

This structure lends itself naturally to an *optimal stopping problem*, where the regulator faces a trade-off between taking immediate regulatory action and waiting for better information. The evolution of $F(P_t)$ can be interpreted as a signal about the degree of uncertainty or severity in the risk environment. If acting today is costly or irreversible (e.g., banning a product or mandating strict compliance), and future updates may significantly reduce uncertainty, then waiting has value—this is the essence of a real option.

In this view, regulatory decisions are not just static threshold choices but dynamic, contingent actions under uncertainty. The real options framework provides a method for valuing this flexibility and integrating precautionary principles without resorting to overly conservative rules.

4.3. Endogenous Policy Design

One may consider making μ and σ functions of policy effort (e.g., monitoring, R&D investment). For example, suppose $\mu(e)$ and $\sigma(e)$ are functions of effort $e \in [0, 1]$, where higher effort reduces uncertainty:

$$\mu(e) = \mu_0 - \alpha e, \quad \sigma(e) = \sigma_0 - \beta e, \quad \alpha, \beta > 0.$$

Then, the regulator may solve an optimization problem that balances the cost of effort with the value of reduced uncertainty:

$$\min_{e \in [0,1]} C(e) + \lambda \cdot \mathrm{Var}(Y_t(e)),$$

where $C(e)$ is the cost of effort and $\mathrm{Var}(Y_t(e))$ is as derived in Section 4.1 with $\mu(e)$ and $\sigma(e)$ substituted. This endogenizes the risk dynamics based on policy input and supports adaptive regulation.

4.4. Distributionally Robust Policy

Compared to a fixed $F(P)$, the GBM formulation provides robustness to model misspecification. It can be extended further to incorporate ambiguity sets over the distribution of

Y, combining GBM with distributionally robust optimization.Thus, another valuable extension is to treat the risk distribution as ambiguous. Instead of assuming a fixed normal distribution, one can define an ambiguity set $\mathcal{P}$ of plausible distributions and solve:

$$\min_{x} \sup_{\mathbb{Q}\in\mathcal{P}} \mathbb{E}_{\mathbb{Q}}[f(x,\xi)]$$

This approach—called Distributionally Robust Optimization—accounts for worst-case expectations over $\mathcal{P}$ and can be used in conjunction with or instead of modeling $F(P_t)$ via GBM. Examples include moment-based ambiguity sets (known mean and variance), Wasserstein balls, or KL-divergence neighborhoods around a nominal distribution.

4.5. Modified Risk Constraint for Irreversibility

To incorporate irreversibility into the Lichtenberg 2006b risk regulation model, we can modify the original safety constraint:

$$\mu_r(x) + Z(\alpha) \cdot \sigma_r(x) \leq m$$

to include irreversibility adjustment term:

$$\mu_r(x) + [Z(\alpha) + \lambda(x)] \cdot \sigma_r(x) \leq m$$

where $\lambda(x)$ represents the irreversibility premium, the additional value from postponing irreversible actions under uncertainty.

- The first term, $\mu_r(x)$, captures the *average or expected level of risk* resulting from the chosen vector of policy instruments $x = (x_I, x_F, x_S)$.
- The second term, $[Z(\alpha) + \lambda(x)] \cdot \sigma_r(x)$, represents the *safety buffer*, which includes:
  - $Z(\alpha) \cdot \sigma_r(x)$: the *statistical safety margin*, corresponding to a chosen confidence level (e.g., 95% confidence implies $Z(\alpha) \approx 1.645$).
  - $\lambda(x) \cdot \sigma_r(x)$: an additional precautionary margin reflecting the *irreversibility of harm*, such as long-term or irreversible environmental damage.
- The sum of the expected risk and the precautionary margin must not exceed m, the *regulatory ceiling*.

One may also modify the constraint by adding $\lambda(x)$ as a standalone precautionary buffer:

$$\mu_r(x) + Z(\alpha) \cdot \sigma_r(x) + \lambda(x) \leq m$$

This structure preserves interpretability, avoids potential overcompensation, and isolates the irreversibility margin as a direct adjustment to the risk threshold.

The irreversibility adjustment to the model tightens the constraint, discouraging irreversible decisions under uncertainty. This formulation also connects to robust control under ambiguity, ensuring decisions remain safe even under model misspecification. It aligns with the logic of Arrow and Fisher (1974): irreversibility justifies delaying action until more information becomes available.

5. Conclusion

This extension bridges probabilistic risk assessment with dynamic decision-making under uncertainty. Modeling $F(P_t)$ as a geometric Brownian motion adds a layer of realism to environmental risk regulation by embedding learning, precaution, and robustness into the safety threshold framework.

The mathematical formulation enables quantification of how risk perception and regulatory thresholds evolve, how waiting has option value, and how policy effort can shape the trajectory of uncertainty. Furthermore, combining this with distributionally robust optimization allows for greater resilience against model misspecification.

This framework is also applicable from the perspective of private investors operating under regulatory uncertainty. Firms facing uncertain policy environments—such as those in biotechnology, energy, or agri-food sectors—can use the real options model design on regulation and use it as an input for the investor decision making model's cost structure to decide when to invest, expand, or delay based on the stochastic evolution of regulatory risk. In this way, the proposed model not only enhances public regulatory decision-making but also supports strategic private-sector behavior under uncertainty.