Liable Without Control: A Risk-Economic Framework for Managing a Waste Generator's Superfund Liability under Listing Uncertainty

Draft manuscript · July 2026

About this document (background added November 2026)

Under the Superfund law (CERCLA), a company that ships hazardous waste to a licensed treatment, storage, or disposal facility can later be named a potentially responsible party and made to pay part of a cleanup if that facility becomes a Superfund site—however carefully the company chose it. This manuscript grew out of Ernie Kee's work extending the risk economics of protective systems, first developed for nuclear power, to that kind of environmental liability, and out of a presentation he gave to the CHWMEG industry group in June 2026. It treats the waste generator's choice of facilities as a problem of minimizing expected loss, separates the chance that a site is listed into technical and administrative parts, and argues for a graduated severity classification and a Price-Anderson-style financial-responsibility regime. The data and code are archived at `https://doi.org/10.5281/zenodo.21502158`.

About the rejection. The draft was submitted to *Reliability Engineering & System Safety* on 22 July 2026 and was rejected by the editor before it was sent out for peer review. No decision letter was received and no reason for the rejection was given, so why the journal declined it is not known. Because the manuscript was never peer reviewed, its arguments and results have not been evaluated by independent referees, and it should be read as an unreviewed working draft. It is not being pursued for journal publication.

The manuscript that follows is the draft as submitted, with the author list removed.

Liable Without Control: A Risk-Economic Framework for Managing a Waste Generator's Superfund Liability under Listing Uncertainty

Abstract

Under the U.S. Comprehensive Environmental Response, Compensation, and Liability Act (CERCLA), a firm that ships waste to a treatment, storage, and disposal facility can be named a potentially responsible party and made to pay a share of a Superfund cleanup under liability that is strict, joint-and-several, and retroactive. The exposure attaches by having sent waste to a site later listed, largely independent of the generator's fault or the care it took in selecting the facility: diligence changes the *probability* of the triggering event but not the *liability rule*. Motivated by this "divorce" of control from liability, we cast the generator's facility-selection and forward-allocation decision as minimization of expected loss under uncertainty, in which a *catastrophe premium*—the added expected cost attributable to listing—ranks facilities and the firm's risk attitude enters through a coherent tail-risk measure. We separate the listing probability into a technical component, reducible with site data and audits, and an administrative component governed by jurisdictional discretion, estimated by survival modeling with jurisdictional frailty and supplemented by structured expert judgment where data are sparse. A prototype on public records shows that data quality, not method, is the binding constraint. A published, graduated severity classification would convert administrative variance into technical variance, and a toxicity-weighted licensed-quantity ceiling could supply the bounded maximum harm that a Price-Anderson-style financial-responsibility regime—authorized but never implemented under CERCLA Section 108(b)—would require. The

work bridges the law-and-economics literature on CERCLA liability with a quantitative, risk-informed treatment of the generator's forward-looking problem.

Keywords: CERCLA, Superfund, potentially responsible party, joint and several liability, risk-informed decision making, expected loss, conditional value at risk, hazardous waste

---

Nomenclature

Abbreviations

CERCLA	Comprehensive Environmental Response, Compensation, and Liability Act of 1980
ECHO	Enforcement and Compliance History Online
EPA	Environmental Protection Agency
PRP	Potentially Responsible Party
RCRA	Resource Conservation and Recovery Act
RCRAInfo	RCRA National Information System
TSDF	Treatment, Storage, and Disposal Facility

Symbols

Symbols are grouped by the section that introduces them; a few are reused with different subscripts, which the groupings disambiguate. The expectation operator $P(\cdot)$ follows the de Finetti convention $P(X) \equiv \mathbb{E}[X]$.

Decision and cost (Section 3)

$\mathcal{A}$	set of disposal facilities (alternatives)
$\alpha,\ \alpha^\star$	a facility; the selected facility
C_α	generator's lifecycle cost of using facility α
χ_α	listing indicator (1 if the facility is listed)
p_α	listing probability, $P(\chi_\alpha = 1)$
u	utility (disutility) defined over cost
$c_\alpha^g,\ c_\alpha^f$	expected cost if not listed / if listed
c_α^p	catastrophe premium, $c_\alpha^f - c_\alpha^g$

Forward allocation (Section 4)

$\mathcal{S}$	index set of facilities (i)
$\mathcal{S}_w$	feasible facilities permitted for and accepting stream w
t, T	planning period; planning horizon
$x_{i,t}$	volume routed to facility i in period t (decision)
y_i	facility-use indicator (decision)
κ_i	qualification / switching cost for site i
$V_i^0, \Delta V_i, V_i$	legacy, planned, and total volume at site i
s_i	generator's equitable share of cleanup at site i
τ_i, Q_i	toxicity weight; total toxicity-weighted volume at site i
K_i	expected cleanup cost at site i
$L^0, \Delta L$	sunk legacy and incremental expected liability
β	discount factor, $(1+r)^{-1}$
$\pi_{i,t}$, w_t, $\overline{x}_{i,t}$	unit disposal price; waste generated; site capacity
CVaR_δ	conditional value at risk at tail level δ
C^{disp}	discounted disposal cost

Listing probability and estimation (Sections 5 and 6)

R_i	regulatory / jurisdictional state of site i
$\mathrm{Var}(p_i)$	variance of the listing probability
T_i	time from a fixed origin to listing
$\lambda_i(t \mid \cdot)$, $\lambda_0(t)$	hazard rate of site i; baseline hazard
$\mathbf{x}_i$	site covariate vector
$\boldsymbol{\beta}$	fixed-effect coefficients (technical drivers)
$b_{R(i)}$, σ_R^2	jurisdiction random effect; frailty variance
$\mathcal{E}_i$	evidence accumulated on facility i
$\pi(p_i \mid \mathcal{E}_i)$, π_0	posterior of p_i; elicited prior
$\mathcal{L}$	likelihood of the evidence

Policy (Section 7)

$\mathbf{q}_i$	measurable site quantities (releases, concentrations, factors)
$g(\cdot)$, τ	severity function; published band cutoff

1. Introduction

The motivating tension is legal and is, from a risk standpoint, peculiar. CERCLA imposes on a waste generator, as an "arranger," liability that is strict, joint-and-several, and retroactive: having sent waste to a site that later becomes a Superfund site can make the generator a PRP and expose it to a share of the cleanup, regardless of fault, of whether its waste caused the release, and—crucially for the decision problem—of the diligence with which

it selected the facility. Care in facility selection is not a statutory defense; it lowers the probability of the triggering event without altering the liability rule. The party therefore controls *which* probability of loss it is exposed to, through site selection and allocation, but neither the magnitude of the loss rule nor the listing decision itself.

The economic analysis of CERCLA liability is, by contrast, well developed, and it has concentrated on two questions that bracket—but do not include—the one posed here. The first is *deterrence*: whether the statute's strict, joint-and-several, and retroactive standard induces efficient *ex ante* care by the parties whose conduct determines whether a release occurs. The answer is famously equivocal. Tietenberg [1] showed that joint and several liability for indivisible harms shifts care incentives in ways that turn on the number of injurers and the apportionment rule; Kornhauser and Revesz [2] identified conditions under which the rule actually *weakens* the incentive to prevent contamination relative to several-only liability; and the wider penalty-design literature treats liability as one instrument among several for internalizing environmental harm, particularly where responsible parties may be judgment-proof [3]. What unifies this strand is its object: the owner or operator whose conduct sets the probability that a site is contaminated in the first place.

The second question is *allocation*: once a site is listed and cleanup costs are incurred, how are they divided among the potentially responsible parties, and with what collateral effects? This literature documents the §113 contribution action, the courts' reliance on open-ended "equitable factors," the settlement dynamics that follow, and the heavy transaction costs and orphan-share problems of multi-party Superfund litigation [4, 2]. Empirical work finds that these features shape cost recovery in practice and spill into land markets: joint and several liability under Superfund measurably discourages the redevelopment of potentially contaminated "brownfield" sites [5, 6]. This strand, too, is concerned with events *after* listing and with the division of a loss that has already been realized.

Neither strand addresses the decision a waste generator must make *before* any release: which facilities to use, and in what proportion, when the probability that a given facility is ever listed is uncertain and the liability that would follow is fixed by a rule the generator cannot influence. This is the inversion noted above—the deterrence literature studies the party that sets the hazard probability, whereas the generator is the party that does not, and can choose only which probability to be exposed to. Both adjacent literatures nonetheless supply ingredients the present development uses. The deterrence results explain *why* a generator cannot safely rely

on the operator's diligence: an under-deterred or judgment-proof operator pushes residual exposure onto the parties the statute can still reach [7]. The allocation results, in turn, characterize the loss a generator would bear if listing occurs—an equitable share of a jointly-and-severally apportioned, transaction-cost-laden cleanup—which is exactly the consequence term in the expected-loss objective constructed below. We take that objective from the protective-system risk economics of Kee et al. [8] and Wortman et al. [9], in which alternatives are ranked by expected cost and a *catastrophe premium* prices the incremental cost of the adverse event, and we adapt it from a regulated operator's choice of protective technology to a generator's choice of disposal facility. Our contribution is threefold: to formulate the generator's prospective, risk-informed facility-selection and allocation problem; to connect it to both the legal structure that creates the exposure and the estimation problem for the listing probability; and to draw from the contrast with the nuclear sector's Price-Anderson regime a concrete reform—a toxicity-weighted licensed-quantity ceiling—that would supply the bounded maximum harm a comparable financial-responsibility structure requires. The remainder of the paper develops the expected-loss objective (Section 3), extends it to a multi-period forward allocation (Section 4), decomposes the listing probability into a technical and an administrative component (Section 5), discusses its estimation from public records (Section 6), and closes with a policy implication (Section 7).

2. The generator's position under CERCLA

Section 107(a) of CERCLA [10] imposes cleanup liability on four classes of parties: the current owner and operator of a facility, the owner or operator at the time of disposal, any person who *arranged for* the disposal or treatment of hazardous substances, and any person who transported such substances to a site of its selection (42 U.S.C. § 9607(a)). A waste generator falls in the third class, as an "arranger." The courts have construed this liability as strict, joint-and-several, and retroactive: it requires no showing of fault or negligence; any single responsible party may in principle be held for the entire cost of remediation where the harm is indivisible; and it reaches conduct that predates the statute [11, 4]. The trigger is the act of having arranged for disposal at a facility from which there is a release, or threatened release, of a hazardous substance that gives rise to response costs—not any failure of care by the generator.

A party that has borne more than its share may seek contribution from other responsible parties under Section 113(f), with courts apportioning

costs among them by open-ended "equitable factors" (42 U.S.C. § 9613(f)). Joint-and-several liability is not, however, automatic. Since *Chem-Dyne* [12] the rule has been that a defendant may escape it by proving the harm divisible—showing "a reasonable basis for determining the contribution of each cause"—and then pays only an apportioned share, a divisibility defense the Supreme Court later affirmed in *Burlington Northern* [13]. Whether it succeeds turns on the facts of the site: in *Bell Petroleum* [14] a single contaminant from non-overlapping operators admitted a clean *volumetric* apportionment, whereas in *Chem-Dyne* itself the commingled, heterogeneous waste of hundreds of contributors did not. Because most sites resemble the latter, courts have applied the divisibility analysis narrowly, and joint-and-several exposure remains the realistic default. The combination of joint-and-several liability with contribution shifts risk onto the solvent and identifiable: the shares attributable to insolvent, defunct, or unlocatable parties—"orphan shares"—are redistributed among those who remain, and the contribution process is itself slow and costly. Empirically, these features dominate observed cost recovery and spill into adjacent land markets [6, 5]. The practical consequence for a generator is that the loss it would bear upon listing is not its physical contribution to the contamination but its *equitable share* of a jointly apportioned, transaction-cost-laden cleanup—a quantity inflated by the orphan shares of co-contributors it neither chose nor can observe.

This exposure sits oddly beside the generator's duties under the Resource Conservation and Recovery Act [RCRA; 15], which governs the ongoing management of hazardous waste "from cradle to grave" (42 U.S.C. § 6901 et seq.). RCRA's obligations are prospective and procedural: the generator must characterize its waste, manifest it, and ship it to a permitted treatment, storage, and disposal facility. Discharging these duties faithfully—selecting a properly permitted facility, and even auditing it—satisfies RCRA but supplies *no defense* to CERCLA liability should that facility later be listed. The statutory defenses to CERCLA are narrow and do not turn on the care of selection: Section 107(b) excuses only releases caused solely by an act of God, an act of war, or an unrelated third party with whom the defendant has no contractual relationship, and the later carve-outs—the de micromis and municipal-solid-waste exemptions, and the small-business and bona fide prospective purchaser provisions added in 2002—are defined by the kind or quantity of waste and by a party's relationship to the site, not by diligence in choosing a facility [11]. A generator that does everything RCRA asks therefore lowers the probability that its chosen facility is mismanaged and releases, but does not alter the rule that would govern its liability if a release

occurs.

This is the divorce of control from liability that motivates the rest of the paper. The party whose conduct actually sets the probability of a release—the facility's owner or operator—is the subject of the deterrence literature reviewed above and may itself be under-deterred or judgment-proof, in which case the residual exposure migrates to the solvent generators the statute can still reach [7, 2, 3]. The generator, for its part, holds a different and narrower lever: it can set neither the listing probability of any facility, nor the liability rule, nor the listing decision, but it can choose *which* facilities to expose itself to, and in what proportion. That choice—made before any release, under uncertainty about which facilities will ever be listed—is a problem of decision under uncertainty, and it is the problem Section 3 formalizes.

3. The decision problem: ranking facilities by expected loss

A waste generator choosing among disposal facilities is making a decision under uncertainty, and the natural object of choice is expected loss. We adopt the de Finetti convention of writing the expectation of a random variable X as $P(X)$ rather than $\mathbb{E}[X]$, which keeps the probability of an event and the expectation of a cost in one compact notation. Let $\mathcal{A}$ index the available facilities. For each alternative $\alpha \in \mathcal{A}$, let C_α denote the generator's lifecycle cost of using that facility—the routine disposal price together with any liability that later materializes—and let $\chi_\alpha \in \{0, 1\}$ indicate whether the facility is ever placed on the National Priorities List, with $p_\alpha = P(\chi_\alpha = 1)$ its listing probability. A risk-averse decision-maker with utility u defined over cost selects

$$\alpha^\star \;=\; \arg\min_{\alpha \in \mathcal{A}} \; P(u \circ C_\alpha)\,, \tag{1}$$

the facility that minimizes the expected disutility of cost. Equation (1) is the von Neumann–Morgenstern criterion written for cost rather than payoff: with a risk-neutral u it reduces to minimizing expected cost, and with a concave u it penalizes dispersion, a property taken up in Section 4.

Conditioning the expected cost on the listing outcome separates the two states the generator cares about. Writing $c_\alpha^g = P(u \circ C_\alpha \mid \chi_\alpha = 0)$ for the expected cost when the facility is never listed and $c_\alpha^f = P(u \circ C_\alpha \mid \chi_\alpha = 1)$ for the expected cost when it is, the law of total expectation gives

$$P(u \circ C_\alpha) \;=\; c_\alpha^g + \left(c_\alpha^f - c_\alpha^g\right) p_\alpha. \tag{2}$$

The increment $c_\alpha^p \equiv c_\alpha^f - c_\alpha^g$ is the *catastrophe premium* of the alternative—the added expected cost attributable to listing—and $c_\alpha^p\, p_\alpha$ is the term that

distinguishes two facilities whose routine prices are alike. It is precisely the quantity a generator trades against price whenever it chooses where to send waste.

The content of that premium is the loss characterized in Section 2. If facility α is listed, the generator's added cost is not its physical contribution to the contamination but its equitable share of a jointly-and-severally apportioned, transaction-cost-laden cleanup—a share inflated by the orphan portions of insolvent co-contributors and fixed by a rule the generator cannot influence. Reputational loss enters here too, to the extent the generator chooses to price it. The decision rule (1) thus consumes, as its consequence term, exactly the liability the legal structure creates; the generator's problem is to weigh that contingent loss, discounted by its probability, against the certain price of disposal.

Two features of (2) organize the rest of the paper. First, the ranking is governed by expected loss, not by hazard. The EPA Hazard Ranking System and third-party facility-review scores measure the hazard a *site* poses to its surroundings; they are informative about p_α and enter the decision as covariates in the estimate of that probability (Section 6), but they are not themselves the decision variable and cannot be substituted for $c^p_\alpha\, p_\alpha$. A facility may score as hazardous yet pose little expected loss to a small contributor, or score modestly yet expose a dominant customer to a large share; only the expected-loss form distinguishes these cases. Second, the listing probability p_α is the one factor in (2) that is both uncertain and partly estimable, which is why its structure and estimation occupy Sections 5 and 6.

The construction is adapted from the protective-system risk economics of Kee et al. [8] and Wortman et al. [9], in which a regulated operator ranks protective designs by expected social cost and a catastrophe premium prices the incremental cost of a severe event. One feature of the adaptation must be made explicit, because it is the same inversion drawn in Section 2. In the protective-system setting the decision-maker *sets* the catastrophe probability p_α through the design it selects, and the apparatus built on the selection rule—locating a socially optimal margin of safety by trading reliability against premium—presumes that control. The waste generator has no such control: it selects *which* p_α to be exposed to, through site selection and allocation, while the level of p_α is set by the operator, the co-contributors, and the regulators. The selection rule (1) and the decomposition (2) carry over intact, since they require only that each alternative carry some listing probability and some consequence; the margin-of-safety machinery does not, and we do not invoke it. What replaces it, for the generator, is the choice of how to

spread waste across facilities over time—the forward-allocation problem of Section 4.

4. Forward allocation and the sunk legacy

The selection rule of Section 3 chooses a single facility, but in practice a generator splits a recurring waste stream across several facilities and over time, so the object of choice is an allocation rather than a single α. Let $\mathcal{S}$ index the available facilities and $t = 1, \dots, T$ the planning periods. Write $x_{i,t} \geq 0$ for the volume routed to facility i in period t, and $y_i \in \{0, 1\}$ for an indicator that facility i is brought into use, incurring a one-time qualification or switching cost κ_i. Each alternative of Section 3 is now a facility $i \in \mathcal{S}$ carrying its own listing probability p_i, expected cleanup cost K_i, and the generator's equitable share s_i of that cleanup.

It is worth being precise about the kind of harm this share apportions. A sudden release—a fire or a tank rupture—is an accident in the classical sense, a discrete stochastic event; but the contamination that typically drives Superfund listing is *chronic*, accumulating with the scale and duration of waste handling rather than issuing from a single mishap. The economics of accident law separates the level of *care* from the level of *activity*, and shows that strict liability, unlike negligence, makes an injurer bear expected harm that rises with its activity [16]; it is that activity-level branch—harm proportional to volume—that the share s_i captures. The listing event itself is neither accident nor activity: it is a legal trigger that bundles the physical harm with the administrative discretion isolated in Section 5, which is why p_i must not be read as an accident probability.

Two accounting facts shape the problem. First, the share a generator would bear at a site grows with the waste it has sent there; a defensible first model for the equitable share is the toxicity-weighted volumetric fraction,

$$s_i(V_i) \;\approx\; \frac{\tau_i \, V_i}{Q_i}, \tag{3}$$

where V_i is the generator's cumulative toxicity-weighted volume at site i, τ_i a toxicity weight, and Q_i the total toxicity-weighted volume managed at the site across all contributors. This volumetric form is more than a modeling convenience: a toxicity- and volume-based division is exactly the kind of "reasonable basis for determining the contribution of each cause" that *Burlington Northern* accepts for apportioning liability rather than imposing it jointly and severally [13]; a volumetric apportionment of exactly this form prevailed in *Bell Petroleum* [14]. A generator that keeps—and

documents—a small share both lowers its expected loss and strengthens its position should apportionment be litigated. Second, that cumulative volume divides into a part the generator can no longer change and a part it is now choosing, $V_i = V_i^0 + \Delta V_i$, with V_i^0 the volume already shipped and $\Delta V_i = \sum_t x_{i,t}$ the volume to be shipped under the plan. Because expected liability is $\sum_i p_i\, P[s_i(V_i)\, K_i]$, it separates into a *sunk* term in V_i^0 alone and an *incremental* term in the decision,

$$\begin{aligned}&\sum_{i\in\mathcal{S}} p_i\, P\big[s_i(V_i^0 + \Delta V_i)\, K_i\big] \\ &\quad = \underbrace{\sum_{i\in\mathcal{S}} p_i\, P\big[s_i(V_i^0)\, K_i\big]}_{L^0\ \text{(sunk legacy)}} + \Delta L(\Delta V).\end{aligned} \tag{4}$$

The legacy term L^0 does not depend on the decision variables and drops out of the optimization—though it should be reported as the enterprise's standing exposure—and only ΔL is optimized. This is the formal content of the "go-forward" principle: past shipments are a fixed liability the firm carries, and the lever it actually holds is the future allocation $\{x_{i,t}\}$.

Collecting the pieces, the risk-neutral program minimizes discounted disposal cost plus incremental expected liability plus switching cost,

$$\begin{aligned}\min_{\{x_{i,t}\},\,\{y_i\}}\ &\underbrace{\sum_{t=1}^{T} \beta^t \sum_{i\in\mathcal{S}} \pi_{i,t}\, x_{i,t}}_{\text{disposal cost}} + \underbrace{\sum_{i\in\mathcal{S}} p_i\, P[s_i(V_i)\, K_i]}_{\text{expected liability}} \\ &+ \underbrace{\sum_{i\in\mathcal{S}} \kappa_i\, y_i}_{\text{switching}},\end{aligned} \tag{5}$$

subject, in each period, to meeting the generated quantity $\sum_i x_{i,t} = w_t$, to facility capacity $0 \le x_{i,t} \le \overline{x}_{i,t}\, y_i$, and to the volume accounting $V_i = V_i^0 + \sum_t x_{i,t}$. Here $\beta = (1+r)^{-1}$ is the discount factor, $\pi_{i,t}$ the unit disposal price, w_t the waste generated in period t, and $\overline{x}_{i,t}$ the capacity of site i. With $T = 1$ and all waste forced to one site, (5) collapses to the single-alternative rule (1).

Two terms in (5) are written compactly but carry operational structure a generator cannot ignore. The unit price $\pi_{i,t}$ is best read as a *delivered* cost—the facility's gate fee together with the cost of transporting the waste to it, which rises with the distance to site i—so that a nearer but nominally

dearer facility can dominate a cheaper, more distant one, and the allocation trades catastrophe premium against gate price and haul jointly. The feasible set is also narrower than $\mathcal{S}$: a facility's permit authorizes only specified waste codes, and not every permitted facility will accept a given generator's particular stream, so the sums in (5) properly run over the subset $\mathcal{S}_w \subseteq \mathcal{S}$ of facilities both authorized for and willing to take stream w. These refinements enlarge the data the program consumes but leave its structure intact; they also bound diversification from the other side, since a thin feasible set $\mathcal{S}_w$ limits how widely waste can be spread and may force the very concentration the liability term penalizes.

The structure of the program (5) yields the paper's most directly usable prescription. Because the equitable share (3) is increasing in the generator's own volume at a site, the expected-liability term is *convex* in that volume: doubling the quantity sent to one facility more than doubles its contribution to expected loss, whereas spreading the same quantity across several facilities does not. Concentration is therefore penalized and the optimum diversifies—the analytic form of the operational maxim not to become any one facility's dominant customer. Diversification is not free: a generator that scatters waste across many facilities pays more in switching costs κ_i and forgoes volume discounts, so (5) trades the liability benefit of spreading against its operational cost.

A second, legal channel reinforces the same prescription. The share (3) is the loss a generator bears only where the harm is *divisible*; where it is not—at sites that commingle heterogeneous wastes from many contributors, as in *Chem-Dyne* [12]—the apportionment defense fails and exposure runs toward the entire K_i under joint-and-several liability, however little the generator sent. The feasibility of apportionment is thus itself a facility attribute, and it belongs in the catastrophe premium: a homogeneous, few-contributor, traceable site carries a smaller premium than a commingled one at the same listing probability. A generator therefore has reason to prefer facilities whose waste streams remain distinguishable, and to document its own volumetric and toxicity contribution prospectively—since the divisibility defense, though available, must be proven from a record that has to exist before any litigation does.

Two refinements matter for a faithful model. First, a risk-neutral objective understates what a generator actually fears, which is not the average outcome but the rare, large one—a single listing that brings a disproportionate share. This is the concavity of the utility u in Section 3, and it is captured by replacing the expected-liability term with a coherent tail-risk measure such

as the conditional value at risk,

$$\min_{\{x_{i,t}\},\,\{y_i\}} \ \mathrm{CVaR}_{\delta}\Big[\, C^{\mathrm{disp}} + \sum_{i\in\mathcal{S}} \chi_i \, s_i(V_i)\, K_i \\ + \textstyle\sum_i \kappa_i y_i \,\Big], \tag{6}$$

where C^{disp} is the discounted disposal cost of the first term of (5) and CVaR_{δ} weights the worst δ-fraction of outcomes, so that allocations with a thin tail are favored. Second, listing outcomes are not independent: facilities under the same authorized state program, the same EPA region, or the same operator tend to be listed together, so the covariance of the $\{\chi_i\}$ belongs in the model, and diversification across correlated sites buys less than a naive count of sites would suggest. The quality of any solution to (5) or (6) is bounded by the quality of the listing probabilities p_i it consumes, which is the subject of Sections 5 and 6.

A final consideration is one of scope rather than refinement. The program (5) is written for a single generator that takes the facilities' listing probabilities p_i, cleanup costs K_i, and continued operation as given; none of these is, in general, exogenous to the generators' collective behavior. If many generators apply the same logic and divert volume away from a dispreferred facility, that facility loses the revenue that sustains it, and a site driven toward closure or insolvency is precisely the judgment-proof operator of Section 2: its share of any eventual cleanup does not disappear but is redistributed as orphan shares onto the contributors who remain, raising their equitable shares s_i even as each acts to lower its own. The forward optimum is therefore a partial-equilibrium prescription for one firm; a fuller account would treat facility viability—and hence the available set $\mathcal{S}_w$, the listing probabilities, and the orphan-share component of s_i—as endogenous to the aggregate of these decisions, a general-equilibrium extension we flag here rather than pursue.

5. The listing probability and its variance

Every quantity in the allocation of Section 4 is known to the generator except the listing probabilities p_i, and the value of any solution to (5) or (6) is bounded by how well those probabilities are estimated. It is therefore worth being precise about where the uncertainty in p_i comes from, because its two sources behave very differently. Conditioning on the regulatory and jurisdictional state R_i in which a facility sits—principally the authorized state

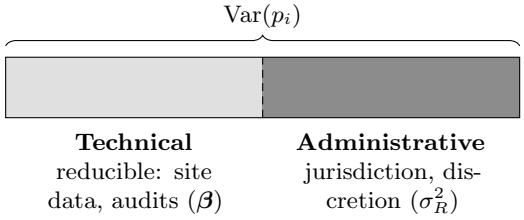


Figure 1: The variance decomposition (7) of the listing probability p_i. The technical component (the $\boldsymbol{\beta}$ term of (8)) is reducible with better site data and audits; the administrative component (the σ_R^2 term), set by jurisdiction and regulatory discretion, lies largely outside the generator's control, and the prototype of Section 6.1 finds the two of comparable magnitude. Greater regulatory specificity (Section 7) would move the dashed boundary so as to shrink the administrative component, converting administrative variance into technical variance.

program and EPA region—the law of total variance splits the uncertainty in the listing probability into two parts,

$$\operatorname{Var}(p_i) = \underbrace{P[\operatorname{Var}(p_i \mid R_i)]}_{\text{technical}} + \underbrace{\operatorname{Var}(P[\,p_i \mid R_i\,])}_{\text{administrative}}. \tag{7}$$

The first term (Figure 1) is physically grounded uncertainty: within a given jurisdiction, whether a particular site is listed depends on pathway susceptibility (depth to groundwater, aquifer vulnerability, proximity to surface water and population), on facility engineering, and on the conduct of the operator and the co-contributors. This component is reducible with better site data and audits—it is the part a facility-review report speaks to, and the part a statistical model can hope to capture. The second term measures how much the listing probability depends on the regulatory state itself, holding the physical site fixed: the same facility, with the same releases, faces materially different odds of reaching the National Priorities List depending on the state in which it sits and the discretion exercised there. This between-jurisdiction term is not small; the prototype of Section 6 finds its spread across jurisdictions comparable to the average listing rate itself.

The administrative term is large because the listing decision, though it rests on a quantitative score, is not a published deterministic function of that score. The Hazard Ranking System produces a number with a threshold near 28.5, but clearing the threshold confers *eligibility* for the National Priorities List, not listing: above it the EPA still chooses whether to propose a site, may defer an eligible site to the Resource Conservation and Recovery Act program, and may weigh state priorities against limited cleanup resources,

while a state may advance its single top-priority site largely irrespective of score. A substantial share of $\mathrm{Var}(p_i)$ is therefore administrative rather than technical—arising from discretion and jurisdiction rather than from the physical site—and is reducible only by greater regulatory predictability, a lever outside the generator's control.

This has a direct consequence for modeling and another for policy. For modeling, p_i should be carried not as a point estimate but as a distribution conditioned on jurisdiction, so that the administrative uncertainty propagates honestly into the allocation rather than being hidden inside a single number; this is the object Section 6 sets out to estimate. For policy, the administrative term is a feature of the regulatory design rather than a natural constant, and could in principle be shrunk—the argument Section 7 takes up.

6. Estimating the listing probability

The allocation of Section 4 and the decomposition of Section 5 both require an estimate of the listing probability p_i, which Section 3 took as given. Estimating it from public records is hard for one dominant reason: listing is rare. Among operating treatment, storage, and disposal facilities the base rate of National Priorities List listing runs well below a percent, so positives are scarce and the class imbalance is severe. The honest consequence, stated at the outset, is that the product of any such estimation is a *calibrated screen and a relative ranking*, not a precise probability for an individual site; absolute probabilities carry wide intervals and should be reported as such.

The natural target is not a static label but a time. Let T_i be the time from a fixed origin—permit issuance, say—to listing, right-censored for the great majority of facilities that have not listed by the observation date. A proportional-hazards model with a jurisdictional frailty captures both the site-level drivers and the jurisdiction effect that Section 5 isolated:

$$\begin{aligned} \lambda_i(t \mid \mathbf{x}_i, R_i) &= \lambda_0(t)\, \exp\!\big(\mathbf{x}_i^{\top}\boldsymbol{\beta} + b_{R(i)}\big), \\ b_R &\sim \mathcal{N}(0, \sigma_R^2), \end{aligned} \tag{8}$$

where $\mathbf{x}_i$ is a vector of site covariates, $\lambda_0(t)$ a baseline hazard absorbing secular trends in listing activity, and $b_{R(i)}$ a random effect for the authorized state or EPA region of facility i. The two pieces of (8) map exactly onto the two terms of (7): the fixed effects $\boldsymbol{\beta}$ are the "problem-site parameters"—the technical drivers and their weights—while the frailty variance σ_R^2 is an estimate of the administrative, between-jurisdiction term. The model thus turns the variance decomposition from a remark into a measurement. For a

cross-sectional snapshot the static analogue is a penalized (elastic-net) logistic model, $\Pr(\chi_i = 1) = \text{logit}^{-1}(\mathbf{x}_i^\top \boldsymbol{\beta} + b_{R(i)})$; the survival form is preferred because it does not discard the censored histories that make up almost all of the data.

The covariates $\mathbf{x}_i$ are the site characteristics the Hazard Ranking System itself weighs, drawn from public sources: pathway susceptibility (depth to groundwater, aquifer vulnerability, proximity to surface water and to population), waste toxicity, and the facility's compliance and corrective-action history as recorded in the RCRAInfo and ECHO systems. These are the same indicators a facility-review report compiles, now used as covariates in an estimate rather than as a score in their own right.

Rarity has a constructive remedy. Where outcomes are too few to identify the model from data alone—the situation here—structured expert judgment can supply the information the data cannot, and Bayesian updating disciplines the combination; this is long-standing practice for rare, high-consequence events in probabilistic risk assessment [17, 18]. A prior is placed on each block of (8)—hydrogeologists informing the pathway coefficients, waste-chemistry specialists the toxicity weights, and regulatory experts the jurisdictional frailty—under a structured protocol in which calibration questions weight experts by demonstrated performance [19], individual judgments are combined by a defensible rule [20, 21], and the process is documented to professional standards [22, 23]. Writing $\mathcal{E}_i$ for the evidence accumulated on facility i, the site estimate is then refined by

$$\pi(p_i \mid \mathcal{E}_i) \;\propto\; \mathcal{L}(\mathcal{E}_i \mid p_i)\, \pi_0(p_i), \tag{9}$$

with π_0 the elicited prior. Crucially, (9) returns a *distribution*, not a point—the jurisdiction-conditioned p_i that the risk-averse objective (6) is built to consume, and that a generator would update each review cycle before re-optimizing its allocation.

6.1. A prototype, and what binds

A prototype exercised this pipeline on the public record, and its main result is a caution about the data rather than a new estimate. The jurisdictional layer of (8) was fit as a conjugate beta-binomial random-effects model on state listing counts: across roughly fifty jurisdictions and about a hundred listings against a pooled rate near 0.004, empirical-Bayes shrinkage put the between-state standard deviation at about the same magnitude as the average rate itself—a direct measurement confirming that the administrative term of (7) is as large as the mean it perturbs. Carrying the pipeline to

the facility level then exposed the binding constraint. Lacking a shared key between the facility inventory and the listing record, facilities were matched by blocked fuzzy name agreement; of some twenty-five thousand inventory facilities checked against the final listings, only a fraction of a percent matched with high confidence, and those matches were dominated by federal facilities whose distinctive names appear verbatim in both lists. Worse, the public inventory is overwhelmingly non-hazardous—publicly owned treatment works and transfer stations dominate it—rather than the RCRA Subtitle C hazardous-facility population the risk question concerns.

The lesson is that the binding constraint on a site-level model is data quality, not method. Before the frailty fit of (8) can be honest, the population must be restricted to commercial Subtitle C facilities, the listing labels hardened with consistent identifiers, and the physical pathway covariates actually joined to each site. Within those limits the model still does real work: it sharpens the technical component of p_i, ranks facilities for deeper review, and yields an explicit, weighted list of risk drivers. What it cannot do is predict the administrative component—the frailty variance σ_R^2 measures the discretion in listing but does not remove it, and no site-feature model forecasts a regulator's decision to advance, defer, or decline a site. The output of the estimation is therefore best carried forward, as Section 4 requires, not as a point value of p_i but as a jurisdiction-conditioned distribution, so that the residual administrative uncertainty propagates into the allocation rather than hiding inside a single number.

7. Policy implication: graduated severity and deterrence

The decomposition of Section 5 carries a policy implication that the modeling of Section 6 sharpens rather than supplies. A large part of the uncertainty a generator faces in p_i is administrative—it lives in the discretionary, between-jurisdiction term of (7)—and that term is not a natural constant but an artifact of how the listing rule is written. The central observation of this section is that greater specificity in the rule would convert administrative variance into technical variance, moving uncertainty out of the part of the problem the enterprise cannot model and into the part it can.

The raw materials for such a rule already exist; what is missing is their binding assembly. CERCLA already defines Reportable Quantities for hazardous substances—the Section 302.4 list of release amounts that trigger a notification duty. The Hazard Ranking System is itself a quantitative composite, concentration-to-benchmark ratios feeding a score with a published cutoff near 28.5. RCRA in turn supplies action levels, land-disposal

treatment standards, and maximum contaminant levels that are quantitative and enforceable. The system is not short of numbers. The gap is twofold: crossing the Hazard Ranking System cutoff confers *eligibility* rather than listing, and there is no published, graduated severity scale that binds a measured level of release to a defined regulatory consequence.

To see what such a scale would buy, suppose listing were a fixed function of measurable site quantities,

$$\chi_i = \mathbf{1}\{\, g(\mathbf{q}_i) \geq \tau \,\},$$

where $\mathbf{q}_i$ collects releases, concentrations, and the pathway and target factors the Hazard Ranking System already scores, and τ is a published band cutoff. Then $P[\, p_i \mid R_i \,]$ would depend on the jurisdiction R_i only through the distribution of the measurable $\mathbf{q}_i$ within it, and the purely *discretionary* component of the administrative term in (7)—the part attributable to a regulator's choice to advance, defer, or decline a site of given measured severity—would vanish. What remained would be measurement uncertainty in $\mathbf{q}_i$, which is technical and reducible with better site data. In the terms of (8), the frailty variance σ_R^2 would shrink toward zero and the fixed effects $\boldsymbol{\beta}$ would carry essentially all of the explainable variation.

This design is not hypothetical in the regulatory world from which the paper's framework is drawn. The Nuclear Regulatory Commission's Reactor Oversight Process sorts quantitative performance indicators and inspection findings into color bands—green, white, yellow, and red—and its Action Matrix ties each band to a prescribed, escalating agency response, so that the consequence of a given measured performance is settled in advance rather than negotiated after the fact; the International Nuclear and Radiological Event Scale, with its levels 0 through 7, does the same for events. Kee et al. [8] and Wortman et al. [9] frame regulatory authority precisely along the axis between *prescriptive* standards—bright, published lines—and *performance-based* standards—outcome targets administered with discretion—and show that the choice between them shapes the margin of safety an operator will hold. A graduated, published environmental severity scale is the prescriptive end of that axis applied to the listing decision.

The advantages compound through the machinery of this paper. The value of the estimation model of Section 6 rises, because administrative noise no longer swamps the technical signal it is built to recover. The allocation of (5) and (6) becomes genuinely actionable, because p_i turns into a forecastable function of observable site quantities rather than a jurisdictional lottery. The precaution calculus is well defined: an operator facing a published band

cutoff knows what margin of safety to hold, which is the Learned Hand condition—the calculus of negligence—made operational [8]. The insurance market prices the exposure more sharply, tightening the market signal a generator can read. And cross-jurisdictional inequity shrinks, because the same measured severity yields the same consequence in every authorized state.

Against these advantages stand serious objections, and intellectual honesty requires giving them their due. Consequence is genuinely site-specific: the same release matters far more over a sole-source aquifer beside a town than in an arid, unpopulated basin, which is why the Hazard Ranking System is a composite of quantity, mobility, and targets rather than a single mass, and why any workable scale must be defined on that composite and conditioned on setting. Sharp cutoffs invite threshold effects, as facilities manage to just below a line. Discretion serves legitimate ends—the regulator's ability to defer eligible sites to RCRA, to prioritize against limited cleanup resources, and to accommodate genuine idiosyncrasy is not merely noise. Disputes migrate rather than disappear: a rule-bound assignment moves the contest onto the measured inputs $\mathbf{q}_i$—sampling, characterization, and the choice of cutoffs. And an irreducible floor remains, since population exposure, co-contributor conduct, and measurement error keep some uncertainty technical but large even under a perfect rule. On balance these considerations point not to a single national threshold—too blunt to be either fair or protective—but to a published, graduated, setting-conditional severity classification tied to defined consequences, with discretion retained only at the genuinely ambiguous margins. It is, in the end, the one lever in the entire problem that lies with the regulator rather than the enterprise, which is reason enough to name it precisely even though pulling it is beyond the generator's reach.

8. Financial responsibility and a Price-Anderson analogue

A sharper version of the same point emerges from comparison with the one mature U.S. regime that both prices and bounds a catastrophic environmental liability: the Price-Anderson Act, which governs nuclear-incident liability and originates in the same protective-system setting from which this paper's framework is drawn. Price-Anderson layers coverage in a defined structure [24]. Each reactor site must carry the maximum primary liability insurance commercially available, currently about $500 million; a mutualized *secondary* layer assesses every covered reactor a retrospective premium—on the order of $158 million each, roughly $16 billion in aggregate per incident—payable

only if the primary layer is exhausted; federal indemnification backstops Department of Energy contractors; and the industry's liability is *capped* at the resulting sum, above which further compensation becomes a matter for Congress. In the notation of Section 3, Price-Anderson bounds the listed-state cost c^f, pre-funds the catastrophe premium $(c^f - c^g)\,p$ through mandatory insurance sized to a defined maximum harm, and supplies a public backstop for the tail beyond it.

The waste-generator's regime is almost the inverse. CERCLA liability is uncapped, unchanneled, strict, joint-and-several, and retroactive; the generator, not merely the facility operator, is exposed; and the only mandatory coverage is the comparatively small RCRA financial-assurance and accident-liability requirement on the facility—on the order of $1 million per occurrence—which neither bounds a Superfund cleanup nor reaches the generator's PRP exposure. Strikingly, the enabling authority already exists: Section 108(b) of CERCLA directs the EPA to require financial responsibility for facilities handling hazardous substances "consistent with the degree and duration of risk." Requiring financial responsibility of this kind is the textbook regulatory answer to a judgment-proof injurer: where the party that controls the hazard may be unable to pay and the probability of an effective suit is low, ex ante regulation complements rather than replaces liability [25, 7]. Yet the agency has declined to exercise it for every industry it has studied—hardrock mining [26] and, subsequently, the chemical, electric-power, and petroleum sectors [27]. The Price-Anderson architecture is, in other words, authorized for the waste setting but unbuilt.

The obstacle is the one identified in Section 5: a Price-Anderson structure presupposes a quantified *maximum* harm to which primary insurance can be sized, and Superfund harm has no such ceiling. A licensing reform would supply it. Suppose each facility were licensed to hold no more than a defined maximum quantity of waste—specified by waste type and weighted by toxicity—and were required to carry primary insurance sized to the maximum credible release that licensed quantity implies. The licensed maximum then plays the role the site limit plays under Price-Anderson: it bounds the insurable loss, makes a mutualized secondary pool feasible, and gives the Section 108(b) requirement a number to attach to. A facility seeking to handle more would have to apply to re-license a higher quantity and *re-insure to the new maximum*, so that any growth in the exposure a site creates is matched, in advance, by growth in pre-funded coverage. The device also closes the loop with the allocation model of Section 4: the licensed ceiling bounds a site's total managed quantity Q_i, and hence the pool from which a generator's equitable share is drawn, converting an open-ended liability into

one whose maximum is known at the time the disposal decision is made.

How far such a scheme could be built within the existing regulatory structure, and how far it would require new legislation, is worth stating precisely, because the two layers of a Price-Anderson regime sit very differently with respect to current authority. The *funded-insurance* layer is largely reachable by administrative action: RCRA permits already license treatment, storage, and disposal facilities, specify their authorized waste codes, and limit unit design capacity, so attaching a toxicity-weighted maximum-quantity ceiling and tying it to a maximum-credible-release insurance requirement extends tools the permitting program already uses, while CERCLA's Section 108(b) supplies independent authority to require financial responsibility "consistent with the degree and duration of risk," and the Reportable Quantities of Section 302.4 and the Hazard Ranking System provide the toxicity and pathway scaffolding such a metric would need. The *liability-shaping* layer is another matter entirely. The cap, the channeling of liability to the insured facility, and the public backstop above it—the features that actually relieve the generator—run directly against CERCLA's foundational design of unlimited, strict, joint-and-several, retroactive liability reaching every contributor, and could not be created by regulation; like Price-Anderson itself, they would require an act of Congress. Two further limits follow from the same contrast: any such regime would be prospective, leaving the legacy exposure of Section 4 untouched, and because RCRA is delegated to authorized states while much audited capacity lies abroad, a national ceiling-and-insurance scheme would interact unevenly with state programs and would not reach foreign facilities at all. The realistic reading, then, is that the funded-insurance and licensed-quantity layer is within reach of existing authority—Section 108(b) already contemplates it—whereas the cap, channeling, and backstop that complete the analogy are a legislative choice rather than an administrative one.

The mechanism inherits the qualifications of the graduated-severity proposal above, and adds two of its own. Quantity bounds harm only imperfectly: consequence depends on toxicity, mobility, hydrogeology, and the proximity of receptors as much as on mass, so a licensed ceiling must be defined per waste type and conditioned on setting—consistent with the Reportable Quantities of Section 302.4 and the pathway and target factors of the Hazard Ranking System—rather than as a single tonnage. Chronic contamination from sustained sub-threshold releases is bounded by the stock at risk only in part, so monitoring and the existing financial-assurance layer remain necessary alongside the cap. Sharp ceilings invite the threshold behavior noted earlier—facilities sized to just below a limit, or split to stay under

it. And capping liability trades deterrence for certainty, a genuine cost in a regime whose unlimited liability is itself a deterrence instrument (Section 2); the cap must therefore be set high enough, and paired with mandatory coverage, to preserve the incentive to prevent releases rather than merely to insure against them. None of these is disqualifying. Taken together they suggest that the absence of a Price-Anderson structure for waste facilities is a policy choice rather than a technical necessity, and that a toxicity-weighted, setting-conditional licensed-quantity ceiling is the instrument that would make such a structure operational.

9. Discussion and limitations

The framework is stated for a waste generator facing U.S. Superfund liability, but its structure is general. Any enterprise whose exposure to a low-probability, high-consequence loss is governed by a party it does not control, and whose own lever is the allocation of business across counterparties of differing and uncertain risk, faces a version of the same problem; the expected-loss ranking of Section 3 and the forward allocation of Section 4 apply wherever a comparable liability regime and a usable base rate exist.

The treatment is complementary to the two CERCLA literatures it draws on rather than competing with them. The deterrence literature [1, 2, 3] asks whether liability induces efficient care by the party that controls the hazard; we take that party's behavior as given and ask how a different party, the generator, should act under the resulting uncertainty. The allocation literature [4, 6, 5] studies how a realized cleanup cost is divided after listing; we use its central object—the equitable share—as the consequence term in a decision made before listing. The contribution is to join these to a prospective, risk-informed decision rule and to the estimation problem that rule depends on. The policy argument of Sections 7 and 8 is, in turn, an instance of the liability-versus-regulation tradeoff: where the controlling party may be judgment-proof and the liability trigger is noisy, ex ante regulation is the natural complement to liability [25].

Several limitations bound the claims. The empirical contribution is constrained by the quality of public data: as Section 6.1 showed, the binding obstacle to a site-level estimate is population definition and label hardening, not the choice of method, so the quantitative results here are illustrative rather than predictive. The listing probability is a narrow lens on a broader exposure—it captures chronic-contamination listing under the U.S. program, and neither the reputational loss a firm may also bear nor the acute events

and non-U.S. facilities outside the Hazard Ranking System's reach—and is best read as one input to judgment rather than a verdict.

A further limitation is normative rather than empirical. The objective throughout is the *generator's* private expected cost, and a rule that minimizes it is silent on who bears the harm it does not internalize. Minimizing private loss can displace hazard onto third parties, most sharply when waste is routed to jurisdictions with weaker regulatory enforcement—including foreign facilities beyond the reach of the U.S. program—where exposed populations may bear uncompensated and involuntary risk, as when delivered waste is informally handled or scavenged. These distributive-justice questions lie outside the private cost calculus formalized here; we examine them separately [28, 29] and flag them so that the framework is not mistaken for a normative endorsement of cost-minimizing disposal.

The policy proposals of Section 7—graduated severity and a licensed-quantity ceiling—are normative, subject to the objections raised there, and, in the parts that would cap or channel liability, require legislation rather than agency action. The framework is therefore best understood as a way to organize a generator's decision, and to locate where better data and clearer rules would most improve it, rather than as a turnkey predictor.

10. Conclusion

The waste generator's position under CERCLA is unusual: it is liable for an outcome it does not control, under a rule its own diligence cannot change. This paper has taken that inversion as its starting point and asked what a generator can nonetheless do. The answer is to stop treating hazard scores as predictions and to rank disposal options by expected loss, in which a catastrophe premium prices the incremental cost of listing; to recognize that the controllable lever is the forward allocation of waste, which a risk-averse, multi-period program spreads so as not to become any facility's dominant customer; and to carry the listing probability not as a point but as a jurisdiction-conditioned distribution, because a large part of its uncertainty is administrative and beyond the firm's reach.

That same administrative uncertainty is the paper's policy theme. Because it is an artifact of how the listing rule is written rather than a fact of nature, greater regulatory specificity—a published, graduated, setting-conditional severity scale, and the toxicity-weighted licensed-quantity ceiling that would let a Price-Anderson-style financial-responsibility regime be built under authority CERCLA already grants—would convert much of it into uncertainty the generator can measure and price. The generator's problem

and the regulator's design are thus two halves of one question: the better the rule bounds and quantifies the harm, the more the prospective decision this paper formalizes can actually be made.

Author contributions

Ernie Kee: Conceptualization, Methodology, Formal analysis, Software, Investigation, Writing – original draft, Writing – review & editing, Visualization. **Karen Miller:** Conceptualization, Validation, Investigation, Resources, Writing – review & editing, Supervision. K. Miller contributed the regulatory and engineering domain expertise on which the framework rests: as a specialist in RCRA and CERCLA practice she supplied the statutory and compliance background of Section 2 and the working understanding of treatment, storage, and disposal facility (TSDF) site characteristics and audit findings that informs the risk covariates of Section 6 and the facility-review discussion throughout. **Hojun Choi:** Methodology, Validation, Writing – review & editing.

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The study uses only publicly available data. The facility, listing, and hazard-ranking records analyzed in the prototype derive from public U.S. Environmental Protection Agency sources—the National Priorities List, the RCRAInfo and ECHO compliance systems, and the CERCLA Section 302.4 list of hazardous substances—and no new data were generated. The derived datasets, data dictionary, and analysis code that support the prototype of Section 6.1 are openly available at `https://doi.org/10.5281/zenodo.21502158` (archived release) and `https://github.com/erniejlkee/prp-listing-data` (working repository).

Acknowledgements

The authors thank Arden Rowell (University of Illinois Urbana–Champaign) for helpful comments and, in particular, for insight into the divisibility analysis of *Chem-Dyne* and *Bell Petroleum* and its implications for apportioning a generator's equitable share; Janet L. Greenberg of GREEN Environmental Consulting, Inc., for her review and comments on the underlying work; and Jeffrey Sacre of CHWMEG for valuable observations on the practical structure of the disposal decision—transportation cost, facility waste-acceptance constraints, and the effect of diverted volume on facility solvency—and on the environmental-justice dimension of waste routing.

Declaration of generative AI and AI-assisted technologies in the writing process

During the preparation of this work the authors used a large language model based AI assistant (Anthropic Claude) to assist with drafting and language editing of the manuscript, with LaTeX and bibliography formatting, and with organizing source material. After using this tool the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

References

[1] Tom H. Tietenberg. Indivisible toxic torts: The economics of joint and several liability. *Land Economics*, 65(4):305–319, 1989.

[2] Lewis A. Kornhauser and Richard L. Revesz. Evaluating the effects of alternative superfund liability rules. In Richard L. Revesz and Richard B. Stewart, editors, *Analyzing Superfund: Economics, Science, and Law*, pages 115–144. Resources for the Future Press, Washington, DC, 1995.

[3] Kathleen Segerson and Tom H. Tietenberg. The structure of penalties in environmental enforcement: An economic analysis. *Journal of Environmental Economics and Management*, 23(2):179–200, 1992.

[4] Richard L. Revesz and Richard B. Stewart. *Analyzing Superfund: Economics, Science, and Law*. Resources for the Future Press, Washington, DC, 1995. edited volume.

[5] Howard F. Chang and Hilary Sigman. The effect of joint and several liability under superfund on brownfields. *International Review of Law and Economics*, 27(4):363–384, 2007.

[6] Howard F. Chang and Hilary Sigman. An empirical analysis of cost recovery in superfund cases: Implications for brownfields and joint and several liability. *Journal of Empirical Legal Studies*, 11(3):477–504, 2014.

[7] Steven Shavell. The judgment proof problem. In *Foundations of Insurance Economics*, pages 341–354. Springer, 1986.

[8] E Kee, M Wortman, V Moiseytseva, F Yilmaz, and D Johnson. Protective systems: Margins of safety, regulatory authority, and the calculus of negligence. In *American Nuclear Society (ANS) International Topical Meeting on Probabilistic Safety Assessment and Analysis (PSA), Pittsburgh, PA*, pages 1366–1370, 2017. ISBN 978-1-51085-180-1.

[9] Martin Wortman, Ernie Kee, and Pranav Kannan. The Role of Protective System Reliability Analysis in the Study of System Safety. In *ASME International Mechanical Engineering Congress and Exposition*, volume 85697, page V013T14A032. American Society of Mechanical Engineers, 2021.

[10] U.S. Congress. Comprehensive environmental response, compensation, and liability act of 1980. 42 U.S.C. §§ 9601–9675, 1980.

[11] Kate R. Bowers. Liability under the comprehensive environmental response, compensation, and liability act (cercla). CRS Report IF11790, Congressional Research Service, 2021.

[12] U.S. District Court, S.D. Ohio. United states v. chem-dyne corp. 572 F. Supp. 802, 1983.

[13] U.S. Supreme Court. Burlington northern & santa fe railway co. v. united states. 556 U.S. 599, 2009.

[14] U.S. Court of Appeals, Fifth Circuit. In re bell petroleum services, inc. 3 F.3d 889, 1993.

[15] U.S. Congress. Resource conservation and recovery act of 1976. 42 U.S.C. §§ 6901–6992k, 1976.

[16] Steven Shavell. *Economic Analysis of Accident Law*. Harvard University Press, Cambridge, MA, 1987.

[17] Stanley Kaplan and B John Garrick. On the quantitative definition of risk. *Risk analysis*, 1(1):11–27, 1981. doi: https://doi.org/10.1111/j.1539-6924.1981.tb01350.x.

[18] Nathan O Siu and Dana L Kelly. Bayesian parameter estimation in probabilistic risk assessment. *Reliability Engineering & System Safety*, 62(1-2):89–116, 1998.

[19] R. M. Cooke. *Experts in Uncertainty: Opinion and Subjective Probability in Science*. Oxford University Press, New York, 1991.

[20] R. T. Clement and R. L. Winkler. Combining probability distributions from experts in risk analysis. *Risk Analysis*, 19:187–203, 1999.

[21] C. Genest and J. V. Zidek. Combining probability distributions: A critique and annotated bibliography. *Statistical Science*, 1:114–148, 1986.

[22] A. O'Hagan, C. E. Buck, A. Daneshkhah, J. Eiser, P. Garthwaite, D. Jenkinson, J. Oakley, and T. Rakow. *Uncertain Judgements: Eliciting Expert Probabilities*. Wiley, Chichester., 2006.

[23] M. A. Meyer and J. M. Booker. Eliciting and Analyzing Expert Judgment: A Practical Guide, NUREG/CR-5424, U.S. Nuclear Regulatory Commission, 1990.

[24] Mark Holt. Price-anderson act: Nuclear power industry liability limits and compensation to the public after radioactive releases. CRS Report IF10821, Congressional Research Service, 2025.

[25] Steven Shavell. Liability for harm versus regulation of safety. *The Journal of Legal Studies*, 13(2):357–374, 1984.

[26] U.S. Environmental Protection Agency. Financial responsibility requirements under CERCLA section 108(b) for classes of facilities in the hardrock mining industry. Final agency action, Federal Register, 2018.

[27] U.S. Environmental Protection Agency. Financial responsibility requirements under CERCLA section 108(b) for facilities in the electric power generation, transmission, and distribution; petroleum and coal products manufacturing; and chemical manufacturing industries. Final agency actions, Federal Register, 2020.

[28] Kevin Nelson and Ernie Kee. Concerns with risk distribution in environmental justice. In *ASME International Mechanical Engineering Congress and Exposition*, volume 87707, page V013T15A004. American Society of Mechanical Engineers, 2023. doi: 10.1115/IMECE2023-114497. URL https://asmedigitalcollection.asme.org/IMECE/proceedings/IMECE2023/87707/V013T15A004/1196294.

[29] Riley Fisher, Ernie Kee, David Johnson, Ha Bui, and Zahra Mohaghegh. Engineering & environmental justice: Protections, hazards, and technological systems. In *ASME International Mechanical Engineering Congress and Exposition*, volume 87707, page V013T15A002. American Society of Mechanical Engineers, 2023.