




Faculty of Security Studies, University of Belgrade
**International Journal of Contemporary
Security Studies (IJCSS)**



Review Article

The Ingratitude Paradox: A Causal Mechanism of Democratic Discounting of Successful Protection

Burak Oktenli^{1*}

¹ School of Continuing Studies, Georgetown University, 111 Massachusetts Avenue NW, Washington, DC 20001, United States.

* Correspondence: bo236@georgetown.edu.

Received: 18 June 2026; Revised: 15 July 2026; Accepted: 25 July 2026; Published: 24 August 2026.

ABSTRACT

Democratic publics evaluate protective action retrospectively, after the action has run its course and its outcome is known. This article specifies a causal mechanism through which aggregate public evaluation of protective action is systematically biased against action that has succeeded. The mechanism, termed the Ingratitude Paradox, operates in five stages: a threat generates demand for protection; the protective action imposes visible procedural costs; the action succeeds and eliminates the threat; the eliminated threat becomes perceptually invisible; and the visible cost is then weighed against the now-invisible benefit, producing a verdict that the protection was unnecessary or disproportionate. The mechanism's defining property is endogeneity: the threat-eliminating step is also the invisibility-generating step, so the tendency is intrinsic to protective success rather than a contingent failure of communication. This article is a theory-building and mechanism-specification study, not an empirical test. It grounds the mechanism in the securitization literature and the availability-heuristic tradition in risk perception, differentiates it systematically from adjacent concepts, demonstrates that it is symmetric with respect to warrant and therefore explanatory rather than justificatory, illustrates its logic through a documented positive case traced across all five stages, specifies scope conditions through a contrastive case in which it did not fire, and states a falsifiable prediction relating the depth of the discounting verdict to the completeness of threat-elimination. That cross-case prediction remains untested; its systematic evaluation is identified as the natural next step. The contribution is a specified, bounded, and testable addition to the study of threat perception and security politics.

KEYWORDS

Threat perception; securitization; democratic accountability; availability heuristic; retrospective accountability; counterfactual evaluation.

1. Introduction

Democratic publics evaluate protective action after the fact. When a government restricts civil liberties to counter a terrorist threat, commits resources to deter an adversary, or imposes public-health measures against a contagion, the public's evaluation of whether the action was warranted is rendered not while the threat is live but after the protective action has run its course and its out-



Academic Editor:
Prof. Dr. Vladimir M. Cvetković
Copyright: © 2026 by the authors.

Oktenli, B. (2026). The Ingratitude Paradox: A Causal Mechanism of Democratic Discounting of Successful Protection. *International Journal of Contemporary Security Studies*, 2(2), 55-68.

come is visible. This temporal structure of evaluation is unavoidable in democratic politics, where accountability is retrospective by design, but it has a consequence that has not been specified as a mechanism: it systematically biases aggregate public evaluation against protective action that has succeeded. This article specifies that mechanism. The retrospective character of democratic accountability is well established in the voting literature (Fiorina 1981; Healy and Malhotra 2013), as is the finding that such retrospection is imperfect and selectively responsive to recent and visible outcomes (Achen and Bartels 2016).

The core claim is that democratic publics, in the aggregate, discount the value of protective action precisely when the action has succeeded, because success eliminates the threat, and the eliminated threat leaves no perceptible trace against which the cost of protection can be weighed. The public that demanded protection while the threat was live comes, once the threat is gone, to regard the protection's cost as excessive and its justification as overstated, not because the public is fickle or ungrateful in any ordinary sense, but because the object that would vindicate the protection has been removed by the protection itself. The mechanism, which this article terms the Ingratitude Paradox, is a structural tendency in retrospective threat perception under conditions of successful protection. The perceptual foundation of this claim is the availability heuristic, under which the ease of recalling instances of an event governs the estimate of its likelihood (Tversky and Kahneman 1973), and the aggregate expression of that heuristic in public risk judgement and its regulatory consequences (Kuran and Sunstein 1999).

The label requires immediate qualification. "Ingratitude" is an analytical term for a perceptual pattern, and it carries no moral judgment of the public. The mechanism does not hold that publics owe gratitude to protective authorities, that retrospective criticism is illegitimate, or that citizens err in scrutinising the cost of protection. On the contrary, the symmetry constraint developed in Section 5 establishes that the mechanism produces the same discounting verdict whether or not the protection was warranted, so that the label names a regularity in perception, not a failing in citizens. Retrospective scrutiny of the cost of protection is a legitimate and expected feature of democratic politics, and the incentives it creates for officials to avoid blame rather than claim credit are themselves well documented (Weaver 1986).

Three further clarifications fix the article's analytic commitments at the outset.

First, the unit of analysis is aggregate public evaluation: the distribution of a democratic public's retrospective assessment of a protective action, as expressed in opinion trajectories, policy support, and the political standing of the protective measure. Individual cognition supplies the micro-foundation of the mechanism, since the availability heuristic operates in individual minds, but the mechanism's claim is pitched at the aggregate level, where individual perceptual shifts compose into a shift in the public verdict. Downstream political consequences such as electoral punishment or policy reversal are treated as derivative effects that may follow from the aggregate shift under further conditions, not as the mechanism's primary object. These levels are distinguished throughout and are not treated as interchangeable. The aggregation of individual risk judgements into a public verdict, and the social processes that amplify or attenuate it, are treated in the risk-perception literature as distinct from individual cognition (Slovic 1987; Kasperson et al. 1988).

Second, the mechanism is a positive account of how retrospective threat perception is structured, not a normative defence of protective authority. It does not claim that protective action is generally warranted or that governments should be shielded from scrutiny. The normative questions that arise when public fear and regulatory response diverge are treated elsewhere and are not settled here (Sunstein 2005).

Third, the mechanism is not a restatement of the observation that prevention is politically thankless. That observation is familiar; what is new is the specification of the mechanism that produces the thanklessness, the demonstration that the tendency is endogenous to success rather than a contingent communication failure, and the derivation of a falsifiable prediction from it. The claim of novelty is deliberately bounded: the article does not assert that no adjacent literature has touched these phenomena, but that the specific perceptual mechanism, its endogeneity, and its scope conditions have not been assembled into a single bounded and falsifiable specification. Section 2.4 differ-

entiate the mechanism from the adjacent concepts with which it might be confused. The empirical asymmetry is documented most directly in disaster policy, where electorates reward relief spending substantially while returning almost no reward for preparedness spending of comparable value (Healy and Malhotra 2009). A related perceptual finding is that knowledge of an outcome reshapes judgement of how foreseeable that outcome was (Fischhoff 1975).

The article situates the mechanism at the intersection of two literatures that have not been connected in this way. The securitization literature, formulated by the Copenhagen School, analyses how issues are constructed as security threats through discourse and practice. The risk-perception literature, developed in the heuristics-and-biases tradition, analyses how publics estimate the probability and severity of hazards. The Ingratitude Paradox connects them: it uses the risk-perception account of how eliminated threats become perceptually unavailable to explain a securitization-relevant outcome, the de-securitization of a threat by its own successful management. The rally effects observed when a threat is salient, and their decay once it is not, illustrate how quickly the public standing of protective action tracks the visibility of the threat rather than its management (Mueller 1970).

The article proceeds as follows. Section 2 sets out the theoretical background, expands the relevant literature beyond securitization and the availability heuristic, and differentiates the mechanism from adjacent concepts. Section 3 states the article's method and defines its key terms. Section 4 specifies the five-stage mechanism, its endogeneity, and the causal transitions between stages, and presents a causal diagram. Section 5 states the symmetry constraint and the falsifiable prediction. Section 6 traces the mechanism through a documented positive case. Section 7 identifies the scope conditions through a contrastive case. Section 8 discusses implications, limitations, and normative safeguards, and concludes.

2. Theoretical Background

2.1. Securitization and Its Inverse

Securitization theory, formulated by Buzan, Wæver, and de Wilde (1998), analyses security not as an objective condition but as a speech act and social process. An actor makes a securitizing move, framing an issue as an existential threat that requires extraordinary measures; an audience accepts or rejects the move; and securitization succeeds when the audience accepts the framing, licensing action outside the bounds of normal politics. The theory has been extended toward a more explicitly causal register. Guzzini (2011) argued that securitization analysis should be understood as a form of causal-mechanistic explanation, identifying the processes through which security framings are produced and take effect. Balzacq (2011) developed a pragmatic account emphasising the practices and contexts through which securitizing moves succeed or fail.

The securitization literature has concentrated on the construction of security objects: how issues come to be framed as threats and what follows. It has paid less attention to the inverse process, the de-construction of a security object, and least of all to the form that concerns this article: the dissolution of a threat's perceived reality by the success of the action taken against it. Desecuritization, the return of an issue from existential threat to normal politics, has been studied primarily as a discursive achievement that actors bring about through reframing (Wæver 1995). The Ingratitude Paradox identifies a different route to desecuritization, one that operates not through reframing but through the perceptual consequences of successful protective action. When protection succeeds, the threat it addressed becomes perceptually unavailable, and the security framing that justified the protection loses its object. This is desecuritization by success rather than by discourse.

The communicative dimension of security bears directly on the mechanism, because it governs whether an eliminated threat's salience can be preserved. Gbambu and Ibrahim (2026) show, in an analysis of the 2025 Iran–Israel rivalry, that security communication does not merely convey threat but can constitute it, amplifying or dampening perceived danger through rhetorical, symbolic, and diplomatic signalling. Their finding matters here in the reverse direction: if communication can con-

struct threat salience where the material referent is ambiguous, then communication is a candidate instrument for preserving the salience of an averted threat whose material referent has disappeared. This reframes communication as a moderator of the mechanism rather than, as an earlier and cruder view would have it, an instrument categorically incapable of affecting retrospective evaluation. Section 4.3 develops this point.

2.2. Risk Perception and the Availability Heuristic

The perceptual substrate of the mechanism has been studied independently in the risk-perception literature. Tversky and Kahneman (1973) identified the availability heuristic, the cognitive procedure by which the perceived probability or frequency of an event is estimated from the ease with which instances of it come to mind. Kahneman and Tversky (1973) showed that such heuristic estimation systematically departs from the normative use of base-rate information. Events that are recent, vivid, or frequently encountered are judged more probable; events that are distant, abstract, or absent from recent experience are judged less probable, independent of their actual base rates. Slovic (1987), in the foundational programme on the psychometrics of risk, showed that public risk judgements diverge systematically from statistical estimates in ways that track the salience, dread, and familiarity of hazards. Slovic, Fischhoff, and Lichtenstein (1982) documented the specific role of availability in distorting the perceived frequency of low-probability high-consequence events.

The application to protective action is direct but requires care, because availability is not governed by incidence alone. A threat that has been eliminated generates no new instances, and by the availability heuristic its perceived probability declines. But the salience of a threat is also sustained by media coverage, elite rhetoric, commemoration, education, direct experience, institutional communication, and partisan mobilisation, any of which can preserve or reactivate availability independently of incidence. The mechanism therefore does not claim that incidence is the sole determinant of salience; it claims that, absent these countervailing supports, the elimination of new instances drives salience down, and that the countervailing supports are themselves variable and often absent. Sunstein (2005), analysing the availability heuristic in security and precautionary policy, showed that threats are over-weighted when recent and salient and under-weighted when prevented or made distant. Mueller (2006), examining counter-terrorism, argued that the visible costs of security architecture are routinely evaluated against an artificially low estimate of the threat the architecture prevents, precisely because successful prevention renders the threat perceptually invisible.

2.3. Adjacent Political and Perceptual Literatures

Because the mechanism concerns retrospective public evaluation, it is adjacent to several literatures beyond securitization and risk perception, and locating it among them both establishes its distinctiveness and moderates its claim to novelty. The literature on retrospective voting (Fiorina 1981) establishes that publics evaluate incumbents on observed outcomes rather than on counterfactuals, which is the accountability structure within which the mechanism operates. The policy-feedback literature (Pierson 1993) shows that policies reshape the political conditions of their own subsequent evaluation, of which the perceptual disappearance of an averted threat is a specific case. Work on political trust and institutional legitimacy establishes the mediating conditions under which publics extend or withdraw support from protective authorities. Research on counterfactual reasoning establishes the general difficulty publics have in valuing non-events, of which an averted harm is the paradigm instance. The Ingratitude Paradox does not displace these literatures; it specifies a perceptual mechanism that operates within the accountability structure they describe, and its contribution is the precise perceptual pathway, not the general observation that non-events are hard to value.

2.4. Differentiation from Adjacent Concepts

The mechanism must be distinguished from a family of related concepts with which it could be conflated. Table 1 sets out the differentiation. The **prevention paradox** (Rose 1985) concerns the distribution of a prevention's benefits across individuals: a measure that benefits a population offers little to each participant. The **preparedness paradox** concerns the interpretation of a non-event: successful preparation makes a threat appear to have been overstated. **Hindsight bias** concerns the retrospective perception that an outcome was predictable; **outcome bias**, the evaluation of a decision by its outcome rather than its ex ante quality; **omission bias**, the differential weighting of harms from action versus inaction. **Counterfactual neglect** names the general failure to credit averted harm. **Retrospective voting** and **blame avoidance** concern the political incentives that follow from outcome-based evaluation. **Policy feedback** concerns the reshaping of political terrain by prior policy. The **"cry wolf" effect** concerns the erosion of credibility across repeated warnings.

Table 1. Differentiation of the Ingratitude Paradox from adjacent concepts.

Concept	Focus	Distinction from Ingratitude Paradox
Prevention paradox	Distribution of benefit across individuals	Concerns per-individual benefit, not perceptual availability over time
Preparedness paradox	Interpretation of a non-event	Describes non-event character; does not specify the perceptual mechanism
Hindsight / outcome bias	Retrospective judgement of predictability / decisions by outcome	General evaluative biases; not endogenous to protective success
Omission bias	Action vs inaction weighting	Concerns act/omission, not eliminated-threat invisibility
Counterfactual neglect	Failure to credit averted harm	Supplies part of the cognitive basis; not the full endogenous pathway
Retrospective voting	Outcome-based accountability	Accountability structure within which the mechanism operates
Policy feedback	Policy reshaping its own evaluation	Describes structural persistence, not the perceptual pathway
'Cry wolf' effect	Credibility erosion across repeated warnings	Concerns repeated false alarms, not successful prevention

The Ingratitude Paradox is distinct from each because it specifies a particular perceptual pathway with a particular endogeneity: the very success of a protective action removes, through the availability heuristic operating on an eliminated threat, the perceptual counterweight against which its cost would be judged. Several of these concepts describe conditions that compound with the mechanism rather than duplicating it: counterfactual neglect supplies part of its cognitive basis, the preparedness paradox describes its non-event character, and policy feedback describes its structural persistence. The contribution is the assembly of these into a single, bounded, endogenous, and falsifiable specification pitched at the level of aggregate public evaluation, which none of the adjacent concepts individually provides.

3. Method and Definitions

3.1. Method

This article is a theory-building and causal-mechanism specification. It does not test a hypothesis on systematic data; it specifies a mechanism, bounds it, and derives from it a prediction that future work can test. The mechanism was developed abductively, by identifying a recurrent and under-ex-

plained pattern in the evaluation of protective action and specifying the perceptual process that would generate it. The relevant literature was selected by tracing the two established traditions on which the pattern draws, securitization and the availability heuristic, and then extending outward to the adjacent political and perceptual literatures within which retrospective evaluation is studied. The cases are used for two distinct purposes and are labelled accordingly: the positive case in Section 6 is an illustrative process-tracing of the five stages, chosen because its threat, cost, success, and retrospective verdict are all unusually well documented; the contrastive case in Section 7 is used to identify scope conditions, chosen because it is a case of successful protection in which the discounting verdict conspicuously did not occur. Neither case is offered as a systematic test, and the cross-case prediction of Section 5 remains untested.

3.2. Definitions

The specification uses the following terms in defined senses. **Protective action** is a measure undertaken by a political authority to prevent, suppress, or deter a threat to a population. **Successful protection** is protective action that causally eliminates or substantially reduces the threat, as distinct from action followed by a decline in the threat that the action did not produce; the distinction is developed at Stage 3 below. **Procedural cost** is the visible, attributable burden the action imposes: civil-liberties restriction, resource expenditure, institutional disruption, or contested control. **Perceptual invisibility** is the condition in which an eliminated threat, generating no new instances, ceases to be salient in public perception. **Discounting verdict** is the aggregate retrospective assessment that a protective action was unnecessary or disproportionate. The **protector** is the identifiable authority that bears the political consequences of the protective action, which may in some cases be plural or diffuse, a point that conditions the mechanism's scope (Section 7).

4. The Five-Stage Mechanism

4.1. Causal Structure

The Ingratitude Paradox operates through a five-stage causal mechanism. At Stage 1, a threat exists that the public perceives as dangerous and that generates political demand for protective action. At Stage 2, the authority takes protective action, imposing procedural costs. At Stage 3, the protective action causally eliminates or substantially reduces the threat. At Stage 4, the eliminated threat becomes perceptually invisible: absent countervailing supports for its salience, the threat's lack of new instances causes its perceived probability to fall below what the historical record warrants, and its prior seriousness is retrospectively discounted. At Stage 5, the procedural costs, which remain visible, are weighed against the now-invisible threat, producing the assessment that the costs were unnecessary or disproportionate.

Table 2 sets out for each stage the event, the perceptual status of the threat, and the perceptual status of the cost. The table makes visible the crossing point at Stage 4, where the threat's perceptual status inverts from present to absent while the cost's status remains constant. It is this asymmetry of perceptual persistence, cost remaining visible while benefit becomes invisible, that drives the Stage 5 verdict.

Table 2. The five-stage mechanism: perceptual status of threat and cost.

Stage	Event	Threat status	Cost status
1	Threat generates demand for protection	Visible, salient	Not yet incurred
2	Protective action imposes costs	Visible	Visible, contested
3	Action succeeds; threat eliminated	Diminishing	Visible

4	Eliminated threat becomes unavailable	Invisible	Visible
5	Cost weighed against invisible benefit	Invisible	Visible; judged excessive

4.2. *The Causal Transitions*

The five stages are connected by transitions that require explicit statement, because the mechanism's validity depends on them rather than on the plausibility of the sequence as a whole.

The transition into Stage 3 is causal and must be established, not assumed. A threat may decline for reasons unrelated to the protective action: an adversary may abandon an aim, an epidemic may exhaust a susceptible population, an economic shift may reduce a flow. The mechanism requires that the protective action caused or materially contributed to the threat's elimination; where the decline is coincidental, the case is not one of successful protection in the defined sense, and the mechanism's later stages describe a misattribution rather than the discounting of a genuine success. Distinguishing causal elimination from coincidental decline is therefore a precondition of the mechanism, not an incidental detail, and it is the point at which the specification is most demanding of evidence.

The transition from Stage 3 to Stage 4 is not a single step but a compound of three separable transitions, each requiring its own assumption. First, causal elimination reduces the incidence of new instances. Second, reduced incidence lowers perceived probability, but only to the extent that the countervailing supports for salience identified in Section 2.2, media, elite rhetoric, commemoration, education, direct experience, are absent or weak. Third, lowered perceived probability translates into retrospective discounting of the threat's prior seriousness. Each transition is defeasible: strong independent monitoring can break the second, and vivid collective memory can break the third. The mechanism claims not that these transitions are inevitable but that, in the absence of the specific institutional supports named in Section 7, they are the default path.

4.3. *The Endogeneity of Invisibility, Reframed*

The mechanism's distinctive feature is that the eliminating step, Stage 3, is also the invisibility-generating step, Stage 4. There is no separate, exogenous force that renders the threat invisible; the threat's invisibility is a perceptual consequence of the protection's success. This endogeneity distinguishes the mechanism from an ordinary communication problem in its default case. If the discounting verdict arose merely from a failure to explain the threat, it could be remedied by better explanation. The verdict arises instead from the protection working.

It does not follow, however, that communication can never mitigate the tendency, and the earlier formulation of this article overstated the point. Better communication cannot recreate the prevented instances themselves. But credible counterfactual evidence, independent monitoring, intelligence disclosure, simulation, and transparent ex post evaluation can preserve the salience of averted harm by other means, and the finding of Gbambu and Ibrahim (2026) that communication can construct threat salience where the material referent is ambiguous implies that it can also sustain salience where the referent has been removed by success. Endogeneity is therefore best understood as a strong tendency conditional on the absence of such measures, not as an unavoidable outcome of every successful protective action. The mechanism predicts discounting as the default; the scope conditions of Section 7 name the institutional provisions that suspend it.

A corollary central to testability follows. A failed protection, one that does not eliminate the threat, does not generate the tendency in the same form: the persistent threat remains visible and continues to supply the perceptual counterweight against which cost is judged. This does not imply that failed protection escapes public criticism. Failed protection typically attracts intense criticism, for incompetence, wasted cost, or false assurance, but that criticism concerns the competence, proportionality, or implementation of the action, whereas the Ingratitude Paradox concerns the retro-

spective judgment of the necessity of protection as such. The mechanism thus separates the evaluation of necessity, which it addresses, from the evaluation of competence, proportionality, legality, and implementation, which it does not. Its distinctive prediction is that the judgment that protection was unnecessary attaches most strongly to the most complete successes.

5. Symmetry and Falsifiability

5.1. *The Symmetry Constraint*

The five-stage mechanism is symmetric with respect to warrant, and this symmetry conditions everything that follows. The mechanism predicts that publics will discount protective necessity regardless of whether the protection was warranted. Where a threat was genuinely inflated before the action, the post-action verdict that the costs were disproportionate is correct rather than structurally distorted, yet the mechanism produces the same perceptual pattern, cost visible and benefit invisible, whether or not the action was warranted. The mechanism describes a regularity in how successful protection is evaluated; it does not adjudicate the substantive correctness of any protective action.

This symmetry blocks an inferential slide the mechanism might otherwise appear to license, from the claim that a verdict is mechanistically produced to the claim that it is unjustified. The slide is invalid: a structural account of how a verdict is reached is silent on whether the verdict is correct. Mueller (2006) shows that publics have sometimes been right in retrospect that pre-action threat narratives were inflated, and Aguirre, Castro, and López (2026), in an analysis of the securitisation of Mexican cartels in United States security discourse, show that threats may be amplified through securitising discourse in ways that make retrospective criticism substantively justified. The present mechanism is consistent with both findings. It explains why, even where a threat was genuinely inflated, the post-action verdict cannot be epistemically trusted on its own terms, because the mechanism produces the discounting verdict in both directions and supplies no internal signal distinguishing the inflated case from the real-and-eliminated case.

The symmetry requires separating three questions that public retrospection conflates: whether the threat was real, whether the action was warranted *ex ante*, and whether the action causally prevented harm. The mechanism bears only on the perceptual dynamics of the third once success has occurred; it deliberately leaves the first two to independent evidence, and a central implication of the article, developed in Section 8, is that this independent evidence must be secured before the acute phase, because the mechanism erodes it afterward.

5.2. *The Falsifiable Prediction*

The mechanism yields a discriminating prediction. Holding the scope conditions of Section 7 constant, it predicts that the depth of the discounting verdict correlates positively with the completeness of threat-elimination at Stage 3: the more complete the protective success, the deeper the subsequent discounting.

Operationalising the prediction requires defining its terms. **Verdict-depth** is the magnitude of the retrospective shift in aggregate public evaluation, measurable as the decline in policy approval, the increase in the share judging the measure unnecessary, or the rise in opposition to renewal, tracked over a defined post-action interval. **Elimination-completeness** is the degree to which the threat's incidence or capability was reduced, measurable through independent pre-action and post-action threat assessment. **Cost-salience**, the rival driver defined below, is the visibility and magnitude of the procedural cost. The **temporal interval** over which the verdict is measured must be specified in any test, since the mechanism predicts that discounting emerges after success is established and may deepen as salience decays; the article does not fix a single interval but requires that any test declare one.

A simple positive correlation between elimination-completeness and verdict-depth would not by itself establish the mechanism, because elimination-completeness may be endogenous to policy in-

tensity: more intrusive actions may both eliminate threats more completely and impose more visible costs. A credible test must therefore control for cost magnitude, prior threat severity, partisanship, political trust, media framing, policy duration, personal exposure, and economic conditions, and must address the possible endogeneity of elimination-completeness to policy intensity through design or instrumentation.

The prediction is discriminating because a rival account yields a different pattern. If the discounting verdict were driven by cost-salience rather than by elimination-completeness, verdict-depth would correlate with cost-salience and be independent of elimination-completeness, controlling for the confounders above. That is a different, measurable pattern, and were it observed the mechanism as specified would be disconfirmed. Cost-salience is not, however, the only alternative: partisan cueing, elite framing, institutional trust, temporal decay, policy fatigue, civil-liberties concern, and attribution uncertainty may each drive retrospective discounting independently, and a complete test would treat these as competing or moderating accounts rather than reducing the field to two.

5.3. A Causal Model

Figure 1 presents the mechanism as a causal diagram, showing the five-stage pathway, the compound Stage 3-to-4 transition, the countervailing supports for salience that moderate it, the rival explanations that a test must exclude, and the confounders that a test must control. The diagram distinguishes the mechanism's core pathway (solid) from moderators (dashed) and from rival drivers (dotted), making explicit that the mechanism is one pathway among several that a test must separate rather than the sole determinant of the verdict.




Figure 1. The Ingratitude Paradox: five-stage causal pathway (solid), salience moderators (dashed), and rival drivers and confounders (dotted).

6. A Positive Case: Y2K Remediation

The mechanism's logic can be traced through a documented positive case in which all five stages are unusually well recorded: the remediation of the Year 2000 computer problem. The case is presented as an illustrative process-tracing, not as a systematic test, and it is chosen because the threat,

the cost, the causal role of the remediation, and the retrospective verdict are each independently documented.

Stage 1: threat and demand. In the late 1990s, the practice of representing calendar years with two digits created a documented risk that date-dependent computer systems, in finance, air traffic control, utilities, and government, would malfunction at the century rollover (United States Senate Special Committee on the Year 2000 Technology Problem 1999). The threat generated broad political demand for protective action, including national legislation and coordinated public and private remediation.

Stage 2: visible cost. Remediation imposed large and well-documented costs. Estimates placed United States expenditure at roughly one hundred billion dollars and worldwide expenditure in the range of three hundred to four hundred fifty billion dollars (Quiggin 2005). The cost was concentrated, attributable, and widely reported, satisfying the procedural-cost condition.

Stage 3: causal success. The critical and contestable transition is whether remediation causally averted harm rather than merely preceding a non-event. The evidence for causal contribution is that systems and jurisdictions that were remediated passed the rollover without major failure, that specific documented faults were identified and corrected in the remediation process, and that the small number of failures that did occur clustered in unremediated systems (Quiggin 2005). This does not establish that every dollar was necessary, and Quiggin argues that much expenditure was misdirected; but it supports the claim that remediation materially contributed to the near-absence of failure, which is the causal condition Stage 3 requires. The case also illustrates the symmetry constraint: the credible argument that the threat was partly inflated does not exempt the case from the mechanism, because the mechanism operates whether or not the action was warranted.

Stage 4: perceptual invisibility. Because the rollover passed almost without incident, the threat generated no vivid instances, and its perceived prior probability collapsed. The public response was studied at the time within the social-amplification-of-risk framework (Kasperson et al. 1988, as applied to Y2K), which documents precisely the perceptual dynamics the mechanism specifies: a threat whose realisation is prevented leaves no experiential trace to sustain its salience.

Stage 5: discounting verdict. The retrospective verdict is well attested. The remediation came to be widely characterised in public discourse as a hoax, a waste, or an overblown scare, and the scholarly retrospective records the same judgment, that the expenditure was, in hindsight, largely regarded as unnecessary (Quiggin 2005). The verdict attached to the completeness of the success: precisely because almost nothing failed, the protection appeared unnecessary. The case thus exhibits the mechanism's signature inversion, in which the completeness of the success is the condition of its retrospective discounting, and it does so with the causal, cost, and perceptual evidence that the specification requires.

The Y2K case is offered as one traced instance, not as confirmation across cases. Its value is to show that the five-stage sequence, including the demanding Stage 3 causal condition, can be documented in a real episode, and to give the falsifiable prediction a concrete referent: a fuller test would place Y2K within a sample of protective episodes scored for elimination-completeness and verdict-depth.

7. Scope Conditions: A Contrastive Case

A mechanism that predicts in all cases predicts in none. The mechanism's content depends on identifying the conditions under which it fires, and a case in which successful protection did not generate the discounting verdict supplies the needed contrast. The international regime for the protection of the stratospheric ozone layer, established by the 1985 Vienna Convention and the 1987 Montreal Protocol, is such a case (Parson 2003). The protective action addressed a documented atmospheric threat through coordinated regulation that imposed substantial industrial-adjustment costs; the action succeeded, ozone-depleting substances were phased out, and the ozone layer has been recovering, yet the success did not generate a public verdict that the costs had been unjustified or the threat overstated.

A caveat is required before the case can bear analytical weight. The claim that the mechanism did not fire requires evidence of public evaluation, not merely the absence of a prominent backlash. The available evidence is that the Montreal regime retained broad public and political legitimacy through and after its success, that it was repeatedly strengthened by subsequent amendments rather than contested as overbuilt, and that public understanding of ozone protection remained positive; the regime's continued expansion is itself weak evidence of a sustained favourable public evaluation. This evidence is suggestive rather than conclusive, and the contrast is accordingly offered as the basis for preliminary propositions rather than established scope conditions.

Three features distinguish the ozone case, and each is proposed as a candidate condition on the mechanism's firing rather than as a demonstrated law. First, the threat trajectory was continuously and independently monitored by atmospheric observation, which decoupled public perception of the threat from the political narrative and bounded the Stage 4 invisibility by an independent evidentiary stream. Second, the costs were borne predominantly by industrial actors who developed commercially viable substitutes, so the cost was not experienced as a recurring, salient burden by the general public whose verdict is at issue. Third, the architecture was treaty-based and multilateral, which diffused the protector role.

The third condition is the least secure and is offered most tentatively. Even within domestic protective policy, responsibility is often shared among governments, agencies, experts, and international bodies, so a "single identifiable protector" is rarely a clean category; the condition is better stated as the degree to which a single protector is identifiable as bearing the political consequences of success, a matter of degree that alters the mechanism's intensity rather than a binary that switches it on or off.

The relationship among the three conditions is likewise proposed rather than settled. The article does not claim that all three are jointly necessary, individually necessary, or merely contributory; it proposes them as factors that plausibly alter the probability and intensity of the mechanism, to be adjudicated by the comparative testing the specification is intended to enable. Deriving firm scope conditions from a single contrastive case would risk overfitting, and the propositions are framed to avoid that error. A useful further contrast is offered by persistent, independently monitored risks such as climate change (Mančić 2025), whose perceptual availability may not disappear after individual protective interventions precisely because continuous monitoring and recurring salience sustain it, which is consistent with the first proposed condition.

Each proposed condition carries a constructive corollary: independent monitoring of the threat, distribution of protective cost away from a salient recurring burden, and multilateral rather than singular custody each attenuate the mechanism, and the ozone regime's apparent escape is attributable to its having combined all three.

8. Discussion and Conclusion

The Ingratitude Paradox is a specified, bounded, and testable causal mechanism, offered as a specification rather than a test. It explains why democratic publics, in the aggregate, discount protective action whose success has eliminated the motivating threat; it locates the explanation in an endogenous relation between success and threat-invisibility conditional on the absence of countervailing supports; and it constrains itself by a symmetry argument that prevents its use as a justification for any protector.

Three implications follow. The first concerns securitization theory: the mechanism identifies a route to desecuritization that the literature's focus on discursive construction has left under-examined, the dissolution of a security object by the perceptual consequences of its own successful management, which carries the political consequence of retrospective delegitimation of the protective authority. The second concerns democratic accountability: standard accounts evaluate protective action on the visibility of its costs and are structurally unequipped to weigh those costs against an averted harm that success has rendered invisible, so that retrospective accountability for successful protection is biased in a specifiable direction, deepest for the most successful actions. The third

concerns deterrence and the general problem of the unobservable non-event: successful deterrence produces a non-event without eliminating the adversary's underlying capability or intention, and Edmond (2026), analysing deterrence effectiveness in the Ukraine context, illustrates the difficulty of evaluating a protective success whose product is the absence of an event. This points to a needed refinement: the mechanism operates differently across threat elimination, suppression, risk reduction, deterrence, and prevention, because a threat that is deterred, latent, recurring, or adaptive remains capable of resupplying its own salience in a way that a fully eliminated threat does not.

Applicability beyond democracies. The mechanism is specified for democratic publics because its object is aggregate public evaluation expressed through open accountability. Retrospective evaluation also conditions authoritarian legitimacy, though through different institutional channels, and the perceptual core of the mechanism, the invisibility of averted harm, is not regime-specific; how it operates where public evaluation is mediated by controlled information environments is a question the present specification raises but does not resolve.

Moderators. The mechanism's intensity is conditioned by moderators that a fuller account and any empirical test should incorporate: institutional trust, political polarisation, direct threat exposure, the media environment, transparency, and the availability of independent evidence. Research on the determinants of security perception (Cvetković et al. 2025a) shows that institutional, socio-economic, and demographic variables structure how publics perceive threat and security, and supplies a template for operationalising these moderators in a future test; work on community-based risk reduction (Cvetković et al. 2025b) similarly identifies institutional trust and risk communication as conditions under which preventive measures retain or lose public legitimacy, which bear directly on whether the Stage 4 transition completes.

Normative safeguards. The mechanism carries a normative risk that must be stated directly: a government could invoke it to dismiss legitimate retrospective criticism as a mere cognitive bias, using the theory to insulate protective authority from accountability. The symmetry constraint forbids this inference at the level of theory, since the mechanism produces the discounting verdict whether or not the protection was warranted and therefore cannot certify any particular protection as justified. But theoretical prohibition is not an institutional safeguard, and the correct response to the mechanism is not deference to protective authority but the construction of independent evidence that survives the perceptual asymmetry. Independent *ex ante* threat assessment, proportionality review, sunset clauses, judicial oversight, public reporting, and *ex post* evaluation that assesses both the reality of the original threat and the causal effectiveness of the intervention are the institutional means by which a public can hold successful protection accountable without either crediting it automatically or dismissing its costs reflexively. The mechanism's practical value is to show why such independent records must be built before the acute phase, since it is precisely the success of protection that later erodes the evidence on which fair evaluation depends.

Limitations. The account's limitations follow from its character as a specification rather than a test. It establishes the mechanism's structure, perceptual foundation, causal transitions, and candidate scope conditions, and derives a prediction, but it does not supply the systematic cross-case data that would confirm or disconfirm the prediction; the cross-case prediction remains untested. The conceptual differentiation of Section 2.4, though systematic, does not exclude the possibility that the mechanism overlaps in part with concepts not considered. The two cases are illustrative and contrastive rather than a sample, and case selection was purposive. The level of analysis is aggregate public evaluation, and the mechanism's translation into individual-level attitude change and into downstream electoral or policy consequences requires further specification and evidence. And the normative-misuse risk, though addressed by the symmetry constraint and the proposed safeguards, cannot be eliminated by theoretical means alone. These limitations mark the specification's boundaries and identify the systematic comparative testing that is its natural next step and the condition of its confirmation.

Statements and Declarations

Author Contributions: The sole author was responsible for conceptualization, analysis, and writing (original draft, review, and editing). **Funding:** This research received no external funding. Con-

flicts of Interest: The author declares no conflict of interest. Data Availability: No new data were created or analysed in this study. Data sharing is not applicable to this article. Use of AI: In preparing this manuscript the author used Claude (Anthropic), a large language model, for language editing only. The argument, source selection, analysis, and conclusions are the author's own; all citations and conceptual claims were independently checked by the author, who takes full responsibility for the content.

9. References

- Achen, C. H., & Bartels, L. M. (2016). *Democracy for realists: Why elections do not produce responsive government*. Princeton, NJ: Princeton University Press.
- Aguirre, J., Castro, E. G., & López, J. T. (2026). Mexican cartels and U.S. national security: Between criminal economies and securitization discourses. *International Journal of Contemporary Security Studies*, 2(1), 61–70.
- Balzacq, T. (2011). *Securitization theory: How security problems emerge and dissolve*. London: Routledge.
- Buzan, B., Wæver, O., & de Wilde, J. (1998). *Security: A new framework for analysis*. Boulder, CO: Lynne Rienner.
- Cvetković, V. M., Aleksova, B., Renner, R., Gačić, J., Ivanov, A., & Milašinović, S. (2025b). Community-based disaster risk reduction: Overcoming barriers to build stronger communities. *International Journal of Disaster Risk Management*, 7(2), 113–130. <https://doi.org/10.18485/ijdrm.2025.7.2.7>
- Cvetković, V. M., Lipovac, M., Renner, R., Stanarević, S., & Raonić, Z. (2025a). A predictive framework for understanding multidimensional security perceptions among students in Serbia: The role of institutional, socio-economic, and demographic determinants of sustainability. *Sustainability*, 17(11), 5030. <https://doi.org/10.3390/su17115030>
- Edmond, C. (2026). The balance of power in Ukraine and deterrence effectiveness against nuclear-armed Russia. *International Journal of Contemporary Security Studies*, 2(1), 1–16. https://doi.org/10.18485/fb_ijcss.2026.2.1.1
- Fiorina, M. P. (1981). *Retrospective voting in American national elections*. New Haven, CT: Yale University Press.
- Fischhoff, B. (1975). Hindsight is not equal to foresight: The effect of outcome knowledge on judgment under uncertainty. *Journal of Experimental Psychology: Human Perception and Performance*, 1(3), 288–299. <https://doi.org/10.1037/0096-1523.1.3.288>
- Gbambu, A. R., & Ibrahim, M. G. (2026). Security communication and the security dilemma: Insights from the Iran–Israel rivalry. *International Journal of Contemporary Security Studies*, 2(1), 17–28. https://doi.org/10.18485/fb_ijcss.2026.2.1.2
- Guzzini, S. (2011). Securitization as a causal mechanism. *Security Dialogue*, 42(4–5), 329–341. <https://doi.org/10.1177/0967010611419000>
- Healy, A., & Malhotra, N. (2009). Myopic voters and natural disaster policy. *American Political Science Review*, 103(3), 387–406. <https://doi.org/10.1017/S0003055409990104>
- Healy, A., & Malhotra, N. (2013). Retrospective voting reconsidered. *Annual Review of Political Science*, 16, 285–306. <https://doi.org/10.1146/annurev-polisci-032211-212920>
- Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological Review*, 80(4), 237–251. <https://doi.org/10.1037/h0034747>
- Kasperson, R. E., Renn, O., Slovic, P., Brown, H. S., Emel, J., Goble, R., Kasperson, J. X., & Ratick, S. (1988). The social amplification of risk: A conceptual framework. *Risk Analysis*, 8(2), 177–187. <https://doi.org/10.1111/j.1539-6924.1988.tb01168.x>
- Kuran, T., & Sunstein, C. R. (1999). Availability cascades and risk regulation. *Stanford Law Review*, 51(4), 683–768. <https://doi.org/10.2307/1229439>

- Mančić, T. (2025). Climate change as a security challenge, risk and threat of the 21st century and its consequences on critical infrastructure. *International Journal of Contemporary Security Studies*, 1(1), 191–204. https://doi.org/10.18485/fb_ijcss.2025.1.1.14
- Mueller, J. E. (1970). Presidential popularity from Truman to Johnson. *American Political Science Review*, 64(1), 18–34. <https://doi.org/10.2307/1955610>
- Mueller, J. (2006). *Overblown: How politicians and the terrorism industry inflate national security threats, and why we believe them*. New York: Free Press.
- Parson, E. A. (2003). *Protecting the ozone layer: Science and strategy*. Oxford: Oxford University Press.
- Pierson, P. (1993). When effect becomes cause: Policy feedback and political change. *World Politics*, 45(4), 595–628. <https://doi.org/10.2307/2950710>
- Quiggin, J. (2005). The Y2K scare: Causes, costs and cures. *Australian Journal of Public Administration*, 64(3), 46–55. <https://doi.org/10.1111/j.1467-8500.2005.00451.x>
- Rose, G. (1985). Sick individuals and sick populations. *International Journal of Epidemiology*, 14(1), 32–38. <https://doi.org/10.1093/ije/14.1.32>
- Slovic, P. (1987). Perception of risk. *Science*, 236(4799), 280–285. <https://doi.org/10.1126/science.3563507>
- Slovic, P., Fischhoff, B., & Lichtenstein, S. (1982). Facts versus fears: Understanding perceived risk. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgment under uncertainty: Heuristics and biases* (pp. 463–489). Cambridge: Cambridge University Press.
- Sunstein, C. R. (2005). *Laws of fear: Beyond the precautionary principle*. Cambridge: Cambridge University Press.
- Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*, 5(2), 207–232. [https://doi.org/10.1016/0010-0285\(73\)90033-9](https://doi.org/10.1016/0010-0285(73)90033-9)
- United States Senate Special Committee on the Year 2000 Technology Problem. (1999). *Investigating the impact of the Year 2000 problem*. Washington, DC: U.S. Government Printing Office.
- Wæver, O. (1995). Securitization and desecuritization. In R. D. Lipschutz (Ed.), *On security* (pp. 46–86). New York: Columbia University Press.
- Weaver, R. K. (1986). The politics of blame avoidance. *Journal of Public Policy*, 6(4), 371–398. <https://doi.org/10.1017/S0143814X00004219>