

Children as Canaries:

Disaggregating the Sentinel Claim into Three Mechanisms of Early Harm

Travis Gilly

ORCID: 0009-0007-2954-6313

Real Safety AI Foundation

Illinois, United States

Working paper, July 2026 (v2)

ABSTRACT

The claim that children function as canaries for harms later reaching adults is intuitive, rhetorically powerful, and used loosely. It is applied indiscriminately to cases sharing only the fact that children were hurt. This paper does not argue that every harm announces itself in children first; harms attaching to adult roles have no pediatric precursor and none should be sought. It argues that a real and recurring pattern in the documented record is being obscured by the imprecision, and that the sentinel claim collapses three mechanisms with different causal structures, different evidence, and different remedies. Mechanism one, sentinel proper, describes harm registering first in children through developmental susceptibility and becoming legible in adults only later: leaded gasoline, where 170 million Americans were exposed as children and roughly 824 million cumulative IQ points and 151 million excess mental disorders arrived decades after the pediatric signal; the Great Smog, where prize cattle sickened before the human toll was counted; radioiodine concentrated through milk at Nevada and Hanford; and swill milk in New York. Mechanism two, targeted selection, describes harm in which children were chosen because they were children: the devshirme, which took roughly 200,000 boys; the Mamluk system, which excluded its own elite's sons and so required perpetual purchase; Canada's residential schools, through which more than 150,000 children passed and at least 4,200 died; Willowbrook; Fernald; and the North Carolina Eugenics Board, which sterilized more than 7,600 people, some as young as ten. Fernald establishes the decisive point: a federal advisory committee found the radiation doses unlikely to have caused physical injury and concluded the studies wronged the boys regardless. Mechanism two is not a dose story. Mechanism three, captive absorption, describes children neither selected nor more susceptible who could not leave: asbestos in roughly 34,800 schools, lead in public housing, Flint, and 4 million lead service lines still in the ground. Only the first is a sentinel effect. The second and third are childism. The paper closes on a recourse asymmetry running from North Carolina, where fewer than 220 of more than 7,600 sterilized people were ever compensated, to the contemporary platform cases, where the barriers defeating child plaintiffs are general rather than child-specific, and therefore will not yield to child-specific advocacy.

Keywords: harm blindness, children's environmental health, history of childhood, childhood studies, sentinel populations, human subjects research ethics, vulnerable populations, environmental justice, childism, developmental vulnerability, lead exposure, residential schools, platform accountability

I. INTRODUCTION

The canary in the coal mine is a good metaphor and a bad one. It is good because it captures something real: a smaller body with a faster metabolism registers a poison before a larger body does, and the bird's distress is information about the air the miners are breathing. It is bad because the miners put the bird in the cage on purpose. The metaphor smuggles in an intentionality that most of the cases it gets applied to do not have, and it obscures the difference between a bird that detects a hazard and a bird that was fed to one.

That ambiguity costs nothing in casual usage. It becomes expensive the moment the claim is asked to do analytical work. Saying children are the canaries of a harm implies a testable structure: the harm appears in children first, it appears in adults later, and the interval between those appearances is a window in which intervention was possible and did not happen. If that structure holds, the failure to act on the pediatric signal is an indictment, because the information existed and was not used. If it does not hold, and the children were hurt for some other reason, the canary framing is a category error that invites an easy rebuttal and takes the surrounding argument down with it.

This paper does not argue that every harm announces itself in children first. That claim would be absurd and no one has made it. Harms attaching to adult roles, to occupational exposure, to purchasing capacity, or to legal capacities minors do not possess do not register in children ahead of adults, and no such early signal should be looked for. The argument here is narrower and more useful: across a documented historical record, a specific pattern recurs in which children registered a hazard first, the signal was available, and it was not acted on until the adult burden arrived decades later. That pattern is real, it repeats, and it deserves a precise name.

Naming it precisely requires separating it from two patterns it is constantly confused with. The claim as usually deployed collapses three mechanisms with different causal structures, different evidentiary requirements, and different remedies. Only one is a sentinel effect. The other two are childism, and they are better grounded under that name. The disaggregation is the contribution of this paper, and the historical record is what forces it.

II. THE SENTINEL CLAIM AND ITS CONFUSIONS

Stated carefully, the sentinel claim is epidemiological. A sentinel population is one in which a hazard becomes detectable earlier, or at lower exposure, than in the general population, because of a property of that population that amplifies the signal. For children the amplifying properties are well characterized: higher intake relative to body mass, greater gastrointestinal absorption, an incompletely formed blood-brain barrier, rapid neurological development creating critical windows in which a small insult produces a permanent deficit, and a longer remaining lifespan across which a latent effect can express.

Those are physiological facts and they generate a real prediction. If a neurotoxicant enters the shared environment, and children encounter it at roughly the ambient concentration adults encounter, the effect will be measurable in children before adults and at lower doses. That is the sentinel effect. It is a claim about detection order and it is falsifiable.

Confusion begins when the label extends to cases where children were harmed but detection order is not the assertion. Two large families exist, and the record is full of both.

In the first, children were selected. Nothing about pediatric physiology explains why the devshirme took boys, why residential schools took children, why Willowbrook infected institutionalized children with hepatitis, or why the North Carolina Eugenics Board sterilized girls as young as ten. What explains it is that children were legally powerless, institutionally available, and in several cases chosen precisely because a person can be shaped before they are formed. Calling these children canaries reverses the moral structure. A canary detects a hazard not aimed at it. These children were aimed at.

In the second family, children were neither selected nor unusually susceptible. They could not leave. Asbestos in school buildings exposed children because the law required them in the building. Lead in public housing exposes children because a child cannot decline a lease. Flint exposed children because a child does not choose a water utility. The exposure is ambient and children absorb a disproportionate share, but the mechanism is captivity, not detection. The child is not a sensor. The child is a resident of a place the adults could leave.

These mechanisms coexist in single cases, which is part of why they collapse. Lead is both sentinel and captivity: children registered the neurotoxicity first, and the children who registered it most were in housing they could not leave. But they remain separable, and separating them matters because the remedies diverge. A sentinel failure is remedied by surveillance and by acting on the pediatric signal. A targeting failure is remedied by prohibition and by dismantling the institution that selects. A captivity failure is remedied by removing the hazard from the places children are required to be, because the children cannot be removed.

III. MECHANISM ONE: SENTINEL PROPER

A. Leaded gasoline

Tetraethyl lead entered gasoline in the 1920s to reduce engine knock, though lead's toxicity was already established and workers making the additive were poisoned and died. In 1924 a cluster of workers at a Standard Oil plant in New Jersey died or suffered severe neurological damage while producing it, newspapers called the product loony gas, and a brief federal review allowed sales to continue. General Motors, Standard Oil, and DuPont created the Ethyl Corporation to market it, and industry figures publicly defended its safety while cheaper and less toxic anti-knock options such as ethanol blends existed (RSAIF MODERN-054).

What makes this the paradigm sentinel case is that detection order is documented rather than inferred. Clinical pediatric lead poisoning was described in the medical literature by 1904 (RSAIF CONTEMPORARY-057). Subclinical cognitive effects in children were reported through the 1970s in populations where ambient exposure was identical for adults on the same streets. The pediatric signal was legible while the adult signal was not. That is what the physiology predicts and what the sentinel claim asserts.

The retrospective accounting is enormous. McFarland, Hauer, and Reuben estimate that more than 170 million Americans alive in 2015, over half the population, had early-childhood blood lead high enough to cause lasting harm, costing roughly 824 million cumulative IQ points; the average loss was about 2.6 points per person, but those born in the 1960s and 1970s, when leaded fuel peaked, lost an estimated six or more points each (McFarland et al., 2022; RSAIF MODERN-054). McFarland and Reuben subsequently estimated that by 2015 the population had gained 602 million General Psychopathology points from childhood lead exposure arising from leaded gasoline, a 0.13 standard deviation increase in liability to mental illness, corresponding to approximately 151 million excess mental disorders, with a 0.64 standard deviation increase in internalizing symptoms, a 0.42 standard deviation increase in attention-deficit and hyperactivity symptoms, and effects most pronounced for cohorts born between 1966 and 1986 (McFarland and Reuben, 2024).

Two features deserve isolation. The exposure measured is childhood; the outcome measured is adult mental health. The harm registered in children and expressed across an adult population decades later. That is the sentinel structure as a finding rather than a metaphor. And the authors state explicitly that the estimate assumes the published lead-psychopathology associations are causal rather than purely correlational (McFarland and Reuben, 2024), an assumption that should be carried forward rather than dropped when the number is quoted.

Independent designs corroborate. Schwaba, Bleidorn, and Hopwood linked historic atmospheric lead from 269

United States counties and 37 European nations to personality data from more than 1.5 million people, finding adults raised in higher-lead counties had less adaptive personality profiles, with the Clean Air Act phase-out as a natural experiment; they report population blood lead fell 93.6 percent between 1970 and 2016, and that the average adult carried 13 micrograms per deciliter in the 1970s, roughly twice what now triggers clinical attention (Schwaba et al., 2021). They note their natural-experiment findings were not fully discriminable from cohort effects, a limitation they state themselves. Lee, Lee, and Warren found older adults who lived as children in cities with the water chemistry required for lead to leach had worse late-life cognition, though not steeper decline (Lee et al., 2022).

The exposure is not historical. Roughly 29 million United States housing units contain lead-based paint hazards, about 2.6 million of them occupied by young children (RSAIF CONTEMPORARY-057). Approximately 590,000 children between one and six had blood lead above 5 micrograms per deciliter in 2016, an estimated 9.2 million lead distribution pipes remain, and non-Hispanic Black children carry 2.8 times the likelihood of elevated blood lead compared with white peers (Pickering et al., 2026). The World Health Organization attributes more than 3.5 million deaths in 2023 to lead exposure from all sources, primarily through cardiovascular disease (RSAIF CONTEMPORARY-057). The United Nations Environment Programme estimates that ending leaded petrol will prevent more than 1.2 million premature deaths annually (RSAIF MODERN-054).

A detail from the corpus record deserves preservation. The additive was the work of Thomas Midgley Jr., who also developed the chlorofluorocarbons later found to destroy the ozone layer, giving a single chemist an outsized role in two global harms (RSAIF MODERN-054).

B. The Great Smog, and the actual canary

In December 1952 a cold spell drove Londoners to burn more coal while an anticyclone produced a temperature inversion trapping still air near the ground. Smoke and sulfur dioxide combined with fog into a dense acrid smog that held for five days. About 4,000 people died in the immediate episode and around 100,000 more were made ill; later analysis attributed roughly 8,000 additional deaths, first assigned to influenza, to the pollution, bringing the total to approximately 12,000 (RSAIF MODERN-188).

The corpus record preserves a detail that the metaphor's users have missed. Prize cattle brought into the city for the Smithfield agricultural show sickened, and some had to be destroyed, an early sign of how toxic the air had become (RSAIF MODERN-188). An actual animal sentinel appears in the paradigm human-sentinel case, and it functioned exactly as the metaphor promises: a differently sized body registered the hazard while the humans around it were still counting the fog as weather.

The human mortality concentrated at both ends of the age distribution. The very young and the elderly died most, while working-age adults breathing the same air survived. This matters for the argument. The mechanism is developmental and situational vulnerability rather than childhood as such, and a framework that only ever finds children is not detecting anything, it is projecting. Later study also suggested a higher risk of childhood asthma among those exposed in infancy or before birth (RSAIF MODERN-188), which extends the latency structure into the exposed cohort's adulthood.

One further feature of the Great Smog belongs to the sentinel argument. The true toll was established only through later study of mortality records (RSAIF MODERN-188). The signal was there in real time and was not read.

The reattribution deserves to be stated carefully, because it is this paper's own standard applied to its own case. The initial government position was that an influenza epidemic explained the elevated mortality in the months following the smog. Retrospective assessment, generating influenza-death estimates by several independent methods, found that only a fraction of those deaths could be attributed to influenza, and that accounting for the remainder would require

an epidemic on the order of twice the case-fatality rate and four times the incidence observed in general practice that winter (Bell, Davis, and Fletcher, 2004). The counter-explanation was not dismissed. It was tested, and it failed. That distinction matters here more than anywhere else in the paper, because every case in this corpus carries an official alternative account, and the Great Smog is the one where somebody did the arithmetic on it.

C. Radioiodine through milk

From 1951 the Atomic Energy Commission conducted atmospheric detonations at the Nevada Test Site, and prevailing winds carried fallout across Nevada, Utah, Arizona, and neighboring states. Reconstruction studies established that milk consumption was the dominant exposure pathway and that children who drank fresh local milk carried the highest thyroid doses (RSAIF MODERN-156). Federal officials publicly minimized the hazard while testing proceeded, and the reassurance did not match the internal understanding of iodine-131 and its concentration in the thyroid, particularly in children.

Hanford ran the same pathway. In December 1949 the Green Run deliberately released iodine-131 from fuel cooled only about sixteen days, to calibrate equipment for detecting Soviet plutonium production. Planners expected roughly 4,000 curies; approximately twice that was emitted. Vegetation sampled near Kennewick registered iodine-131 at about 600 times the level then considered tolerable. Public-health materials record that drinking contaminated milk was the major pathway by which iodine-131 reached people, concentrating in the thyroid and especially affecting children (RSAIF MODERN-160).

These are mechanism-one cases in clean form, and they refine the criteria usefully. The amplifying property is not body mass or the blood-brain barrier. It is a dietary pattern, children drink more milk, interacting with an organ in a particular developmental state, a small growing thyroid avid for iodine. No one aimed radioiodine at children. The children were not selected and were not confined anywhere adults were not.

Chernobyl ran the pathway a third time, on another continent and in another decade, and the independent replication is what converts a pattern into a mechanism. Iodine-131 intake accounted for more than 90 percent of the thyroid dose across the affected regions, contaminated cow's milk was the dominant route, and individual thyroid doses reached 42 Gy, varying with the age of the person and with milk consumption habits (Drozdovitch, 2021). The amplification is measured rather than asserted. Age-dependent thyroid dose factors for iodine-131 run from 36 in newborns to 1.4 in adults, so one release delivers roughly twenty-six times the dose to the youngest body it reaches (Reiners, 2011). The epidemiology follows the arithmetic. Childhood thyroid cancer incidence in Belarus rose from roughly 0.5 per million before the accident to roughly 40 per million in girls and 20 per million in boys, while no significant increase appeared among those exposed above the age of twenty (Reiners, 2011). The same record notes that iodine blocking was not offered to the population and children were not prohibited from drinking contaminated milk (Reiners, 2011). That places Chernobyl beside Nevada and Hanford in the sense this paper cares about. The signal was legible in advance, the protective step was available, and it was not taken.

The compensation record shows what happens when a sentinel signal is finally conceded. The Radiation Exposure Compensation Act arrived in 1990, decades after exposure, offering fixed administrative payments in place of contested tort claims; the July 2025 reauthorization raised downwinder payments to 100,000 dollars and widened eligibility to entire states, including New Mexico, that earlier versions had excluded (RSAIF MODERN-156). That exclusion is itself part of the documented harm, because families exposed by the same weather systems were for decades told they lived outside the compensable line. No verified count of fatal cancers attributable to Nevada fallout exists, and the Hanford Thyroid Disease Study produced results many families found inconclusive (RSAIF MODERN-160).

D. Persistence across generations: Agent Orange

Between 1961 and 1971, Operation Ranch Hand sprayed nearly 19 million gallons of herbicide over Vietnam, including at least 11 million gallons of Agent Orange, to strip forest cover and destroy crops (RSAIF CONTEMPORARY-062). The active problem was a contaminant. Agent Orange carried dioxin, which does not break down quickly, and it accumulated in soils, sediments, the food chain, and human tissue, concentrating at the airbases where the herbicide was stored, mixed, and loaded.

Agent Orange extends the sentinel structure along an axis the other cases only gesture at: the exposed cohort and the affected cohort are separated not by years but by a generation. The Vietnamese Red Cross has estimated that about 3 million people were affected, including at least 150,000 children born with serious birth defects (RSAIF CONTEMPORARY-062). This paper carries those figures as the advocacy-linked estimates the corpus record says they are, not as a settled count, and the reason to carry them at all is structural rather than quantitative. Dioxin's persistence means a parent's exposure becomes a child's developmental injury, and the child registers a harm the parent absorbed. That is the same geometry the maternal-transmission drug cases would show, arriving here through soil and the food chain rather than through the placenta. The transmission has a documented biological route. Exposure to dioxin during the window in which primordial germ cells undergo epigenetic reprogramming can carry an environmental condition into the next generation through the male germline, and the raised incidence of birth defects among the offspring of veterans exposed to Agent Orange is consistent with that pathway (Pilsner, Parker, and Sergeyev, 2017). The route is worth naming precisely because it does not require the child to have been anywhere near the herbicide.

The case also marks a boundary the sentinel argument has to respect. The corpus record is explicit that attributing specific illnesses and birth defects to dioxin across large populations and long periods is difficult, while the measured contamination at particular airbases and the list of conditions recognized for exposed veterans rest on firmer ground (RSAIF CONTEMPORARY-062). Two kinds of evidence sit side by side, one firm and one contested, and an honest sentinel argument reports both rather than collapsing them.

E. Swill milk, and the limits of the mechanism

The milk pathway has an older case that tests the boundary. As New York expanded before the Civil War, distilleries multiplied and dairymen built stables beside them to use the hot grain mash. Cattle were confined by the hundreds in filthy sheds, fed distillery slop, and kept in production as disease set in. The milk was thin and bluish, so dealers whitened it with plaster of Paris, thickened it with starch and eggs, and tinted it with molasses. It reached poor families who could not afford country milk, and it was fed above all to infants (RSAIF MODERN-204).

The most-cited figure, that as many as 8,000 infants a year died in the city from causes tied to contaminated milk, comes from an 1858 New York Times editorial. The corpus record is explicit that this was an alarm raised by journalists and reformers, not a verified toll, and that no registration system of the period could isolate swill milk as a cause (RSAIF MODERN-204). The paper carries it with that caveat attached, because a sentinel argument that quotes an advocacy number as a finding has forfeited the standard it is asking others to meet.

The mechanism assignment is also mixed, and instructively so. Infants were most exposed, but the corpus record states why: cheap loose milk was what the poorest households could buy (RSAIF MODERN-204). That is economic captivity, not physiological susceptibility. The case belongs to mechanism three at least as much as to mechanism one, and pretending otherwise would be exactly the error this paper is written against.

The scholarly record on the urban milk supply supports the corpus account and extends it. Brodie (1991) identifies the structural cause: as cities grew, dairy farms were pushed many miles from the consumers of milk, and uncooled milk carried over that distance through the hot summer months raised the chance of milk-borne disease reaching infants. The

hazard was a product of urban form and transport, and it was encountered by whoever drank the milk, which satisfies the condition of ambient exposure that mechanism one requires.

The remedy history carries the sharper lesson, and it is the reason this case earns its place despite failing the assignment test. Bhatia, Krieger, and Subramanian (2019) trace what followed. The first milk stations opened in New York in 1893, funded by the philanthropist Nathan Straus, and by 1910 there were 297 stations across 38 cities. In 1912 the New York City Health Department mandated pasteurization of all milk coming into the city, well ahead of comparable measures elsewhere, and federal legislation made pasteurization mandatory nationwide in the 1920s. The authors report Lee's 2007 finding that the decline in infant mortality was inversely correlated with the cleaning of the market milk supply between 1840 and 1940, across a period that also saw breastfeeding decline and offered no medical treatment for infantile diarrhea. The intervention worked.

What it also did is the part this paper needs. Quoting Nathanson, Bhatia and colleagues observe that constructing infant mortality as a problem of bad milk was attractive to public health officials because it promised a simple prophylactic against infant deaths, obviating the need for the fundamental environmental and behavioral reforms that had proved so difficult to accomplish. They report Meckel's account of the sequel: once the structural milk interventions were in place, infant welfare activity shifted from milk reform to maternal reform, and the physicians of the period divided between those who called for sanctions on mothers who did not nurse and those who argued that insufficient family income, rather than laziness or ignorance, was what ended breastfeeding.

That is a failure mode the sentinel argument has to survive, and it has nothing to do with detection. A pediatric signal can be read correctly, acted on, and still serve as a substitute for the structural change the signal was evidence of. The milk was cleaned. The tenements were not. Swill milk therefore tests the mechanism twice: once on assignment, where economic captivity competes with physiological susceptibility, and once on remedy, where a real success absorbed the political energy of the question it had opened. Both tests are worth failing in public, which is why the case stays in.

F. What makes a case a sentinel case

The cases supply the criteria. A case belongs to mechanism one when four conditions hold. The exposure is ambient or incidental rather than directed, so children and adults encounter the hazard without anyone choosing children. The pediatric effect is detectable before the adult effect, and the gap is attributable to a specified physiological or behavioral property rather than to differential exposure. The adult effect subsequently materializes, which converts the pediatric signal from a curiosity into a warning that was ignored. And the interval is long enough that acting on the signal would have prevented the adult burden.

Applying those criteria narrows the field, and the narrowing is the point. The discipline of excluding cases that fail them is what gives the remainder its force.

The second condition carries most of the weight, and it is not this paper's invention. Environmental health already has a name for the property it points at. A window of susceptibility is a bounded developmental interval during which an exposure produces an effect it would not produce outside that interval, and the construct is established well enough to have generated statistical machinery built specifically to locate such windows in cohort data (Liu, Bobb, and Lee, 2017). It has been mapped across the life course for breast cancer, where prenatal development, puberty, pregnancy, and the menopausal transition each behave as distinct windows (Terry, Michels, and Brody, 2019), and applied to inorganic arsenic, where the fetus and young children are identified as particularly sensitive subpopulations (Bommarito and Fry, 2016).

The same literature disciplines the claim, and the discipline is worth carrying. Dentine analysis of early-life manganese exposure found that timing rather than dose drove the association, with prenatal levels behaving differently

from postnatal ones and the direction of effect reversing between them (Claus Henn, Austin, and Coull, 2018). A sentinel argument that treats more exposure as straightforwardly worse is reaching past its evidence. What the mechanism licenses is narrower and stronger than that: at a specifiable time, in a specifiable system, a given exposure does something to a child that it does not do to an adult.

The sentinel concept itself is also not this paper's invention, and confronting the existing literature does more for the argument than borrowing its metaphor did. Environmental and public health has a formal apparatus for sentinel surveillance, built on three stated conditions: the sentinel shares an environment with the population of concern, the sentinel responds to the agent in an analogous way, and clinical signs manifest in the sentinel before they manifest in that population (Gubernot, Boyer, and Moses, 2008). Those conditions are close cousins of the four this paper derived from its cases, which is reassuring rather than embarrassing. The literature also supplies the founding example. The iconic canary is invoked directly in the field, alongside the ataxic dancing cats of Minamata, a warning sign that was not recognized in time to prevent the human epidemic, and the 1962 lead poisoning of cattle and horses near a Minnesota smelter, which did prompt the state health department to survey the surrounding human population (Reif, 2011). The animal sentinel record already contains both outcomes this paper is tracking: the signal read, and the signal missed.

More useful is the reason the field gives for why sentinels work, because the paper has been asserting it without knowing the argument existed. Animal sentinels are valuable in part because of their relative freedom from concurrent exposures, from confounding, and from exposure misclassification; in human studies, smoking, alcohol, and occupational exposure can obscure the effect of a community exposure (Reif, 2011). Restricted daily mobility over a shorter lifespan makes their exposure assessment more accurate still. Read that list against a child. A child does not smoke, does not drink, holds no job with its own exposures, and does not range far from home. The pediatric signal is not cleaner because childhood is delicate. It is cleaner because a child's exposure history is short, local, and uncontaminated by the adult confounders that make community effects hard to see in the population that can vote.

The literature also disciplines the paper twice. A systematic review of animal sentinel studies found no evidence-based guidelines for using sentinel data in human health decisions, significant limitations in the methods employed, and clear linkages to human health frequently absent altogether (Rabinowitz, Gordon, and Holmes, 2005). Even in the paradigm case, the inference from sentinel to population is more often asserted than established. That locates what is on offer here with more precision than a claim to originality would. The four conditions overlap the field's three, and the overlap is the point: they are not new. What the field reports it does not have is a rule for deciding when a sentinel signal should move a decision about the population, and that is what the conditions, together with the assignment procedure that follows them, are for. The same review reports that animal sentinel studies were difficult to locate in the literature at all, because the concept had no adequate index term. A category of evidence that has no keyword is hard to accumulate, which is a visibility gap operating on the scholarship rather than on the exposed.

The second discipline is the harder one, and it goes to this paper's title. The working definition in the field is that a sentinel is any non-human organism that reacts to an environmental contaminant before the contaminant reaches people (Gubernot, Boyer, and Moses, 2008). Non-human, by construction. The concept was built to describe something whose distress is informative rather than important on its own account, and that is the whole reason the exclusion is there. To call a child a canary is therefore to apply a term whose definition places the bearer outside the population whose harm counts. That is not a defect of the metaphor. It is mechanism two, arriving through the vocabulary rather than through an institution, and it is the reason this paper's first section is about the confusions in the claim rather than about its power.

IV. MECHANISM TWO: TARGETED SELECTION

The second family is not a detection story. It is a selection story, and the instrument comes from childism rather than epidemiology.

A. Institutional capture

The devshirme took Christian boys, mostly from Balkan villages, removed them from their families, converted them, and trained them for military, palace, or administrative service. Roughly 3,000 boys were taken annually at its height, drawn from about one in forty Christian households in the levied areas; a surviving register from 1603 to 1604 records 2,604 boys taken in a single levy; and estimates put the total at approximately 200,000 boys between about 1380 and 1648 (RSAIF MEDIEVAL-013). Families sometimes hid or maimed their sons to protect them. The ablest rose very high, becoming Janissary officers, provincial governors, admirals, or grand viziers, and defenders of the system pointed to that mobility. The foundation was compulsion.

The Mamluk system purchased boys from non-Muslim regions, converted them, trained them in military households, and incorporated them into armed service; the Arabic term means owned. The sultanate ruled Egypt and Syria for 267 years (RSAIF ANCIENT-106). Its structural feature is the one that matters here: the state reproduced itself chiefly through new purchases rather than hereditary succession, since sons of mamluks were generally excluded from the elite military caste. That choice tied state power directly to slave markets and made continued acquisition of children a permanent requirement of the regime rather than an occasional practice (RSAIF ANCIENT-106). A system that must keep buying children in order to exist did not stumble into harming them. The corpus record states the analytical point precisely: the prestige of the outcome and the coercion of the entry point existed together throughout, and advancement does not cancel coercion.

The historiography presses harder on that point than the corpus record does, and the paper is obliged to say so. Ottoman archival evidence puts the devshirme total at roughly 200,000 across three centuries, which matches the corpus figure and works out to something on the order of 667 children a year across the entire region, far below the rate at the system's height and far below what Balkan nationalist historiography has claimed (Kirmizialtin, 2007). The same survey records that historians treat the levy as a major avenue of social advancement, that Bosnian and Albanian Muslims petitioned to have their own sons recruited even though the levy was formally restricted to non-Muslims, and that Ottoman records contain instances of non-Muslim families bribing officials to add a son's name to the list (Kirmizialtin, 2007). Some families competed for the selection.

That is a harder fact than the corpus record anticipates, and it does not dissolve the mechanism. It sharpens it. Mechanism two identifies children selected as a resource by an institution, and an institution's success in making its selection desirable is a property of the institution rather than a defense of it. What the historiography does establish is the fact the mechanism requires: the child levy was the one form of conversion the Ottoman state unquestionably enforced as central policy, in a field where the extent of forced conversion is otherwise contested (Kirmizialtin, 2007). The boys were legally slaves of the sultan and could nonetheless rise to the highest offices of the state (Wiesner-Hanks, 2008). Both were true simultaneously, and a framework that can hold only one of them is not describing this case.

The method is documented in the same terms the framework uses. Children were removed from their familiar cultural and family environment, converted, educated in dedicated institutions, and then placed in military bodies or in the court and bureaucracy, a system of slave administration and slave soldiery with a continuity running from the earliest Turkish presence in Anatolia into the Ottoman period proper (McNair, 2010). The Mamluk structure confirms the acquisition requirement from the other side. On the traditional account, the military-slave elite monopolized political authority for more than 250 years while excluding every other social group, including its own descendants

(Van Steenberg, Wing, and D'hulster, 2016). An elite that bars its own children from inheriting has to obtain someone else's.

B. Assimilation by removal

Three national systems ran the same logic with the purpose stated openly, and together they establish that mechanism two is a policy instrument rather than an aberration.

Canada's residential schools took more than 150,000 First Nations, Inuit, and Metis children. The National Centre for Truth and Reconciliation has documented at least 4,200 who died while attending, and the Truth and Reconciliation Commission concluded the real number is higher, potentially far higher, because of poor record-keeping, unnamed burials, and destroyed files (RSAIF COLONIAL-118). From 1920 an amendment to the Indian Act made attendance compulsory, giving removals the force of law. The last federally run school closed in 1996.

The United States operated or supported a parallel system from 1819 to 1969. The Department of the Interior's 2024 final report documented at least 973 Native children who died, identified at least 74 marked and unmarked burial sites, and updated the official list to 417 federal boarding schools across 37 states or then-territories. It estimated that the United States made available more than 23.3 billion dollars in inflation-adjusted 2023 terms for the boarding-school system and related assimilation policies between 1871 and 1969, and identified more than a thousand additional institutions advancing similar goals (RSAIF MODERN-139).

Australia removed Aboriginal and Torres Strait Islander children under protection and welfare laws that made officials their legal guardians. The national inquiry concluded that between one in three and one in ten Indigenous children were forcibly removed between about 1910 and 1970, with total estimates commonly ranging from about 10,000 to 100,000, a spread that reflects records kept inconsistently across many agencies and decades (RSAIF MODERN-176). The selection criterion was explicit: policy often targeted children of mixed descent, on the belief that they could be absorbed into white society while Indigenous identity was suppressed. The 1997 *Bringing Them Home* report described the removals as a gross violation of human rights and argued they met the international definition of genocide. The inquiry also found that few Indigenous families escaped the impact.

Notice what the Canadian compulsory-attendance amendment does to the mechanism. It converts a targeting case into one that also runs through captivity, and the conversion was a legislative choice made in 1920.

The economic history supplies something the moral record does not, and it is uncomfortable in a way worth stating plainly. Causal analysis of the Canadian schools finds that the average school had substantial effects on both cultural and economic assimilation (Feir, 2016). The policy achieved the aim its architects stated. That finding does not soften the case; it is the case. Mechanism two describes children used as instruments toward an end, and an instrument that works is still an instrument. The same analysis reports suggestive evidence that the effect does not hold for the most abusive schools, which is recorded here rather than omitted, because a paper that only cites the findings flattering to its thesis is running the selection it claims to diagnose.

What the policy did not do is stay inside the generation it processed. Despite the human-capital gains recorded among attendees, subsequent generations show lower educational attainment, reversing the ordinary relationship between parents' and children's human capital in a pattern consistent with cultural detachment and the breakdown of family relationships (Jones, 2024). The instrument outlasted its operators.

C. The floor that keeps rising

These three systems share a property that no sentinel case has, and it is a methodological finding rather than a moral one.

Canada's documented death total of at least 4,200 has continued to grow as previously unexamined records are reviewed, and the Commission treats it as a floor rather than a ceiling (RSAIF COLONIAL-118). The United States figure moved from a lower number in the 2022 first volume to at least 973 in the 2024 final volume, and the institution count from 408 to 417, as additional records were located; the report describes the death total as a minimum and researchers expect the true number to be considerably higher (RSAIF MODERN-139). At Tuam, the 796 recorded child deaths were established not by a state inquiry but by a local historian, Catherine Corless, working alone, who found death records for hundreds of children and no corresponding burial records (RSAIF MODERN-175).

In a targeting case, the documented figure is not a measurement of the harm. It is a measurement of how hard somebody looked. That is a direct consequence of the mechanism: an institution that selects children also controls the records of what it did to them, and the count therefore tracks investigative effort rather than injury. Sentinel cases do not behave this way. The lead estimate can move up or down with method, but it is generated by external instruments, blood surveys and gasoline consumption records, that the exposing party did not administer. Nobody had to excavate anything to find it.

This has a practical implication that runs against the reflex. When a targeting figure rises, the correct inference is not that the harm grew. It is that the earlier figure was an artifact of the perpetrator's own bookkeeping, and that the current figure probably still is.

D. Institutional care as the harm

Tuam supplies the sharpest version. The Bon Secours Mother and Baby Home operated from 1925 to 1961, and Corless documented death records for 796 babies and infants with no corresponding burial records, roughly one child death every two weeks across the home's operating life (RSAIF MODERN-175). A 2017 test excavation confirmed substantial quantities of juvenile human remains in underground chambers consistent with disused sewage or septic structures, turning a documentary finding into physical evidence. Full excavation began in 2025.

The national picture is worse than the individual scandal. The Commission of Investigation into Mother and Baby Homes, reporting in 2021 on institutions operating between 1922 and 1998, found that about 9,000 children died across the 18 homes it examined, roughly 15 percent, about one in seven, of the children who passed through, and about double the infant-mortality rate of the general population in the same years. The Commission concluded that in the decades before 1960 these homes did not save the lives of children born outside marriage and in fact appear to have reduced their chances of survival (RSAIF MODERN-175).

That finding deserves to be stated plainly. An institution whose stated purpose was the care of children produced worse survival than no institution. The harm was not a failure of the system; it was the system operating.

The criminological literature names the wider structure, and the name is the mechanism. Ireland ran an architecture of containment: Industrial Schools, Magdalene Laundries, Mother and Baby Homes, and psychiatric hospitals, most operated collaboratively by the Church and the state, forming a network that grew until by the 1950s roughly one percent of the country's population was subject to some form of containment (Black and Ring, 2022). The same account records that at least 15 percent of infants died in the homes, and that inside the containment the selection continued: babies of mixed-race heritage were typically placed in Industrial Schools rather than adopted or fostered (Black and Ring, 2022). Captivity was the outer structure and targeting operated within it, which is why this case sits in mechanism two while looking like mechanism three from the outside.

The Magdalene Laundries ran the adjacent machinery for the mothers. Approximately 10,000 women and girls are known to have entered between 1922 and 1996, referred by courts, police, industrial and reformatory schools, health and social services, priests, or their own families, often for being poor, unmarried and pregnant, considered wayward, or

simply without anyone to protect them (RSAIF MODERN-174). The 2013 McAleese report found that about 26.5 percent of known entry routes were referrals made or facilitated by the state. Women worked unpaid in commercial laundries that took real contracts and generated real revenue for the orders holding them. The last laundry closed in October 1996.

E. Medical experimentation

At Willowbrook State School, researchers deliberately exposed newly admitted children with intellectual and developmental disabilities to hepatitis, beginning in 1956 and continuing about fourteen years. More than fifty children aged roughly five to ten were documented in the special unit, with successive groups enrolled across the span (RSAIF MODERN-151). The justification is the part worth preserving: researchers noted that a child entering Willowbrook had close to a one-in-two chance of contracting hepatitis within the first year, and used that figure to argue deliberate exposure changed little. The figure measured the institution's own neglect. The neglect became the argument for experimenting on the children it was neglecting. In 1969 Willowbrook held about 6,200 residents in buildings intended for about 4,000, and parents seeking admission were sometimes told space was available through the research unit when the general institution was full, which placed consent under a pressure close to compulsion.

That defense has a name, and the naming matters because it is what makes the pattern legible across cases. The claim that the researchers merely observed what would have happened anyway was advanced under the heading of a study in nature, and it was interrogated and rejected on the ground that the researchers were not watching a natural process but participating in an institutional one (Rothman, 1982). The move generalizes past medicine. A harm the institution produces is reclassified as a background condition the institution merely found, and the reclassification converts a perpetrator into a witness. Willowbrook entered the professional record through Beecher's survey of unethical clinical research, the document that made the practice visible to the field rather than only to the institutions running it (Beecher, 1966). Children were subsequently designated a uniquely vulnerable class of research subject, and that designation dates to the early 1970s (Laventhal, Tarini, and Lantos, 2012). The protection postdates the studies that established the need for it, which is the ordering this paper keeps finding.

In Guatemala between 1946 and 1948, researchers from the United States Public Health Service, working with the Pan-American Sanitary Bureau and Guatemalan agencies, deliberately exposed prisoners, soldiers, psychiatric patients, sex workers, and children to syphilis, gonorrhea, and chancroid. Roughly 5,500 Guatemalans were involved in the wider program, about 1,300 were deliberately exposed, and at least 83 died, though records do not always establish causation (RSAIF MODERN-152). Methods included arranged sexual contact and, where that did not reliably transmit, direct inoculation onto the genitals, face, or through spinal punctures. The Presidential Commission concluded the studies violated ethics by the standards of their own time and titled its report *Ethically Impossible*. The lead investigator, John Cutler, later took part in Tuskegee. The records sat unexamined until a historian found them in 2010.

At Fernald State School, researchers associated with MIT fed boys breakfast cereal containing radioactive iron and calcium tracers, recruiting them through a Science Club offering outings, extra food, and baseball tickets. Quaker Oats had a commercial interest in demonstrating its cereal's nutritional properties against a competitor's. Parents received letters describing a nutrition study with special diets and blood tests; the letters did not state that radioactive materials would be fed to their children. Accounts range from roughly 50 to about 90 boys, one frequently cited figure being about 74 (RSAIF MODERN-165).

Fernald establishes the decisive point of this section. No death or specific illness has been reliably attributed to the low tracer doses. The federal Advisory Committee on Human Radiation Experiments concluded that while the exposures were unlikely to have caused physical injury, the studies had nonetheless wronged the boys by using them without disclosure or genuine consent (RSAIF MODERN-165). Mechanism two is not a dose story. The wrong is

complete where the toxicology is null. Mechanism one cannot make that claim; a sentinel case with no dose effect is not a case at all. This is the cleanest available demonstration that the two mechanisms are not variants of one thing.

The Committee reached that conclusion after sixteen public meetings, more than two hundred interviews, and review of roughly 400,000 documents, delivering a report of 925 pages (Faden, 1996). The scale is worth recording because it answers the objection that findings of this kind rest on advocacy. It also locates Fernald precisely. The federal apparatus that eventually reviewed it had already produced separate reports on research involving children and on research involving those institutionalized as mentally infirm, as two distinct protected categories (Capron, 2017). Fernald and Willowbrook sit where the categories intersect. The intersection is not an artifact of classification. It is where selection was easiest.

F. The child as the instrument

Family separation at the United States border is the contemporary case that clarifies what mechanism two is, and it belongs here rather than in the captivity family where its detention setting would seem to place it.

Under the 2018 zero tolerance policy, adults were prosecuted for crossing the border and their children were taken into separate custody and treated as unaccompanied minors even though they had arrived with a parent. Separation was used as a deterrent, on the theory that the prospect of losing one's children would discourage others from coming (RSAIF CONTEMPORARY-201). That sentence is the whole mechanism. The harm to the child was not a side effect of enforcing a policy against adults. The harm to the child was the policy, aimed at adults who had not yet arrived. The child is the lever.

This is the purest instrumental use of a child in the corpus, and it is worth sitting with, because instrumentalism is exactly what the sentinel argument risks doing rhetorically. Here it is operational. A child's suffering was assigned a communicative function.

The Department of Homeland Security estimated that Border Patrol separated 3,014 children during the zero tolerance period, but the Inspector General could not confirm the total because of technology and data failures, and later accounting identified additional children separated before the formal policy (RSAIF CONTEMPORARY-201). The American Civil Liberties Union reported that a 2023 settlement would cover an estimated 4,500 to 5,000 children and parents. The tracking failure has its own structure: the adult and the child entered different custody systems run by different agencies, one for criminal defendants and one for children, with no shared record linking them. The government could not always say which child belonged to which parent.

The rising-floor property appears again, exactly as the mechanism predicts. The count is 3,014, then unconfirmable, then 4,500 to 5,000. The institution that separated the children also kept, or failed to keep, the records of having done it.

G. Commercial targeting

The commercial version of selection is less physically brutal than the institutional one and structurally identical, and the infant formula case is its fullest documented run.

Beginning in the 1950s and intensifying through the 1970s, formula manufacturers led by Nestlé identified the Global South as a growth market as Western birth rates declined, and marketed infant formula as modern, scientific, and superior to breastfeeding in regions where clean water was unavailable, poverty made adequate preparation impossible, and the product could consume 30 to 50 percent of a poor family's income (RSAIF CONTEMPORARY-064). The selection machinery was explicit. Company salespeople dressed in nurse uniforms, the milk nurses, visited new mothers in hospitals and homes; free samples were distributed at birth, timed to the window in which lactation is established; and the biology performed the lock-in, because a mother whose milk never came in has no alternative when the free

tin runs out. The corpus record preserves the mortality differential: bottle-fed babies in developing countries were 14 to 25 times more likely to die than breastfed babies, and the World Health Organization has estimated 1.5 million infant deaths preventable annually through exclusive breastfeeding, a global figure to which marketing-driven formula dependence contributed and which is carried here as the estimate it is (RSAIF CONTEMPORARY-064).

The accountability arc is the recourse asymmetry in miniature. The 1974 War on Want pamphlet *The Baby Killer* drew a libel suit from Nestlé against activists who reprinted it, which the company won on the title while the court found the underlying criticism valid. The boycott that began in 1977 became the longest consumer boycott in history. The International Code of Marketing of Breast-milk Substitutes was adopted by the World Health Assembly in 1981 on a vote of 118 to 1, and it was made voluntary; violations have been documented in virtually every country studied across the four decades since, and no corporate liability for the deaths has ever been established (RSAIF CONTEMPORARY-064). The remedy was drawn as a code without an enforcement mechanism, which is the same shape as a remedy drawn narrower than the harm.

The published record fills in both the vote and the interval. The single negative vote was cast by the United States (Brady, 2012). The Code was drafted against a documented rise in mortality, malnutrition, and diarrhoea among very young infants in the developing world, and the obstacle to enforcing it has never been obscure: annual sales of breast-milk substitutes exceed 31 billion dollars, which leaves the industry little incentive to restrict its own advertising (Brady, 2012). The condition has a name that predates the Code by nearly a decade. Commerciogenic malnutrition was coined in 1972 to describe malnutrition produced by commercial promotion rather than by scarcity, and the term survives because the thing it names did (Lutter, Hernandez-Cordero, and Grummer-Strawn, 2022).

Measurement in 2022, forty-one years after adoption, found that 64 percent of more than three thousand pregnant women and mothers across eight countries had encountered promotion of Code-covered products within the previous six months, that roughly one in five mothers of infants under six months had been advised by a health care provider to feed their child something other than breast milk, and that the marketing tactics had grown more sophisticated rather than less (Lutter, Hernandez-Cordero, and Grummer-Strawn, 2022). The same study found violations tracking the strength of national legislation, with the countries holding the strongest Code laws showing the fewest violations. The voluntary instrument did not fail for want of evidence. It failed because it was voluntary, and where a state made it binding it worked better. That is the enforcement gap stated as a measurement rather than as a complaint.

Juul sits in the same family with different chemistry: a replacement market identified in adolescents, and flavor, device, and delivery tuned to it. Youth e-cigarette use rose from 1.5 percent in 2011 to 20.8 percent in 2018, and by the end of that year Juul held 76 percent of the retail market; the product was marketed as a cessation aid and is rarely used for that purpose by young people (Fadus, Smith, and Squeglia, 2019). By 2019, 27.5 percent of United States high school students reported vaping within the previous thirty days, more than two and a half times the 2017 figure (Carroll and Cheng, 2020).

The design record decides the assignment, and the reasoning is worth showing because the case briefly resembles a captivity case. Pod devices represent an evolution in design that increases the efficiency of nicotine delivery, paired with flavor appeal, sleek form, and ease of concealment (Lee, Rees, and Yossefy, 2020). Concealability is a design choice aimed at a population that must hide its use inside a building it is not free to leave. That fact invites an assignment to mechanism three, and the invitation should be declined. Juul holds no one anywhere. The child carries the product in. What concealability demonstrates is targeting sophisticated enough to design around a constraint its target already lives under, which makes the school a background condition the selection exploits rather than the structure delivering the harm.

The distinction is not pedantry, and the decision procedure earns its keep here. If the mechanism were captivity, the remedy would run to the school. Because the mechanism is targeting, the remedy runs to the design and to the

marketing that carried it, a strategy built on social media and affiliate channels, which is to say on precisely the surfaces where age is least verifiable (Lee, Rees, and Yossefy, 2020). A platform that identifies minors as a growth segment and calibrates reward architecture to their developmental stage runs the identical selection with a feed instead of a formula tin.

H. Reproductive coercion

North Carolina used a standing Eugenics Board to authorize sterilization of people labeled feeble-minded, mentally ill, epileptic, delinquent, poor, or socially undesirable. More than 7,600 people were sterilized between 1929 and 1974, third in the nation behind California and Virginia (RSAIF MODERN-147). Unlike programs confined to hospital inmates, North Carolina let county welfare departments petition for sterilization of people living in the community, extending reach into ordinary households. The board reviewed petitions on paper and approved most, so a small body of officials ordered irreversible surgery on people they never met.

The program sterilized people as young as ten, and the great majority of operations were carried out after 1945, precisely as other states wound their programs down and eugenics lost scientific credibility elsewhere (RSAIF MODERN-147). Its racial focus shifted over time, with a rising and eventually majority share of those sterilized being Black in a state where Black residents were a minority. Welfare workers sometimes presented sterilization as a condition of continued public assistance. Nial Ruth Cox was sterilized as a teenager after her family was warned it could lose assistance. For many survivors, the sterilization was discovered only in adulthood, when they tried to have children or requested their own medical records.

The scholarly record corroborates the corpus on the two facts that matter most here. The Board sterilized people as young as ten, and the selection ran through social workers rather than physicians, which is what produced the disproportionate share of Black women among those sterilized (Chatters, Taylor, and Schulz, 2022). A caseworker's discretion, exercised on a form, is the administrative instrument this mechanism keeps producing. The same literature records that those sterilized, and the harms they carried, have been rendered invisible (Chatters, Taylor, and Schulz, 2022). That is not a metaphor. It is a description of what follows when the population able to testify is the population least equipped to be heard.

The comparative record supplies a fact that cuts against reading North Carolina as an outlier in cruelty. Three states have established compensation programs: North Carolina, Virginia, and California. Utah, whose sterilizations ran across six decades and three institutions, has offered no acknowledgment and no restitution (Tabery, Novak, and Sarafraz, 2023). North Carolina is not the worst case of redress in the record. It is among the best available, which is what makes the arithmetic in the section on recourse worth reading twice.

I. Childism as the correct diagnosis

Young-Bruehl's account of childism as a prejudice with its own structure, rather than as an incidental byproduct of adult convenience, is the right instrument for all of these cases (Young-Bruehl, 2012). What they share is not a physiological property of children but a belief about children: that they are raw material, that their incompleteness is a resource, and that the absence of a formed self is a licence rather than a claim to protection. The Mamluk case makes it structural, the devshirme administrative, the three assimilation systems explicit, Tuam actuarial, Willowbrook and Guatemala medical, North Carolina bureaucratic, and the formula campaign commercial. None of it is epidemiology.

The definition is worth stating in full because it reads as a description of mechanism two rather than as a theory about it. Childism is the prejudice that rationalizes or justifies acts and omissions harming children, failing to meet their needs and protect their rights, resting on the belief that children are property and can, or even should, be controlled,

enslaved, or removed to serve adult needs (Savage, 2023, citing Young-Bruehl). The term is older than the 2012 book: two American physicians used it in 1975 to name the automatic presumption of superiority of any adult over any child, and the consequent precedence of adult needs, desires, hopes, and fears over the child's (Savage, 2023). Controlled, enslaved, or removed to serve adult needs is not an analogy for the Mamluk, devshirme, and residential-school cases. It is an inventory of them.

Recent philosophy strengthens the ground under this and supplies the structural argument the corpus keeps demonstrating. A social category becomes morally relevant when three conditions hold: individuals are externally ascribed to it, its properties are reified through the social construction of its definition, and the category itself predetermines the normative treatment owed to those ascribed. Childhood satisfies all three (Brando, 2022). The consequence is the part that matters here. Adults always decide how children should be treated, and the structural asymmetry survives good treatment: a child who is treated kindly, given a voice at home and at school and in public affairs, remains dependent on adult will to decide whether those privileges are granted at all (Brando, 2022). Benevolence does not dissolve the asymmetry, it exercises it. That is mechanism three stated as moral philosophy rather than as a building, and it explains why captivity in this paper is a structural relation rather than a physical one.

The developmental literature reaches the same place through oppression theory, and names the mechanism outright. Applying Young's five conditions, anti-youth ageism is analyzed as marginalization, powerlessness, cultural imperialism, exploitation, and violence, with exploitation defined as being used to benefit those in power (Wray-Lake, Rottenberg, and Kennedy, 2025). Being used to benefit those in power is mechanism two in five words, arrived at independently and from another discipline. The same account records that oppression is enacted through tyrannical power and also manifests through unquestioned norms and practices carried out by well-meaning people (Wray-Lake, Rottenberg, and Kennedy, 2025), which is why the corpus contains so few villains and so many committees. The terminology is unsettled across these literatures, which use adultism, juvenile ageism, childism, adultcentrism, and youth oppression for overlapping phenomena, and the disagreement is itself informative: a prejudice with five competing names and no settled one is a prejudice that has not been consolidated as an object of study. The disagreement is not merely lexical. Childism proper is rooted in ethics and moral philosophy and works largely through rights discourse, mirroring feminism and antiracism in seeking child-inclusive social processes and structures; but the field acknowledges an internal objection to that framing, that a politics of inclusion fails to problematize what the formerly excluded category is being included into (Burman, 2022). This paper inherits that objection rather than answering it. Every mechanism here describes children deposited somewhere by a structure adults built, and a remedy consisting of admission to the structure is not obviously a remedy.

One further strand bears directly on the section that closes this paper. The philosophy of childhood has long treated childhood as a form of temporary disability, defined by its inabilities and its need of correction, and has concentrated its analysis of children's status on children's lives as instrumental to the ends of society or of their parents (Ben-Porath, 2003). Against that, Ben-Porath argues that adults' obligations, rather than children's rights, are the appropriate basis for adult society's relations with children. This paper does not need to resolve that debate, but the recourse asymmetry gives it a sharp edge. A right requires a claimant who can assert it. An obligation does not. Where the barrier is general and falls hardest on those least able to press a claim, a framework built on adult obligation routes around the exact structure that defeats a framework built on the child's right.

Wall's work does something different and is needed for a different purpose (Wall, 2010). Where Young-Bruehl diagnoses the prejudice, Wall supplies the affirmative ground on which a child's claim rests without routing through the child's future adulthood. That distinction is load-bearing for mechanism one, not two. The reason the sentinel argument is dangerous even when true is that it makes a child's injury matter for what it predicts about adults. It is instrumental. It says: attend to the children because the children are telling you what will happen to you. Wall is what prevents that

from becoming the whole case. The child's injury is a harm to a person now, and its sentinel value is an additional fact about the world rather than the source of its moral weight. A paper that runs the sentinel argument without that anchor has argued that children matter as instruments, and it will deserve what it gets.

V. MECHANISM THREE: CAPTIVE ABSORPTION

The third family is most often mistaken for the first, because the epidemiology looks similar. Children appear disproportionately in exposure counts. The explanation is location.

A. Aberfan

Merthyr Vale Colliery had tipped waste onto the slopes above the village of Aberfan for roughly half a century. Seven spoil tips stood directly above the village. Tip 7, begun in 1958 and the only one in use by 1966, stood about 111 feet high and held roughly 297,000 cubic yards of spoil, including about 30,000 cubic yards of tailings, the fine waste that behaves like quicksand when wet. In contravention of the National Coal Board's own procedures, the tip was partly built on ground from which springs emerged. Those springs had been marked on Ordnance Survey and Geological Society maps since 1874 (RSAIF MODERN-207).

On the morning of 21 October 1966, after weeks of rain, the saturated tip slid. At about 9:15 it came down as a fast-moving slurry, destroyed about twenty houses and a farm, and engulfed Pantglas Junior School, where roughly 240 children had just come out of assembly. The valley was under low cloud, so nobody saw it. Everybody heard it, and it was moving too fast to outrun. It was the last day before half-term. It killed 144 people: 116 children, most between seven and ten, and 28 adults including five teachers. Roughly half the village's children died.

Aberfan is the cleanest available demonstration that a high child death count can carry no information whatsoever about susceptibility. Nothing about a seven-year-old made that child more vulnerable to a landslide than an adult. The children died because of a timetable, in a building attendance at which was compulsory, beneath a hazard their parents had reported and could not remove.

The warning record is what makes it a case rather than an accident. The tip complex had already slid in 1944 and 1963, with physical evidence visible to the naked eye, though the Board spent days at the inquiry denying the 1963 slide occurred. The Pantglas area had flooded severely at least 11 times between 1952 and 1965, residents reporting black water that left a greasy residue. The borough council corresponded with the Board between July 1963 and March 1964 specifically about the danger from coal slurry tipped at the rear of the Pantglas schools. A witness told the tribunal that in the four months before the collapse the toe of Tip 7 moved downward by roughly 20 to 30 feet (RSAIF MODERN-207).

The tribunal sat for 76 days, the longest inquiry of its type in British history to that point, taking evidence from 136 witnesses and hearing 2.5 million words. It concluded that the disaster could and should have been prevented, that blame rested upon the National Coal Board, and that there was a total absence of tipping policy, which was the basic cause. It characterized the failure as one not of wickedness but of ignorance, ineptitude, and a failure in communications (RSAIF MODERN-207).

B. Schools

Asbestos was built into American schools through much of the twentieth century as pipe and boiler insulation, sprayed fireproofing, and ceiling and floor tiles. A 1984 EPA survey estimated roughly 34,800 schools contained friable asbestos-containing material, potentially exposing about 15 million students and about 1.4 million school employees

(RSAIF MODERN-193). Children were exposed because attendance is compulsory.

This case requires an honest complication that the corpus record supplies and the sentinel framing would suppress. The disease that surfaces is long-latency, taking twenty to fifty years, and it surfaces especially among the custodial and maintenance workers whose jobs bring them into direct contact with the material (RSAIF MODERN-193). The children are the larger exposed population; the adults are where the disease record is. That is the reverse of the sentinel structure. It is a mechanism-three case precisely because it is not a mechanism-one case, and reading it as sentinel would invert what the evidence says.

The record contains a second complication worth keeping. Early responses sometimes stripped asbestos out in ways that released more fibers than leaving it intact would have, which is why current guidance emphasizes managing sound material in place (RSAIF MODERN-193). A remedy applied to the wrong mechanism can increase the harm.

The risk-assessment literature sharpens both complications, and it cuts against this paper in the process, which is why it is reported rather than omitted. Applying the risk-assessment criteria used for the federal Superfund program, one assessment found that removals had little effect on risk to school populations, that risks to teachers and students in asbestos-containing school buildings were approximately equivalent to the general public's exposure to ambient asbestos and below the levels regulators typically treat as concerning, and that abatement itself raised risk to workers and to nearby individuals, with careless routine building maintenance generating the greatest worker risk of all (Chrostowski, Foster, and Anderson, 1991). Judged by those criteria, taking no action would have been preferred to removal in a majority of cases.

Two things follow and only one of them helps. The abatement finding supports what this paper claims about remedies aimed at the wrong mechanism. The exposure finding undercuts the premise that school asbestos delivered substantial measured harm to the children held in the buildings. The case therefore earns its place in mechanism three through its structure rather than through a demonstrated pediatric injury: compulsory attendance placed roughly 15 million children in buildings containing the material, and every decision about what to do with that material was made by adults who could leave. That is a weaker claim than the sentinel framing would make and it is the claim the evidence supports. It also rests here on a single assessment, now more than three decades old, in a literature where risk estimates have been contested; it is carried as a complication, not as a finding.

C. Housing

Lead paint is a hazard fixed in place. It remains on the walls and woodwork of millions of older buildings, layered under newer coats, released slowly across the life of the building through deterioration, friction on windows and doors, renovation, and deferred maintenance (RSAIF CONTEMPORARY-057). Roughly 29 million housing units carry lead-based paint hazards, about 2.6 million occupied by young children. Families who rent, who cannot compel repairs, and who cannot afford to move are least able to escape a hazard fixed in the walls around them.

Public housing sharpens this. Federal enforcement alleged that the New York City Housing Authority, the country's largest public-housing landlord, made false statements to HUD and the public about lead-paint compliance and intentionally deceived inspectors by concealing conditions rather than fixing them. The settlement required the City to provide 1.2 billion dollars in additional capital funding over five years and 200 million dollars annually thereafter until the problems were fixed; the authority houses about 400,000 residents; and a 2019 HUD agreement installed a federal monitor, reflecting a judgment that the agency could not be trusted to fix the conditions itself (RSAIF CONTEMPORARY-203).

The children in those apartments were not selected and were not more susceptible than the adults to a lie. They were exposed because a landlord that is also a government agency misrepresented its compliance, and residents cannot

independently verify the safety of their own homes.

D. Water

In April 2014, under state-appointed emergency management, Flint switched its drinking water to the Flint River to save money without adding corrosion control. The corrosive water stripped protective coatings from lead pipes. The emergency-management structure that prioritized cost savings also concentrated authority in appointees not accountable to local voters (RSAIF CONTEMPORARY-154). Residents reported discolored, foul-smelling water almost immediately and were treated as anecdotal. The 2021 settlement totaled 626.25 million dollars covering more than 90,000 residents, with the majority directed to children. A Legionnaires' outbreak linked to the switch caused at least 12 deaths. Criminal charges against officials were largely unwound on procedural grounds.

Flint contains a finding that belongs to this paper's central distinction. Officials disputed residents' complaints for months. What they could not dispute was a local pediatrician's data showing rising lead levels in children's blood (RSAIF CONTEMPORARY-154). The children's bodies were the detection instrument, not for a harm that would later reach adults, but for a present harm the adults were being told was not happening. That is a third thing, and the vocabulary does not currently have a name for it: not sentinel-for-the-future, but sentinel-for-the-truth. The children's blood was the only evidence the denial could not absorb. The published analysis found the proportion of children with elevated blood lead rising after the source change, with the increase concentrated in socioeconomically disadvantaged neighborhoods (Hanna-Attisha, LaChance, and Sadler, 2016).

Flint also predates itself, and that is the captivity point in its sharpest form. Audit and compliance records show children in Flint Community Schools drinking from lead plumbing inside the school buildings years before the switch made the city a national story, and the water was already corrosive before 2014 (Davis, 2016). The crisis has a start date because that is when the adults noticed. The captivity has no start date. It is the building.

Washington ran the same failure a decade earlier and larger. In an effort to reduce disinfection byproducts, the water system switched to chloramine without adequately accounting for its effect on corrosion control, and lead that had been stable inside the pipes for years began to dissolve. A study of 2,342 children under six found that 361 of them, 15.4 percent, lived in homes with lead service lines, and that 19 of those children, 5.3 percent, had elevated blood lead, tying the exposure to the pipes rather than to other sources (RSAIF CONTEMPORARY-184). The CDC concluded the association was strongest from 2001 through June 2004, the window when chloramine ran without adequate corrosion control. As in Flint, early messaging understated the risk.

The Washington record also supplies the second instance of a pattern this paper has already met once. Partial replacement of a lead service line, swapping out only part of it, can jar loose additional lead and briefly raise exposure (RSAIF CONTEMPORARY-184). The engineering literature is less forgiving than the word briefly suggests. Sampling at more than a hundred residences found that partial replacement more than doubled lead release from premises plumbing in the short term and did not reduce it in the long term; six months after the work, 27 percent of first-draw samples exceeded the federal action level, against 13 percent before it (Trueman, Camara, and Gagnon, 2016). Full replacement, by contrast, cut lead levels by more than half within three days. The mechanism is understood rather than suspected: joining lead pipe to copper creates a galvanic cell, and bench measurement puts the resulting lead release at between 1.1 and 16 times the level from an intact lead pipe alone (Triantafyllidou and Edwards, 2011).

The asbestos record showed early removal sometimes releasing more fibers than leaving material intact (RSAIF MODERN-193). Twice now, a remedy applied without regard to the mechanism has increased the harm it was meant to end. That is not a coincidence of two domains. It is what happens when the response is aimed at the visible hazard rather than at the structure producing the exposure. In the water case the published finding is stronger than the corpus

record states. The partial remedy did not trade a short-term rise for a long-term gain. It produced the rise and then delivered no gain at all.

Beneath Flint sits the national condition. The EPA estimates about 4 million lead service lines still serve properties in United States communities, each a path for lead to reach drinking water. The October 2024 Lead and Copper Rule Improvements require inventory and replacement within ten years (RSAIF CONTEMPORARY-192). Split ownership between utility and property owner means a single pipe can have a public and a private half with no party responsible for the whole, which has determined how quickly replacement proceeds. Every line still in the ground is a standing risk, and because corrosion control can fail, treatment is not a permanent substitute for removal.

E. Newark, and the proof that captivity is a choice

Newark is the most important case in this section, because it is the counterfactual, and a mechanism with no counterfactual is a description rather than a claim.

Newark faced the same failure: corrosion control failed as water chemistry shifted, lead rose above the federal action threshold, officials were slow to acknowledge it, residents were told to rely on filters whose performance was then called into question, and the city moved to bottled water. Litigation by the Natural Resources Defense Council pressed the city to treat the water and disclose the risk rather than manage it out of view (RSAIF CONTEMPORARY-186).

Then Newark did the thing every other city says cannot be done. It announced a program to replace approximately 18,000 lead service lines over 24 to 30 months, and as the inventory grew it ultimately replaced all of its roughly 23,000 known lead service lines in less than three years (RSAIF CONTEMPORARY-186). The mechanism by which it did so is the finding: Newark replaced the privately owned portion of service lines at public expense, removing the cost barrier that had stalled replacement everywhere else.

Recall what the national record says that barrier is. Split ownership between utility and property owner means a single pipe can have a public and a private half with no party responsible for the whole, and resolving who pays for the private segment has determined how quickly replacement proceeds in many cities (RSAIF CONTEMPORARY-192). Newark resolved it in one budget decision and cleared 23,000 pipes in under three years.

That is the whole argument of mechanism three, delivered as a natural experiment. The children in Flint, in Washington, and in the 4 million properties still on lead lines are not captive because captivity is a physical fact about childhood. They are captive because of a property-law allocation that one city chose to override and the others did not. Captivity is a policy setting. It has a dial, Newark turned it, and the dial is still where it was everywhere else.

F. Camp Lejeune, and the scale a captive population can reach

For roughly three decades the drinking-water systems serving Marine Corps Base Camp Lejeune were contaminated with trichloroethylene, tetrachloroethylene, vinyl chloride, and benzene, from on-base sources and an off-base dry cleaner. It reached barracks, base housing, schools, and workplaces. Federal estimates indicate as many as one million service members, family members, and civilian workers may have been exposed between 1953 and 1987 (RSAIF MODERN-155).

The captivity here is unusually legible. A child at Camp Lejeune lived there because a parent was stationed there. There was no other water. Among the documented harms are miscarriages, childhood cancers, and birth defects, and the 2012 statute creating health-care eligibility was named for a Marine's daughter who died of leukemia.

The concentrations are known only because they were reconstructed. Measured data were limited or unavailable across much of the contamination period, so the federal agency modeled monthly mean levels from groundwater flow,

contaminant transport, and water-distribution simulation in order to have anything to study at all (Maslia, Aral, and Ruckart, 2016). What the reconstruction produced is a scale worth stating in units. At the Hadnot Point treatment plant, reconstructed trichloroethylene reached a maximum monthly average of 783 micrograms per liter against a federal maximum contaminant level of 5, with one measured value at 1,400, across the period from 1953 to 1984. At Tarawa Terrace, tetrachloroethylene reached a maximum monthly average of 183 against the same limit of 5, and exceeded that limit continuously from 1957 to 1987 (Maslia, Aral, and Ruckart, 2016). A captive population drank water at more than a hundred times the federal limit for three decades, and the record of what they drank had to be simulated after the fact.

The epidemiology is real and narrower than advocacy summaries suggest, which is the distinction this paper keeps insisting on. A case-control study of children born on base between 1968 and 1985 found an odds ratio of 4.1 for neural tube defects with any first-trimester benzene exposure, with a confidence interval running from 1.4 to 12.0, and a monotonic exposure-response relationship for trichloroethylene (Ruckart, Bove, and Maslia, 2013). Associations with childhood hematopoietic cancers were weaker, with confidence intervals spanning one, and no association appeared for oral clefts. The authors describe their intervals as wide and their results as suggesting rather than establishing the associations, and the paper carries the finding at that strength.

One methodological detail belongs to the argument rather than to the methods. The interviews on which the study rests were conducted twenty to thirty-seven years after the births, which the authors note likely produced recall error and missing data (Ruckart, Bove, and Maslia, 2013). The delay is not a limitation of the research. The delay is the thing under study, and it degraded the evidence that might have measured it.

G. Captivity, not detection

Mechanism three is where the sentinel framing does the most damage. If a child in Flint is described as a canary, the implicit claim is that the child's lead level is informative about what will happen to adults in Flint. That is not what the case shows. It shows a class of people who could not choose their water being poisoned by people who chose it for them. Reframing that as a warning to adults centers the wrong population and is also false, since the adults drank the same water and their harm was direct rather than predicted.

VI. MIXED CASES AND THE DISCIPLINE OF ASSIGNMENT

Real cases mix mechanisms. Assignment is a judgment about which structure predominates, not a partition, and a framework that cannot say how it assigns is not a framework. This section states the decision procedure and then tests it on cases that resist it.

A. The decision procedure

Four questions separate the mechanisms, and they are asked in order.

First, were children chosen? If the harm fell on children because someone selected children, the case is mechanism two, and no feature of dose or susceptibility changes that. Fernald settles this at the null end: a case can be mechanism two with no physical injury at all.

Second, if children were not chosen, does a physiological or behavioral property of childhood make the harm register in children before adults at the same ambient exposure? If so, and if the adult burden later materializes, the case is mechanism one. The property must be specified, lead absorption, thyroid iodine uptake, respiratory vulnerability, milk consumption, not assumed from the fact that children were hurt.

Third, if neither selection nor differential susceptibility explains the pattern, were children present in disproportionate number because they could not leave or could not choose? Then the case is mechanism three, and the injury count is a fact about location rather than about biology.

Fourth, if none of the three applies, it may not be a children-as-canaries case at all, and the honest answer is to say so rather than to force it.

B. Cases that carry two mechanisms

Lead is the clearest genuine mix. It is sentinel on detection order and captivity on distribution. The sentinel reading explains why the signal was available in the 1970s; the captivity reading explains why the burden falls 2.8 times harder on Black children (Pickering et al., 2026) and why 2.6 million young children live in units with lead hazards (RSAIF CONTEMPORARY-057). Both are true, and the assignment records which reading answers which question rather than picking one.

Residential schools mix targeting and captivity, and the corpus preserves the exact moment the second was bolted onto the first: the 1920 amendment to the Indian Act making attendance compulsory (RSAIF COLONIAL-118). Before 1920, selection. After 1920, selection plus legal captivity. The mechanism changed by statute on a date.

Agent Orange mixes sentinel detection with intergenerational transmission, which is why it sits in mechanism one with a caveat rather than cleanly. Swill milk mixes sentinel susceptibility with economic captivity, and the corpus is explicit that infants were most exposed because cheap loose milk was what the poorest households could buy (RSAIF MODERN-204), which is mechanism three wearing mechanism one's clothes.

C. The case that is none of them

Boeing is in this section precisely because it is not a children-as-canaries case, and running the procedure on a case that fails it is how the procedure earns trust.

The 737 MAX crashes killed 346 people across Lion Air Flight 610 and Ethiopian Airlines Flight 302, traced to an automated system, MCAS, that could command repeated nose-down movement from a single angle-of-attack sensor, certified with inadequate pilot disclosure (RSAIF MANU-001). Run the four questions. Were children chosen? No. Did a property of childhood make the harm appear in children first? No; the passengers were a cross-section, and a child on those aircraft died at the same rate as an adult in the next seat. Were children present in disproportionate number because they could not leave? No; the aircraft is not a population children occupy disproportionately. The case fails all three, and the correct output is that it is a corporate-harm case with no sentinel, targeting, or captivity structure at all.

That matters because the pull to annex a vivid case is strong, and Boeing is vivid. A framework that claimed it would be a framework that claims everything, which is the disease this paper was written to treat. Boeing belongs in the corpus. It does not belong in this paper's argument, and saying so is the discipline working.

D. Two standing temptations

The first is to read any case with a high child exposure count as a sentinel case, which the asbestos-in-schools record refutes directly: 15 million children exposed, and the disease surfacing in the maintenance workers (RSAIF MODERN-193).

The second is to reach for child labor as evidence that children register economic harm first. They do not. The 1900 census recorded approximately 2 million children working; in Southern textiles that year, 25,000 of nearly 100,000 workers were children under sixteen; by 1904 about 50,000 children worked in Southern mills, roughly 20,000 under

twelve (RSAIF MODERN-199). Those children are not an early signal of a harm that later reached adult millworkers. They are the labor, and the harm is the same harm in a smaller body. Reading it as sentinel would be the exact category error this paper is written against.

The child-labor record supplies one more thing the paper needs. Congress passed the Keating-Owen Act in 1916; the Supreme Court struck it down in *Hammer v. Dagenhart* in 1918; a later tax on child-labor products was also invalidated; durable federal limits arrived only with the Fair Labor Standards Act of 1938, and the Court reversed its earlier reasoning in *United States v. Darby* in 1941 (RSAIF MODERN-199). Twenty-five years between the first federal attempt and the durable one, and the obstacle was doctrinal, not evidentiary. Everyone could see the children. The photographs existed. That interval is the historical rhyme with Section VIII, where the obstacle is again doctrine rather than proof.

VII. THE CONTEMPORARY DIGITAL CASES

The mechanisms carry forward, and the assignment works. It is worth doing carefully, because the child-safety field's habitual move is to call all of it one thing, which is the same error that produced the canary confusion in the first place.

A. Captivity: the platform performs the pairing

Omegle connected strangers in one-to-one video chatrooms under a tagline inviting users to talk to strangers. It required no personal identifying information, designating participants only as You, Stranger 1, and Stranger 2, so users chatted with complete anonymity. It did not restrict children from the service (RSAIF TECH-038).

The captivity here is not spatial, and that is what makes it instructive. A child could close the browser. What the child could not do was choose the counterparty, because the platform performed the matching. The exposure was not to a place the child could not leave but to a selection the child did not make and could not see. That is captive absorption with the geometry rotated: the constraint is on the choice rather than on the exit.

A district court permitted a product design claim to proceed, reasoning that the duty to design a reasonably safe product is fully independent of a defendant's role in monitoring or publishing third-party content, and analogizing to *Lemmon v. Snap*, where liability attached to a negligently designed feature rather than to any user's post (*A.M. v. Omegle*, 2022). The amended pleading alleged that by advertising to both children and adults while randomly pairing users, the operator unreasonably created a foreseeable risk, and that by matching minors with adults while knowing predators were likely to use the platform, it caused the abuse. Omegle ceased operations in November 2023, under litigation and reputational pressure rather than by regulatory order (RSAIF TECH-038). No court ever ruled on whether the pairing architecture was defective.

B. Targeting: reward architecture aimed at a developmental stage

The video game addiction wave alleges that publishers borrowed mechanics from gambling and behavioral psychology, variable-reward loot systems, battle passes, streak maintenance, and reward schedules calibrated to maximize session length, and that their own telemetry identified which mechanics produced compulsive use in young players (RSAIF TECH-036). If that allegation is true, the structure is selection in Young-Bruehl's sense with a modern delivery system: a developmental stage identified as a resource and tuned toward.

The Instagram record supplies the version with internal documentation attached. Reported internal research found that among teen girls who already felt bad about their bodies, 32 percent said Instagram made them feel worse, and that 13 percent of British and 6 percent of American teen users who reported suicidal thoughts traced those thoughts to

Instagram (RSAIF TECH-010). More than forty states have joined actions alleging that the company misled the public about youth harms and used addictive designs affecting children.

The companion chatbot cases are the newest instance and the one closest to the historical targeting pattern, because the product simulates a relationship. The documented deaths in the litigation record include a 13-year-old girl in 2023 and a 14-year-old boy in 2024, alleged by their families to be connected to companion chatbot interactions, and those remain allegations rather than findings of causation (RSAIF TECH-030). A federal court allowed key claims to proceed past dismissal in May 2025 (*Garcia v. Character Technologies*, 2025).

C. Neither: the reporting failure

Platform retention of child sexual abuse material after notice is not a sentinel case, not a targeting case, and not a captivity case. It is a straightforward operational failure, and its place in this paper is what happened when the children tried to use the remedy.

Two minors alleged they had been groomed and coerced as boys into producing explicit videos, that the videos were later posted by third parties, and that when they reported the material as child sexual abuse material, the platform declined to remove it. One submitted identification to prove he was a minor, and the support process responded that no violation of policy had been found and no action would be taken. The videos remained available for roughly a further week, coming down only after a federal agent intervened; an earlier citizen complaint identifying the posting account's links as child pornography had produced no action (RSAIF TECH-037; *Doe v. Twitter*, 2021).

A child identified himself, proved his age, reported his own abuse through the mechanism built for exactly that, and was told by an automated review that it did not violate policy. Whatever else the child-safety apparatus is, it is not a system that failed to be told.

D. What the assignment is for

The distinction is not academic bookkeeping. It determines what a plaintiff must prove and what a legislature would have to change.

The captivity reading of Omegle points at the pairing architecture, which is why the design-defect theory reached past Section 230 in Oregon. The targeting reading of the game and platform cases points at internal knowledge, which is why the Instagram internal research and the MDL's knowledge-based duty theory are where that litigation lives. The reporting-failure reading of the retention case points at nothing structural at all, which is why the only claim that survived was the narrow beneficiary theory, and why it required proving the platform knew.

Three cases, three mechanisms, three different doctrinal seams, and one shared outcome, which is the subject of the next section.

VIII. THE RECOURSE ASYMMETRY

The historical record and the contemporary record converge on a single finding, and it is the most important thing in this paper.

North Carolina states it plainly. More than 7,600 people were sterilized. In 2013 the legislature reserved 10 million dollars for compensation. More than 780 people applied. Fewer than 220 were certified and paid, because eligibility was limited to operations authorized by the Eugenics Board itself, and many people sterilized under related county or institutional authority fell outside the definition and received nothing (RSAIF MODERN-147). The gap between

7,600 and 220 is not a funding failure. It is a definitional one. The remedy was drawn narrower than the harm, and the drawing was done by the same state.

The redress literature has a name for this, and it is not incompetence. Trade-offs between accessibility, cost, and evidentiary rigor are hard-wired into the instruments used to assess historic redress claims, which is why programs in Australia, Canada, Ireland, and Sweden draw criticism of the same shape regardless of which government runs them (Winter, 2020). The exclusion is a property of the tool rather than a failure to wield it well. That is a stronger claim than bad faith, and a harder one to fix.

The pattern repeats, and each instance fails on a different technicality, which is the point. Nevada downwinders waited until 1990 for an administrative remedy, and families exposed by the same weather systems were for decades told they lived outside the compensable line (RSAIF MODERN-156). Fernald survivors received about 1.85 million dollars in 1998, divided among the class, in late adulthood, long after many responsible parties had died (RSAIF MODERN-165). The Guatemalan claim against the United States was dismissed on sovereign immunity (RSAIF MODERN-152). The religious orders that ran the Magdalene Laundries and profited from the unpaid labor contributed little or nothing to the redress scheme, and eligibility rules excluded some who had suffered (RSAIF MODERN-174). Flint's criminal prosecutions were largely unwound on procedural grounds (RSAIF CONTEMPORARY-154).

Camp Lejeune is the asymmetry at its largest measured scale. Congress removed the legal barriers in 2022 and created an administrative route with fixed awards from 100,000 to 550,000 dollars. More than 400,000 claims were filed. By mid-2026, settlement offers exceeded roughly 795 million dollars and payments exceeded roughly 571 million, and fewer than one percent of claimants had been paid (RSAIF MODERN-155). The statute worked. The barrier moved from doctrine to throughput, and the outcome did not change, because a large share of those affected are elderly or already dead. A remedy that arrives after the claimant does not is not a remedy; it is an accounting entry.

The theory of remedy explains why the arithmetic keeps resolving this way. A monetary payment is said to restore a victim to the position they held before the wrong, and in practice it falls short, because the sum is calculated against harm and suffering while the wrong is a separate thing: undoing the harm does not undo the wrong, and one cannot repay a wrong the way one repays a debt (Daly, 2025). Money may be necessary and never sufficient. That is a limit on what any remedy could achieve, and it is not the argument this paper is making. This paper's claim is smaller and harder to answer. These remedies did not fall short of an unreachable ideal. They fell short of their own stated terms.

Aberfan is the limit case, because there the causation was never in doubt. A tribunal found that a public body caused 144 deaths through a documented total absence of policy. No employee was prosecuted, sacked, or demoted; one notably unsatisfactory witness had been promoted by the time Parliament debated the report; the Board was never prosecuted; and its chairman offered a resignation whose refusal letter he had already agreed the wording of. The Board initially offered 50 pounds per victim, raised to 500 after public outrage, then took 150,000 pounds from the public disaster fund to help pay for removing the remaining tips, so that money donated for the bereaved financed the safety work the Board had failed to perform. It was repaid in 1997, thirty-one years later, without interest (RSAIF MODERN-207).

Aberfan disposes of the objection that these outcomes reflect evidentiary difficulty. Everyone knew what happened. The tribunal said so in 1967. The remedy still did not arrive.

The transitional-justice literature states the harder version. Examining Ireland's two decades of inquiries and redress schemes, Gallen finds an unwillingness on the part of both Church and State authorities to accept the fundamental transformation that addressing systemic historical abuse would require, and concludes that across the mechanisms the state designed, power remained outside the hands of victim-survivors, with the state's own practices marginalizing them and risking new forms of harm and distress (Gallen, 2020). He reaches for Balandier's proposition that power's most effective move is to permit itself to be contested ritually, and so consolidate. This paper does not need to endorse that

claim about motive in order to record that the pattern it predicts is the pattern the corpus contains.

The contemporary cases run the same structure through different doctrine. In the video game litigation, one court found the Section 230 analysis of prior platform decisions persuasive and barred the claims (*Angelilli v. Activision Blizzard*, 2025). Another reasoned that liability would broadly chill protected expression, noting, on authority developed to protect minors, that minors are entitled to a significant measure of First Amendment protection (*Courtright v. Epic Games*, 2025). A doctrine built to protect children's access to expression was applied to defeat children's claims about the machinery built around it.

In the Omegle appeal, the Eleventh Circuit affirmed dismissal, holding the statutory exception permitting civil sex trafficking claims against platforms requires actual knowledge, not constructive knowledge, of the underlying incident (*M.H. v. Omegle*, 2024). It held separately that the child pornography claim failed because the complaint did not allege the operator ever possessed or accessed the recordings. It described the allegation that the operator paired a child with a stranger knowing she was a minor at risk of becoming a victim of child pornography as disturbing, and held it insufficient. The design argument was never reached, having been raised first in a reply brief. The architecture that produced the harm, anonymity, is the architecture that makes actual knowledge unprovable. The design supplies its own defense.

The redress scholarship arrives at the same structure from the other direction. Schemes built to widen access have been found to widen it at the cost of limited coverage, restrictive eligibility criteria carried over from the litigation they were meant to replace, and no independent route of appeal, with the result that the scheme serves the institutional and financial interests of the body that created it (Farrell and Devaney, 2007). And where survivors' own disadvantages are what obstruct their claims, the impulse to design around those disadvantages produces what one of the field's founders called the great temptation of paternalism: the same facts that justify the program are then used to justify restricting the people it exists to serve (Winter and Wilkinson, 2024).

The reflex is to say the law fails children. That framing is available, it is emotionally correct, and it is analytically wrong in a way that determines what happens next. The same immunity operates identically on adults defrauded by advertisements they were shown (*Calise v. Meta*, 2024). The barrier is general. It was not built to exclude children and does not select for them. That is not an idiosyncrasy of this paper's reading. It is a recognized structural feature of remedial systems built on interpersonal justice. A system requiring a victim to identify a responsible party and successfully pursue them makes compensation contingent on that pursuit succeeding, and the contingency does not fall evenly; the search for interpersonal justice can end by leaving the most disadvantaged victims with nothing (Papayannis, 2023). The mechanism is neutral on its face and regressive in operation, and the neutrality is what makes it durable.

The statutory architecture explains why this particular immunity is so difficult to move. Section 230 grants platforms broad protection from liability for content, and its operation is definitional rather than merely procedural: the safe harbour permits a provider to curate content without acquiring the legal status of a publisher, which fixes the platform as a benign intermediary supplying access rather than an actor with a position on what it circulates (Bossio, Flew, and Meese, 2022). A defendant classified as infrastructure is not declining to answer for a harm. It has been placed outside the category of things that answer.

The child-specific statute the field reflexively reaches for demonstrates the problem rather than escaping it. The dominant federal instrument, the Children's Online Privacy Protection Act, reaches only children under 13 and requires that a platform have actual knowledge that children are using its service, and it is accordingly regarded as outdated and largely ineffective (Ingber and Su, 2024). Both limits are definitional. The age cutoff excludes the adolescents the targeting cases are about, and the actual-knowledge requirement makes a platform's own ignorance the condition of its protection. That is the structure the Eleventh Circuit applied in the Omegle appeal, reached by a different route, and it is the structure North Carolina used. The remedy is available to whoever the definition admits, and the definition was

drafted by the party that would otherwise pay.

That is worse, not better. A child-specific barrier yields to a child-specific statute, which is the remedy the field reflexively reaches for. A general barrier that falls hardest on the population least able to route around it has no child-shaped piece to remove. North Carolina is the historical proof: the eligibility definition that excluded most of the sterilized was not aimed at them either. It was drawn for administrative convenience, and they fell outside it, and the effect was the same as malice would have been.

The children are not being singled out by the doctrine. They are being left where the doctrine's general shape deposits them, which is exactly where mechanism three said they would be: unable to leave a structure the adults built and can exit at will.

One line of scholarship suggests the general barrier is not the end of the analysis, and it is carried here because it cuts toward an optimism the rest of this section does not. Attorneys general in more than forty states across the political spectrum have brought actions against social media companies, alongside a consolidated federal multidistrict litigation and suits filed by school districts, tribal nations, advertisers, and news providers; and the closest precedents, tobacco and opioids, produced not only financial payouts but changes to marketing rules, access and age restrictions, and substantially better documentation of adverse health effects (Lubin, Liu, and Yarnell, 2025). If that reading is right, remediation need not run through the compensation of individuals at all, which is precisely where every mechanism in this paper predicts it will fail. The same authors record the countervailing fact: firms have reduced trust-and-safety capacity and restricted external oversight even as the litigation has expanded. This paper takes no position on which tendency prevails. It observes only that the one route around a general barrier that the historical record supports is a remedy that does not require each child to prove their own case.

IX. LIMITATIONS

The sentinel argument is retrospective and so is its confirmation. Lead is legible now because the adult burden arrived and could be measured against childhood exposure estimated from historical gasoline consumption and survey data. That method is unavailable prospectively. Anyone claiming a present-day sentinel signal is making a prediction, and this paper supplies no procedure for validating one before the adult burden appears.

The population estimates carry the causal assumption their authors state (McFarland and Reuben, 2024). Low-exposure evidence is mixed: Vester and colleagues, with a prospective birth cohort of 262 mother-child dyads, found cumulative childhood lead exposure associated with decreased school-age IQ in unadjusted analyses but not after adjusting for maternal IQ, household income, prenatal vitamin use, home environment, and maternal serum cotinine, concluding those factors may confound the association at low exposure (Vester et al., 2025). That does not disturb the population estimates, which concern a far higher exposure regime, but it marks the claim's boundary.

The pattern is a pattern, not a law. Harms attaching to adult roles produce no early signal in children, and none should be sought: algorithmic rent-setting reaches people who hold leases (*In re RealPage*, 2023), automated hiring screening reaches applicants over forty (*Mobley v. Workday*, 2024), tenant screening reaches rental applicants (*Louis v. SafeRent*, 2023), and pornography consumption is not a hazard that announces itself in a pediatric cohort ahead of an adult one. Nothing here predicts otherwise, and a framework that only ever finds children is not detecting anything.

Several historical figures are contemporary estimates rather than counts, and this paper marks them as such rather than laundering them: the 8,000 infant deaths attributed to swill milk come from an 1858 newspaper editorial (RSAIF MODERN-204); no verified death total exists for Nevada fallout (RSAIF MODERN-156); Fernald enrollment ranges from about 50 to about 90 (RSAIF MODERN-165); Stolen Generations totals range from about 10,000 to 100,000 (RSAIF MODERN-176); and the Guatemala death figure of at least 83 comes with the Commission's caveat that records

do not always establish whether a given death was caused by the experiments (RSAIF MODERN-152).

The rising-floor problem cuts both ways and should be stated against the argument as well as for it. If the targeting counts are artifacts of investigative effort, then this paper cannot claim to know the magnitude of mechanism two, only its structure. The 4,200 in Canada, the 973 in the United States, and the 796 at Tuam are floors whose distance from the true number is unknown and unknowable from the records the institutions kept. A framework that draws confidence from a number it has just characterized as a measurement of somebody's diligence is not entitled to that confidence.

The contemporary analysis rests on litigation records, which document what was alleged and what a court did with it. A dismissal on immunity grounds is not a finding that a product is safe. This paper treats those dismissals as evidence about recourse and not about the products, and that restraint cuts against the argument as often as it supports it.

X. CONCLUSION

The canary metaphor earns its keep in one place. When a hazard is released into shared air, water, or milk, and the children begin to show effects the adults do not yet show, the children are telling the adults something true about the exposure. Leaded gasoline is that case: 170 million children exposed, 824 million IQ points, 151 million excess mental disorders arriving in the adult population decades after the signal was legible in the pediatric one. The Great Smog is that case, and it came with an actual animal sentinel that nobody read either.

Everywhere else the metaphor does damage. Where children were chosen, calling them canaries launders a decision into an accident. Fernald settles it: a federal advisory committee found the doses harmless and the studies wrong anyway. No sentinel case can say that. Where children could not leave, calling them canaries points attention at the adults who could. Aberfan settles that one: 116 children died at 9:15 on a Friday morning because of a timetable, and nothing about their bodies enters the explanation at any point.

The disaggregation tells an advocate which argument to make. Sentinel case: the argument is surveillance, and the interval in which a known signal went unused. Targeting case: the argument is the selection and the belief that licensed it, and the dose is beside the point. Captivity case: the argument is the hazard's presence where attendance is compulsory, and the remedy runs to the hazard, because the child cannot be removed.

Newark is the reason to believe any of this is actionable. It had the same failed chemistry, the same slow denial, the same useless filters, and then it replaced 23,000 lead service lines in under three years by paying for the private half of the pipe. The obstacle that every other city calls structural turned out to be a line item. Captivity is a policy setting with a dial on it, and one city turned it.

The contemporary record suggests the third argument matters most and the field is least equipped to make it. The barriers now standing between a child and a remedy were not built with children in mind. They are general, and they will not move in response to a demonstration that children are uniquely deserving, because they never denied it. North Carolina drew an eligibility line for administrative convenience and 7,000 sterilized people fell outside it. Section 230 was drawn for a different purpose entirely and children fall outside it. The children are simply in the room, and the door does not open from their side.

ACKNOWLEDGMENTS

The author discloses the use of speech-to-text software (Wispr Flow) and Anthropic's Claude as assistive technology for a diagnosed neurodevelopmental condition. All intellectual contributions, research questions, theoretical arguments, methodological decisions, and conclusions are solely the author's own.

Correspondence concerning this article should be addressed to Travis Gilly, Real Safety AI Foundation, Illinois, United States. Email: t.gilly@ai-literacy-labs.org.

REFERENCES

- A.M. v. Omegle.com LLC, 614 F. Supp. 3d 814 (D. Or. 2022).
- Angelilli v. Activision Blizzard, Inc., 781 F. Supp. 3d 691 (N.D. Ill. 2025).
- Asbestos Hazard Emergency Response Act of 1986, 15 U.S.C. §§ 2641–2656.
- Beecher, H. K. (1966). Ethics and clinical research. *New England Journal of Medicine*, 274(24), 1354–1360. <https://doi.org/10.1056/NEJM196606162742405>
- Bell, M. L., Davis, D. L., and Fletcher, T. (2004). A retrospective assessment of mortality from the London smog episode of 1952: The role of influenza and pollution. *Environmental Health Perspectives*, 112(1), 6–8. <https://doi.org/10.1289/ehp.6539>
- Ben-Porath, S. R. (2003). Autonomy and vulnerability: On just relations between adults and children. *Journal of Philosophy of Education*, 37(1), 127–145. <https://doi.org/10.1111/1467-9752.3701009>
- Bhatia, A., Krieger, N., and Subramanian, S. V. (2019). Learning from history about reducing infant mortality: Contrasting the centrality of structural interventions to early 20th-century successes in the United States to their neglect in current global initiatives. *The Milbank Quarterly*, 97(1), 285–345. <https://doi.org/10.1111/1468-0009.12376>
- Black, L., and Ring, S. (2022). Historical gendered institutional violence: A research agenda for criminologists. *Journal of Contemporary Criminal Justice*, 39(1), 17–37. <https://doi.org/10.1177/10439862221138669>
- Bommarito, P. A., and Fry, R. C. (2016). Developmental windows of susceptibility to inorganic arsenic: A survey of current toxicologic and epidemiologic data. *Toxicology Research*, 5(6), 1503–1511. <https://doi.org/10.1039/c6tx00234j>
- Bossio, D., Flew, T., and Meese, J. (2022). Australia’s News Media Bargaining Code and the global turn towards platform regulation. *Policy and Internet*, 14(1), 136–150. <https://doi.org/10.1002/poi3.284>
- Brady, J. P. (2012). Marketing breast milk substitutes: Problems and perils throughout the world. *Archives of Disease in Childhood*, 97(6), 529–532. <https://doi.org/10.1136/archdischild-2011-301299>
- Brando, N. (2022). The moral relevance of social categories: Analysing the case of childhood. *European Journal of Philosophy*, 31(1), 195–208. <https://doi.org/10.1111/ejop.12780>
- Brodie, B. (1991). Baby’s milk: A source of trust between mothers and nurses. *Public Health Nursing*, 8(3), 161–165. <https://doi.org/10.1111/j.1525-1446.1991.tb00749.x>
- Burman, E. (2022). Child as method and/as childism: Conceptual-political intersections and tensions. *Children and Society*, 37(4), 1021–1036. <https://doi.org/10.1111/chso.12615>
- Calise v. Meta Platforms, Inc., 103 F.4th 732 (9th Cir. 2024).
- Capron, A. M. (2017). Building the next bioethics commission. *Hastings Center Report*, 47(S1), S4–S9. <https://doi.org/10.1002/hast.710>
- Carroll, A. E., and Cheng, E. R. (2020). Vaping and youth—first, do no harm. *JAMA Pediatrics*, 174(7), 645. <https://doi.org/10.1001/jamapediatrics.2020.0277>
- Chatters, L. M., Taylor, R. J., and Schulz, A. J. (2022). The return of race science and why it matters for family science. *Journal of Family Theory and Review*, 14(3), 442–462. <https://doi.org/10.1111/jftr.12472>

- Chrostowski, P. C., Foster, S., and Anderson, E. (1991). Human health risks associated with asbestos abatement. *Risk Analysis*, 11(3), 465–481. <https://doi.org/10.1111/j.1539-6924.1991.tb00632.x>
- Claus Henn, B., Austin, C., and Coull, B. A. (2018). Uncovering neurodevelopmental windows of susceptibility to manganese exposure using dentine microspatial analyses. *Environmental Research*, 161, 588–598. <https://doi.org/10.1016/j.envres.2017.12.003>
- Clean Air Act 1956, 4 & 5 Eliz. 2 c. 52 (UK).
- Communications Decency Act § 230, 47 U.S.C. § 230.
- Connecticut Fair Housing Center v. CoreLogic Rental Property Solutions, LLC, 369 F. Supp. 3d 362 (D. Conn. 2019).
- Courtright v. Epic Games, Inc., 795 F. Supp. 3d 1156 (W.D. Mo. 2025).
- Daly, K. (2025). Justice: Word, idea, practice. *Criminology*, 63(4), 707–743. <https://doi.org/10.1111/1745-9125.70009>
- Davis, K. M. (2016). False assurances: The effects of corrosive drinking water and noncompliance with lead control policies in Flint, Michigan. *Environmental Justice*, 9(4), 103–108. <https://doi.org/10.1089/env.2016.0009>
- Doe v. Twitter, Inc., 555 F. Supp. 3d 889 (N.D. Cal. 2021).
- Drozdovitch, V. (2021). Radiation exposure to the thyroid after the Chernobyl accident. *Frontiers in Endocrinology*, 11, 569041. <https://doi.org/10.3389/fendo.2020.569041>
- Erznoznik v. City of Jacksonville, 422 U.S. 205 (1975).
- Faden, R. (1996). The Advisory Committee on Human Radiation Experiments. *Hastings Center Report*, 26(5), 5–10. <https://doi.org/10.2307/3528463>
- Fadus, M. C., Smith, T. T., and Squeglia, L. M. (2019). The rise of e-cigarettes, pod mod devices, and JUUL among youth: Factors influencing use, health implications, and downstream effects. *Drug and Alcohol Dependence*, 201, 85–93. <https://doi.org/10.1016/j.drugalcdep.2019.04.011>
- Fair Labor Standards Act of 1938, 29 U.S.C. §§ 201–219.
- Farrell, A., and Devaney, S. (2007). Making amends or making things worse? Clinical negligence reform and patient redress in England. *Legal Studies*, 27(4), 630–648. <https://doi.org/10.1111/j.1748-121X.2007.00061.x>
- Feir, D. (2016). The long-term effects of forcible assimilation policy: The case of Indian boarding schools. *Canadian Journal of Economics*, 49(2), 433–480. <https://doi.org/10.1111/caje.12203>
- Gallen, J. (2020). Transitional justice and Ireland’s legacy of historical abuse. *Éire-Ireland*, 55(1–2), 35–67. <https://doi.org/10.1353/eir.2020.0002>
- Gubernot, D. M., Boyer, B. L., and Moses, M. S. (2008). Animals as early detectors of bioevents: Veterinary tools and a framework for animal-human integrated zoonotic disease surveillance. *Public Health Reports*, 123(3), 300–315. <https://doi.org/10.1177/003335490812300310>
- Hammer v. Dagenhart, 247 U.S. 251 (1918).
- Hanna-Attisha, M., LaChance, J., and Sadler, R. C. (2016). Elevated blood lead levels in children associated with the Flint drinking water crisis: A spatial analysis of risk and public health response. *American Journal of Public*

- Health*, 106(2), 283–290. <https://doi.org/10.2105/AJPH.2015.303003>
- In re RealPage, Inc., Rental Software Antitrust Litigation (No. II), 709 F. Supp. 3d 469 (M.D. Tenn. 2023).
- Indian Act, R.S.C. 1927, c. 98 (Can.), as amended 1920.
- Ingber, A. S., and Su, C. C. (2024). Protecting children of the TikTok era: A discourse analysis in the absence of law. *Policy and Internet*, 17(2). <https://doi.org/10.1002/poi3.431>
- Jones, M. (2024). The intergenerational legacy of Indian residential schools. *Demography*, 61(6), 1871–1895. <https://doi.org/10.1215/00703370-11679677>
- Kirmizialtin, S. (2007). Conversion in Ottoman Balkans: A historiographical survey. *History Compass*, 5(2), 646–657. <https://doi.org/10.1111/j.1478-0542.2007.00420.x>
- Laventhal, N., Tarini, B. A., and Lantos, J. D. (2012). Ethical issues in neonatal and pediatric clinical trials. *Pediatric Clinics of North America*, 59(5), 1205–1220. <https://doi.org/10.1016/j.pcl.2012.07.007>
- Lead and Copper Rule Improvements, 89 Fed. Reg. 86,282 (Oct. 30, 2024).
- Lee, H., Lee, M., and Warren, J. R. (2022). Childhood lead exposure is associated with lower cognitive functioning at older ages. *Science Advances*, 8(45). <https://doi.org/10.1126/sciadv.abn5164>
- Lee, S., Rees, V. W., and Yossefy, N. (2020). Youth and young adult use of pod-based electronic cigarettes from 2015 to 2019: A systematic review. *JAMA Pediatrics*, 174(7), 714. <https://doi.org/10.1001/jamapediatrics.2020.0259>
- Liu, S. H., Bobb, J. F., and Lee, K. H. (2017). Lagged kernel machine regression for identifying time windows of susceptibility to exposures of complex mixtures. *Biostatistics*, 19(3), 325–341. <https://doi.org/10.1093/biostatistics/kxx036>
- Louis v. SafeRent Solutions, LLC, 685 F. Supp. 3d 19 (D. Mass. 2023).
- Lubin, N., Liu, Y., and Yarnell, A. (2025). Social media harm abatement: Mechanisms for transparent public health assessment. *Annals of the New York Academy of Sciences*, 1549(1), 171–184. <https://doi.org/10.1111/nyas.15345>
- Lutter, C., Hernández-Cordero, S., and Grummer-Strawn, L. M. (2022). Violations of the International Code of Marketing of Breast-milk Substitutes: A multi-country analysis. *BMC Public Health*, 22(1). <https://doi.org/10.1186/s12889-022-14503-z>
- Maslia, M. L., Aral, M. M., and Ruckart, P. Z. (2016). Reconstructing historical VOC concentrations in drinking water for epidemiological studies at a U.S. military base: Summary of results. *Water*, 8(10), 449. <https://doi.org/10.3390/w8100449>
- McFarland, M. J., and Reuben, A. (2024). Contribution of childhood lead exposure to psychopathology in the US population over the past 75 years. *Journal of Child Psychology and Psychiatry*, 66(5), 659–666. <https://doi.org/10.1111/jcpp.14072>
- McFarland, M. J., Hauer, M. E., and Reuben, A. (2022). Half of US population exposed to adverse lead levels in early childhood. *Proceedings of the National Academy of Sciences*, 119(11). <https://doi.org/10.1073/pnas.2118631119>
- McNair, D. M. (2010). Historical and psychological origins of child soldiering in Ba'athist Iraq. *Digest of Middle East*

- Studies*, 19(1), 37–48. <https://doi.org/10.1111/j.1949-3606.2010.00004.x>
- M.H. v. Omegle.com LLC, 122 F.4th 1266 (11th Cir. 2024).
- Mobley v. Workday, Inc., 740 F. Supp. 3d 796 (N.D. Cal. 2024).
- M.P. v. Meta Platforms, Inc., 127 F.4th 516 (4th Cir. 2025).
- Papayannis, D. M. (2023). The morality of compensation through tort law. *Ratio Juris*, 36(1), 3–25. <https://doi.org/10.1111/raju.12369>
- Pickering, E. V., Fu, X., and Melaram, R. (2026). Groundwater chemistry and children’s blood lead levels: A county-wise analysis in the United States. *GeoHealth*, 10(1). <https://doi.org/10.1029/2025GH001670>
- Pilsner, J. R., Parker, M., and Sergeyev, O. (2017). Spermatogenesis disruption by dioxins: Epigenetic reprogramming and windows of susceptibility. *Reproductive Toxicology*, 69, 221–229. <https://doi.org/10.1016/j.reprotox.2017.03.002>
- Rabinowitz, P., Gordon, Z., and Holmes, R. (2005). Animals as sentinels of human environmental health hazards: An evidence-based analysis. *EcoHealth*, 2(1), 26–37. <https://doi.org/10.1007/s10393-004-0151-1>
- Radiation Exposure Compensation Act of 1990, 42 U.S.C. § 2210 note.
- Reif, J. S. (2011). Animal sentinels for environmental and public health. *Public Health Reports*, 126(1_suppl), 50–57. <https://doi.org/10.1177/00333549111260S108>
- Reiners, C. (2011). Clinical experiences with radiation induced thyroid cancer after Chernobyl. *Genes*, 2(2), 374–383. <https://doi.org/10.3390/genes2020374>
- Rothman, D. J. (1982). Were Tuskegee and Willowbrook ‘studies in nature’? *Hastings Center Report*, 12(2), 5. <https://doi.org/10.2307/3561798>
- Ruckart, P. Z., Bove, F. J., and Maslia, M. (2013). Evaluation of exposure to contaminated drinking water and specific birth defects and childhood cancers at Marine Corps Base Camp Lejeune, North Carolina: A case-control study. *Environmental Health*, 12(1), 104. <https://doi.org/10.1186/1476-069X-12-104>
- Savage, L. (2023). Corporal punishment: Why the intentional use of violence against children is still acceptable. *Australian Journal of Social Issues*, 59(3), 667–675. <https://doi.org/10.1002/ajs4.274>
- Schwaba, T., Bleidorn, W., and Hopwood, C. J. (2021). The impact of childhood lead exposure on adult personality: Evidence from the United States, Europe, and a large-scale natural experiment. *Proceedings of the National Academy of Sciences*, 118(29). <https://doi.org/10.1073/pnas.2020104118>
- Tabery, J., Novak, N. L., and Sarafraz, L. (2023). Victims of eugenic sterilisation in Utah: Cohort demographics and estimate of living survivors. *The Lancet Regional Health—Americas*, 19, 100436. <https://doi.org/10.1016/j.lana.2023.100436>
- Terry, M. B., Michels, K. B., and Brody, J. G. (2019). Environmental exposures during windows of susceptibility for breast cancer: A framework for prevention research. *Breast Cancer Research*, 21(1). <https://doi.org/10.1186/s13058-019-1168-2>
- Triantafyllidou, S., and Edwards, M. (2011). Galvanic corrosion after simulated small-scale partial lead service line replacements. *Journal of the American Water Works Association*, 103(9), 85–99. <https://doi.org/10.1002/j.1551-8833.2011.tb11535.x>

- Trueman, B. F., Camara, E., and Gagnon, G. A. (2016). Evaluating the effects of full and partial lead service line replacement on lead levels in drinking water. *Environmental Science and Technology*, 50(14), 7389–7396. <https://doi.org/10.1021/acs.est.6b01912>
- United States v. Darby, 312 U.S. 100 (1941).
- Van Steenberg, J., Wing, P., and D’hulster, K. (2016). The Mamlukization of the Mamluk Sultanate? State formation and the history of fifteenth century Egypt and Syria: Part I — Old problems and new trends. *History Compass*, 14(11), 549–559. <https://doi.org/10.1111/hic3.12357>
- Vester, A., Xu, Y., and Newman, N. (2025). Cumulative childhood lead exposure estimation and school-age IQ in a prospective birth cohort. *Environmental Health*, 25(1). <https://doi.org/10.1186/s12940-025-01244-w>
- Wall, J. (2010). *Ethics in light of childhood*. Georgetown University Press.
- Wiesner-Hanks, M. E. (2008). Do women need the Renaissance? *Gender and History*, 20(3), 539–557. <https://doi.org/10.1111/j.1468-0424.2008.00536.x>
- Winter, S. (2020). Tools for tragedy: Procedures for assessing historic redress claims. *Politics and Policy*, 49(1), 162–185. <https://doi.org/10.1111/polp.12385>
- Winter, S., and Wilkinson, M. (2024). Evaluating paternalism in redress programs. *Politics and Policy*, 52(6), 1399–1416. <https://doi.org/10.1111/polp.12635>
- Wray-Lake, L., Rottenberg, J., and Kennedy, H. (2025). Anti-youth ageism: What it is and why it matters. *Child Development Perspectives*, 19(3), 172–178. <https://doi.org/10.1111/cdep.12540>
- Young-Bruehl, E. (2012). *Childism: Confronting prejudice against children*. Yale University Press.

APPENDIX: HARM CORPUS CASE RECORDS

Real Safety AI Foundation, *Harm Corpus*. Case records cited by identifier:

ANCIENT-106, Mamluk Slave Soldiers.

COLONIAL-118, Canada Residential Schools.

CONTEMPORARY-057, Lead Paint.

CONTEMPORARY-154, Flint Water Crisis.

CONTEMPORARY-186, Newark Lead in Water.

CONTEMPORARY-062, Agent Orange in Vietnam.

CONTEMPORARY-064, Nestlé Infant Formula Marketing.

CONTEMPORARY-184, Washington DC Lead in Water.

CONTEMPORARY-201, Immigration Detention and Family Separation.

CONTEMPORARY-203, Asbestos and Lead in Public Housing.

MEDIEVAL-013, Ottoman Devshirme.

MODERN-139, U.S. Indian Boarding Schools.

MODERN-054, Leaded Gasoline.

MODERN-147, North Carolina Eugenics Board.

MODERN-151, Willowbrook Hepatitis Studies.

MODERN-152, Guatemala Syphilis Experiments.

MODERN-155, Camp Lejeune Contaminated Water.

MODERN-156, Nevada Test Site Downwinders.

MODERN-160, Hanford Green Run and Fernald Exposure.

MODERN-165, Fernald and Wrentham Science Club Experiments.

MODERN-174, Magdalene Laundries in Ireland.

MODERN-175, Tuam Mother and Baby Home.

MODERN-176, Stolen Generations in Australia.

MODERN-188, Great Smog of London.

MODERN-193, Asbestos in Schools and Public Buildings.

MODERN-199, Child Labor in U.S. Mills, Mines, and Canneries.

MODERN-204, Swill Milk in New York.

MODERN-207, Aberfan Colliery Spoil Tip Collapse.

TECH-010, Instagram Teen Mental Health.

TECH-030, Companion Chatbots and Vulnerable User Harm.

TECH-036, Video Game Addiction Litigation.

TECH-037, Twitter CSAM Retention.

TECH-038, Omegle Random Pairing and Child Predation.