

Appendix A. A Formal Model of Externality Correction under Producer Lobbying

Formal addendum to the working paper “The Peace Dividend Mandate”

Christopher Louis · Peace Economy Institute

September 2026

Abstract

This note supplies the formal core underlying the three theoretical extensions of the working paper “The Peace Dividend Mandate.” A single model of a defense contractor, a procuring government, and an accountable electorate yields, as connected results rather than separate assertions: the failure of linear Pigouvian correction against an amplified externality (§A.2); the necessity and sufficiency of a convex, conflict-indexed burden (§§A.2–A.3); the inversion of the producer’s lobbying incentive under that burden, including in the high-conflict regime where a static instrument inverts back toward war (§A.3); and the structural role of cost pass-through in correcting a second, nested information asymmetry (§A.4). The active-conflict ratio is derived as a calibrated Pigouvian rate rather than chosen for rhetorical force (§A.5). The amplification assumption on which §§A.2–A.3 rest is examined and weakened against the empirical record in §A.6; the scope conditions of §A.7 mark precisely where each result depends on empirical magnitudes not yet established; and §A.8 records how the pass-through result is carried in Section V of the working paper. The propositions are offered for formal scrutiny and refutation; they are claims about structure, and structure is the only thing that can be proven in advance of implementation.

A.1 Setup and notation

We model the interaction of three parties: a representative defense contractor (the *firm*), the *government* that procures from it, and the *citizens* to whom that government is accountable. Let $H \geq 0$ be a scalar index of the harm-generating activity the firm’s products sustain: jointly the scale and lethality of weapons output and the conflict intensity it underwrites. We call H the *externality level*.

The firm. The firm’s operating margin (revenue net of production cost) from activity at level H is $m(H)$, with $m'(H) > 0$: at the margin, escalation is profitable. We say the externality is in the *amplified regime* on an interval when m is strictly convex there, $m''(H) > 0$: the marginal profitability of harm is itself rising in harm, as when wartime scarcity lifts both volumes and unit margins.

Social harm. The external social cost of activity at level H is $D(H)$, with $D'(H) > 0$. Write $\lambda(H)$ for the *externality multiplier*: the marginal external social cost $D'(H)$ expressed per dollar of the firm's weapons spend at conflict level H . Section A.5 calibrates it from the documented record.

The correction. The Mandate imposes a compliance burden $B(H)$ on the firm. We compare two functional forms:

- a *static* (linear) Pigouvian charge, $B(H) = \tau H$, with constant marginal burden $B'(H) = \tau$;
- the Mandate's *dynamic ratio*, $B(H) = \rho(H)s(H)$, where $s(H)$ is the firm's weapons spend and $\rho(H)$ the conflict-indexed ratio with $\rho'(H) > 0$. Because ρ rises with conflict, the marginal burden $B'(H)$ rises with H under the mild curvature conditions stated in Proposition A.2, so that B is convex where a linear charge is not.

Pass-through. A share $\varphi \in [0, 1]$ of the burden is recovered in procurement prices and so borne by the government; the firm retains $m(H) - (1 - \varphi)B(H)$. We treat φ as a structural parameter of the procurement relationship rather than a free choice, and in §A.4 show it is the hinge on which the entire correction turns.

A.2 The amplified externality and the failure of linear correction

Throughout this section the firm bears the burden in full ($\varphi = 0$); pass-through is introduced in §A.3.

Definition A.1. A negative externality is *amplified* on an interval $\mathcal{H}$ if the producer's operating margin is strictly increasing and strictly convex in the externality level there: $m'(H) > 0$ and $m''(H) > 0$ for all $H \in \mathcal{H}$. Equivalently, the producer's marginal profit from harm, $m'(H)$, is itself increasing in harm.

Remark A.1. In the textbook externality the producer profits from *output*, and harm is a fixed byproduct of output; facing a per-unit-harm charge, the producer has every incentive to *reduce* harm intensity. The amplified case is the one in which harm intensity is itself the profitable margin (a more destructive product commands a higher price) so that $\partial(\text{profit})/\partial(\text{harm}) > 0$ directly. The distinction is not a matter of the *magnitude* of a Pigouvian rate but of its *functional form*, as the next two results show. To the author's knowledge this class has not been separately distinguished in the externality literature; the claim of novelty is recorded here as an invitation to identify prior treatment, not as a settled negative.

Proposition A.1 (Linear correction fails under amplification). *Let m be strictly convex on $[0, \bar{H}]$ and let the correction be linear, $B(H) = \tau H$. Then firm profit $\pi(H) = m(H) - \tau H$ is strictly convex, so it attains its maximum at a corner of $[0, \bar{H}]$. The firm selects maximal harm $\bar{H}$ whenever*

$$\frac{m(\bar{H}) - m(0)}{\bar{H}} > \tau.$$

Deterring the firm below $\bar{H}$ therefore requires $\tau \geq (m(\bar{H}) - m(0))/\bar{H}$, a threshold fixed by the average return to harm over the entire admissible range, not by local marginal damage. As $\bar{H}$ grows, the required rate rises without bound whenever m' does.

Proof. Since $\pi''(H) = m''(H) > 0$, π is strictly convex and its maximum over the compact interval $[0, \bar{H}]$ is attained at an endpoint. Comparing endpoints, $\pi(\bar{H}) \geq \pi(0)$ iff $m(\bar{H}) - \tau\bar{H} \geq m(0)$, i.e. iff $(m(\bar{H}) - m(0))/\bar{H} \geq \tau$. The deterrence threshold and its divergence in $\bar{H}$ follow at once. $\square$

The standard Pigouvian prescription (set the rate equal to marginal damage) presupposes a concave private objective with an interior optimum, and neither survives amplification. A linear instrument either fails to bind, leaving the firm at the ceiling, or must be pitched at the average return across the whole range, an extreme and range-dependent rate. What is required instead is a correction whose *marginal* burden rises with harm.

Proposition A.2 (Convex correction restores an interior optimum). *Let B be twice differentiable with $B''(H) > m''(H)$ for all $H \in [0, \bar{H}]$. Then $\pi(H) = m(H) - B(H)$ is strictly concave with a unique maximizer $\hat{H}$, interior whenever $m'(0) > B'(0)$ and $m'(\bar{H}) < B'(\bar{H})$, characterized by $m'(\hat{H}) = B'(\hat{H})$, and $\hat{H}$ falls as B is made uniformly steeper. The dynamic ratio $B(H) = \rho(H)s(H)$ generates such convexity whenever $\rho'(H) > 0$, $s'(H) \geq 0$, and ρ and s are weakly convex, with at least one of ρ'' , s' , s'' strictly positive.*

Proof. $\pi'' = m'' - B'' < 0$ gives strict concavity and hence a unique maximizer. The boundary-sign conditions place it in the interior, where the first-order condition reads $m' = B'$. Raising B' pointwise lowers π' pointwise, shifting $\hat{H}$ to the left. Finally $B'' = \rho''s + 2\rho's' + \rho s''$, which is strictly positive under the stated conditions, so B is convex. $\square$

A.3 The lobby inversion principle

So far H has been the firm's own choice. But the level of conflict is also shaped politically, and the firm can spend to move it: through threat inflation, opposition to arms control, and the financing of favorable analysis. We now ask which *direction* of political effort raises the firm's payoff.

Definition A.2. The firm's *lobbying direction* at externality level H is the sign of $d\pi/dH$. A correction mechanism *inverts lobbying* at H if $d\pi/dH < 0$, so that the firm's privately optimal political effort is directed at *reducing* H .

Proposition A.3 (Lobby inversion). *With firm payoff $\pi(H) = m(H) - (1 - \varphi)B(H)$, the firm's lobbying direction is the sign of $m'(H) - (1 - \varphi)B'(H)$, and lobbying inverts at H iff*

$$(1 - \varphi)B'(H) > m'(H).$$

Under the static charge $B'(H) = \tau$, inversion holds only on the set $\{H : m'(H) < (1 - \varphi)\tau\}$. Because m' is increasing in the amplified regime, this set is an interval $[0, H^)$, so inversion fails for all $H > H^*$, that is, in precisely the high-conflict region. Under the dynamic ratio, $B'(H)$ is increasing; if the schedule satisfies $(1 - \varphi)B'(H) > m'(H)$ across the escalation range, lobbying inverts throughout it, including at the highest conflict levels.*

Proof. Differentiating π gives the sign condition directly. In the static case $(1 - \varphi)\tau$ is constant and m' strictly increasing, so $\{m' < (1 - \varphi)\tau\} = [0, H^*)$ with $m'(H^*) = (1 - \varphi)\tau$. In the dynamic case, B' increasing together with the stated inequality yields $d\pi/dH < 0$ pointwise on the range. $\square$

This is the political-economy payoff, and it rhymes with Proposition A.1: the same convexity that defeats a linear *tax* at the production margin defeats a flat *tax* at the lobbying margin, and the same dynamic ratio repairs both. A correction that climbs with conflict is not adopted for emphasis; it is the unique form that keeps the firm's political incentive pointed at peace in the very regime where a static instrument inverts it back toward war.

Remark A.2 (Defensive baseline). One does not want inversion to bite at every H : legitimate defensive provision should remain viable. Let H_d denote the defensive baseline of Section 3.4a of the working paper. The schedule is properly calibrated when $(1 - \phi)B'(H) \leq m'(H)$ for $H \leq H_d$, defense is not penalized at the margin, and $(1 - \phi)B'(H) > m'(H)$ for $H > H_d$: escalation beyond it is. The marginal burden then crosses the marginal margin exactly at the defensive threshold, and the ratio tiers of Section 3.3 of the working paper are a discrete approximation to a schedule whose crossing point is H_d .

A.4 The nested information failure, and why pass-through completes the correction

The pass-through that complicates the inversion of §A.3 is, properly understood, the instrument that closes a second market failure. We make the point precise.

Definition A.3. A *nested information failure* obtains when an externality is sustained by two stacked private-information relationships. At the lower layer the producer holds information the intermediate principal cannot verify: here, the contractor's true costs, capabilities, and the genuine necessity of procurement, the asymmetry documented empirically by Stiglitz and Bilmes.¹ At the upper layer the intermediate principal holds information the ultimate principal cannot verify (here, the government's knowledge of the true cost and conduct of conflict, hidden from citizens). The relationships are *nested* because the intermediate party is the lower relationship's principal and the upper relationship's agent.

Let $g(H)$ denote the government's own payoff from conflict capacity (security, deterrence, and the domestic political returns to a credible posture) and suppose it is independently increasing in the relevant range, $g'(H) > 0$.

Proposition A.4 (Single-layer corrections are insufficient). *Let the externality be supported by a nested failure with $g'(H) > 0$. A correction acting only at the lower layer (reforming the procurement contract so the firm internalizes $B(H)$) leaves the upper-layer asymmetry intact: the government's choice over H still solves $\max_H g(H) - (\text{procurement cost})$, with the social cost still unobserved by citizens, so the externality persists through the uncorrected layer. Symmetrically, a transparency reform at the upper layer that does not price harm at the lower layer leaves the contractor's incentives unchanged. Restoring constraint requires action at both layers at once.*

Proof. A lower-layer reform alters the firm's payoff but neither the government's objective nor the citizens' information set: the government still chooses H to maximize $g(H)$ net of procurement

¹Joseph E. Stiglitz and Linda J. Bilmes, *The Three Trillion Dollar War: The True Cost of the Iraq Conflict* (New York: W.W. Norton, 2008).

cost under unchanged upper-layer observability, so its optimum in H —and the externality it sustains—survives the reform. Conversely, upper-layer transparency without lower-layer pricing leaves the firm’s first-order condition in H unchanged. No single-layer instrument therefore removes the distortion; the corollary that follows shows how a lower-layer burden, passed through, reaches both layers at once. $\square$

Corollary A.5 (Pass-through is the upper-layer instrument). *Suppose the lower-layer burden $B(H)$ is passed through to procurement prices at share ϕ . Then ϕ governs both the locus of lobby inversion and the strength of the upper-layer correction.*

(i) Locus of inversion. *As $\phi \rightarrow 0$ the firm bears the burden and inverts when $B'(H) > m'(H)$ (Proposition A.3). As $\phi \rightarrow 1$ the burden ceases to bear on the firm at all: its retained payoff no longer depends on B , its lobbying direction reverts to the uncorrected sign of $m'(H)$, and the bite transfers to the government, whose payoff $g(H) - \phi B(H)$ inverts when $\phi B'(H) > g'(H)$. For intermediate ϕ both margins are partially corrected.*

(ii) Closure of the upper layer. *The same pass-through writes the priced social cost into the public budget as an entry increasing in H , rendering at the citizen layer the war cost previously hidden: exactly the observability the upper-layer asymmetry lacked.*

Hence the pass-through that weakens firm-level inversion is the very instrument that closes the second information layer.

Proof. The inversion conditions in (i) are Proposition A.3 applied to the firm’s retained payoff $m - (1 - \phi)B$ and to the government’s payoff $g - \phi B$, respectively; as $\phi \rightarrow 1$ the firm’s condition $(1 - \phi)B' > m'$ fails for every H (the left side vanishes while $m' > 0$), so the correction no longer bears on the firm at the lobbying margin, and the government’s condition reduces to $B' > g'$. For (ii), the procurement price rises by $\phi B(H)$, a public-account entry increasing in H ; observability at the citizen layer is restored to the extent $\phi > 0$. $\square$

The apparent tension—that pass-through dissolves the firm’s incentive to lobby for peace—is the model reporting that the correction is *dual-layer*. At low ϕ the mechanism acts on the contractor; at high ϕ it acts on the government and simultaneously makes the cost of conflict legible to the electorate, while the firm’s residual interest in conflict, no longer offset by a retained burden, is the object of the anti-gaming provisions of Section 3.4 of the working paper and of the pivot logic of its Section V. There is no value of ϕ at which the externality is left uncorrected; there is only a continuum along which the corrected party, and the channel of correction, shift from firm to electorate. The consequences for the working paper’s Section V are taken up in §A.8.

A.5 Calibration: the active-conflict ratio as a Pigouvian rate

The Pigouvian benchmark sets the marginal burden equal to the marginal external cost. Per dollar of the firm’s weapons spend, that marginal external cost is the externality multiplier $\lambda(H)$. Setting the ratio equal to it, $\rho(H) = \lambda(H)$, makes the dynamic burden $B(H) = \lambda(H)s(H)$ a Pigouvian correction in ratio form.

The active-conflict value follows from the documented record rather than from choice. World military expenditure reached \$2.89 trillion in 2025, an eleventh consecutive annual rise, while the Institute for Economics and Peace puts the total economic impact of violence in the same year at a record \$21.8 trillion, a gross ratio above seven to one.² Two measurement cautions keep the calibration conservative: the IEP total is purchasing-power adjusted where the expenditure figure is not, and it counts military and security spending itself among its components, so the external cost per dollar of spend (the object λ measures) lies below the gross ratio. The working paper accordingly documents the defensible range at five to seven to one, and the schedule takes the conditional multiplier in active major conflict at the lower bound of that range: $\rho \approx 5$, the 5:1 tier. The figure is therefore not selected for force; it is the conservative end of the documented external-cost ratio, applied as a rate.

Because the multiplier is conditional on conflict intensity (harm per dollar is higher in active war than in stable peace) $\lambda(H)$, and with it $\rho(H)$, rises with H . The tier schedule of Section 3.3 of the working paper (from 0.5:1 at verified long-term peace through 5:1 at active major conflict) is then a step approximation to the continuous Pigouvian schedule $\rho(H) = \lambda(H)$.

Combining $\rho = \lambda$ with the inversion requirement of Proposition A.3 gives the design's empirical content. The schedule supports inversion on the escalation range iff

$$(1 - \varphi) [\lambda'(H)s(H) + \lambda(H)s'(H)] > m'(H), \quad H > H_d.$$

In words: the conditional multiplier and weapons spend must grow, with conflict, at least as fast as the firm's marginal margin. The inequality is testable, and §A.7 records it as a scope condition rather than an assumption.

A.6 On the amplification assumption

Everything in §§A.2–A.3 rests on the externality being amplified: on the producer's marginal return to conflict rising with conflict. This is the binding empirical claim of the framework and the point at which informed objection will concentrate, the more so because the diagnosis of the war economy as a market failure descends from Stiglitz's own work. We state what the assumption requires, weaken it to a condition the evidence supports, and show that where even the weaker condition is in doubt the mechanism survives through a different channel.

The evidence bears on the assumption in three layers, and they do not all point the same way. *Demand and revenue rise sharply and unambiguously.* In 2024 the combined arms revenue of the hundred largest producers reached a record \$679 billion, a rise the Stockholm International Peace Research Institute attributes directly to the wars in Ukraine and Gaza and to elevated geopolitical tension; the response is steepest for firms closest to active conflict, with one Central European shell producer recording a single-year revenue increase of 193 percent and European and Ukrainian

²Stockholm International Peace Research Institute, *Trends in World Military Expenditure, 2025*, SIPRI Fact Sheet (Stockholm: SIPRI, April 2026); Institute for Economics and Peace, *Global Peace Index 2026* (Sydney: IEP, June 2026), reporting \$21.81 trillion in purchasing-power-parity terms, equivalent to 10.5 percent of global GDP.

producers as a group posting increases of 13 and 41 percent, respectively.³ *Asset markets price conflict as raising expected profit.* The cleaner causal signal is the equity response to escalation, a share price being the market's estimate of discounted future profit: an event study of 370 defense and aerospace firms in seventeen countries finds positive abnormal returns for most of them in the days around the February 2022 invasion, strongest among larger firms and firms based in NATO member states;⁴ over the three years that followed, the STOXX Europe Aerospace and Defence index rose by roughly 170 percent, and the most conflict-exposed European producer, Rheinmetall, rose roughly tenfold.⁵ That the repricing strengthens with the severity and proximity of conflict is consistent with a marginal return that rises, rather than merely stays positive, as conflict intensifies, though equity returns alone do not establish it. *Realized margins are the complication.* The structure of procurement works against a convex margin: a large share of major contracting is cost-plus, under which the government reimburses allowable cost and pays a fee that federal acquisition regulation caps at 10 percent of estimated cost for most work and 15 percent for research and development, so per-contract margins are deliberately constrained and do not balloon with demand.⁶ Industry margins have nonetheless drifted upward over two decades (the Department of Defense's own contract finance study records a move from roughly 7–9 percent in 2000–2009 to 11–13 percent in the following decade) and prime contractors are found to out-earn comparable commercial firms on assets employed,⁷ while fixed-price and surge work carries uncapped upside and, as recent multi-billion-dollar write-downs show, uncapped downside.⁸ The honest reading is that the firm's marginal return to conflict operates mainly through volume and mix rather than through an escalating fee on the regulated core.

The strong claim—that the marginal private return to conflict is positive—is therefore well supported. The strict claim of §A.2—that the margin is globally convex—is not established, and is in tension with the cost-plus cap. The framework should not rest on it, and need not.

Definition A.4. Let H_d be the defensive baseline and $[H_d, \bar{H}]$ the escalation range. The externality

³Stockholm International Peace Research Institute, *The SIPRI Top 100 Arms-Producing and Military Services Companies, 2024*, SIPRI Fact Sheet (Stockholm: SIPRI, December 2025). The shell producer is Czechoslovak Group (Czech Republic), whose arms revenue rose 193 percent to \$3.6 billion; arms revenues of European companies outside Russia rose 13 percent to \$151 billion, and those of Ukrainian companies 41 percent to \$3 billion.

⁴ATM Adnan, Md Tapan Mahmud, and Abdullah Al Mamun, “The Impact of Geopolitical Crisis on Defense and Aerospace Stock Performance: Evidence from the Russia–Ukraine War,” *Accounting Research Journal* 39, no. 1 (2026): 122–145.

⁵Reuters, “Analysis: ‘Beyond Fundamentals’: Is Europe’s Arms Race Priced In?,” March 4, 2025 (STOXX Europe Aerospace and Defence index up approximately 170 percent since the invasion); dpa, “Germany’s Rheinmetall Sees Share Price Up Tenfold Since Ukraine War,” February 26, 2025 (Rheinmetall at €968 against €96.8 on February 23, 2022).

⁶Federal Acquisition Regulation 15.404-4(c)(4)(i), implementing 10 U.S.C. § 3322(b) and 41 U.S.C. § 3905(b): under cost-plus-fixed-fee contracts the fee may not exceed 10 percent of estimated cost, or 15 percent for experimental, developmental, or research work.

⁷Department of Defense, Office of Defense Pricing and Contracting, *Defense Contract Finance Study* (Washington: DOD, April 2023), drawing on research by the University of Virginia’s Darden School of Business: defense companies “increased their operating margins from between 7 percent and 9 percent in the 2000–2009 timeframe to between 11 percent and 13 percent in the decade that followed,” and “returns on assets in defense exceed those of the commercial analogs.”

⁸For example, Lockheed Martin recorded \$1.6 billion of program losses in the second quarter of 2025, principally on a classified fixed-price aeronautics program: Lockheed Martin Corporation, “Lockheed Martin Reports Second Quarter 2025 Financial Results,” news release, July 22, 2025.

exhibits *regime amplification* if the firm's marginal return is increasing on the escalation range, $m''(H) > 0$ for $H \in [H_d, \bar{H}]$, with no restriction on m below H_d .

Regime amplification is strictly weaker than the global convexity assumed in §A.2. It permits the marginal margin to be flat or even declining on the defensive and cost-plus core $[0, H_d]$, the behavior the procurement evidence suggests, and asks for acceleration only where it is plausible: in active major conflict, where surge procurement, emergency supplementals, and a shift toward uncapped work raise volume and mix, and so realized profit, faster than proportionally. It is a claim about the slope of profit in the escalation regime, not about the regulated baseline.

Proposition A.6. *Under regime amplification the operational results survive on the escalation range. The static-correction failure of Proposition A.3 obtains for $H > H_d$: a flat charge feasible at the defensive baseline is overtaken by the rising marginal return on $[H_d, \bar{H}]$, so firm lobbying turns back toward conflict in the high-conflict regime. A dynamic ratio with B' increasing and $(1 - \varphi)B'(H) > m'(H)$ on $[H_d, \bar{H}]$ restores inversion throughout that range. Neither conclusion constrains m on $[0, H_d]$.*

Proof. The argument of Proposition A.3 is local to the range on which m' is increasing and is independent of the behavior of m elsewhere; restricting attention to $[H_d, \bar{H}]$ and applying it there gives both claims. □

The case for the dynamic form thus rests on acceleration in the escalation regime alone, which the evidence renders plausible, rather than on convexity everywhere, which it does not.

Suppose even this is doubted, and the cost-plus structure flattens realized profit so thoroughly that the firm's marginal return is near-constant in active conflict. The mechanism does not fail; it relocates. A high cost-plus share is precisely a high pass-through φ : the government bears the cost risk and the marginal burden. By Corollary A.5, as $\varphi \rightarrow 1$ the burden no longer bears on the firm and the correction operates on the procuring government and, through the public budget, on the electorate. The institutional feature that most weakens the firm-margin channel is the same feature that drives the system into the regime where the pass-through channel dominates. The framework is robust across the empirical uncertainty rather than hostage to it: where firm profit is amplified, the lobby-inversion arm of §A.3 operates; where cost-plus insulates it, the budget-legibility arm of §A.4 does. There is no configuration of the margin in which the externality is left uncorrected.

The decisive test is neither revenue nor equity return but realized profit. An estimate of the defense sector's operating profit as a function of conflict intensity, with the cost-plus and fixed-price mix modeled explicitly, would establish whether $m'' > 0$ holds on the escalation range; revenue overstates the response because of the fee cap, and equity returns conflate expected profit with shifting risk premia. Until then, regime amplification is properly carried as a maintained hypothesis—supported, not proven—and the robustness result above is what licenses building on the framework in the meantime.

A.7 Scope conditions and failure modes

Each result above holds under stated conditions; the mechanism is “foolproof” only in the sense that its failure conditions are explicit and bounded. We list them as such.

- **Amplification.** The dynamic-ratio results assume the externality is amplified. Section A.6 weakens this from global convexity to *regime amplification* (acceleration on the escalation range alone), reviews the mixed empirical record, and shows the mechanism is robust even where the weaker condition fails, through the pass-through channel of Corollary A.5. The case rests on the weaker condition, not the strict one.
- **Incidence.** Inversion requires the burden to fall on an actor with a stake in the outcome: $(1 - \varphi) > 0$ for firm inversion, $\varphi > 0$ for government inversion. If pass-through lands on a party that neither internalizes nor is accountable for H (undifferentiated future taxpayers with no present voice, for instance) neither margin inverts. The mechanism is only as strong as the accountability of whoever bears φB .
- **Marginal versus average multiplier.** The calibration $\rho = \lambda$ treats the documented cost ratio as marginal and causal. If the observed ratio is an average, or partly reflects reverse causation, the implied rate is mis-set and the schedule should be re-derived from a marginal estimate.
- **Salience of legibility.** The upper-layer arm of Corollary A.5 assumes that budget entries translate into democratic constraint. If voters do not respond to the now-visible cost, the second layer is priced but not corrected.
- **Jurisdictional completeness.** The model assumes B attaches to the relevant production. Reincorporation or jurisdictional flight (Section 3.4b of the working paper) voids it; market-access conditionality is the assumed closure, and absent it the burden is escapable.
- **Capture is a separate margin.** Inversion redirects lobbying aimed at the *conflict environment*; it does not address lobbying aimed at weakening the certification authority *itself*. That is the distinct task of the anti-capture architecture (Section 3.2 of the working paper), for which the model is no substitute.

A.8 Correspondence with Section V of the working paper

A natural reading of the lobby-inversion result is that the firm bears the burden and is thereby driven to lobby for peace. Corollary A.5 shows that this reading cannot survive the pass-through the paper itself specifies: at high φ the burden no longer bears on the firm, and the corrected actor is the procuring government, through the public budget. Section V of the working paper (“The Economic Logic: Why Compliance Becomes the Path of Least Resistance”) stands on ground the model ratifies, and the correspondence is worth recording precisely.

The paper’s Section 3.1 states the pass-through openly, and in the model’s terms: the obligation is recovered through procurement pricing, the cost of conflict becomes legible in public budgets, and accountability locates with the elected governments that procure. That is Corollary A.5(ii) in prose. The paper’s Section IV describes the Mandate as a dual-layer correction, addressing I_1 by contractor internalization and I_2 through the very pass-through that critics read as a weakness; that is Proposition A.4 and Corollary A.5 exactly, and the model supplies the formal warrant for the paper’s claim that the apparent bug is the second correction. And Section V’s compliance-to-pivot

logic, in which the contractor enters the peace-infrastructure market it is already funding, operates at every value of φ : the market the Mandate creates exists whether or not the firm bears the burden, so the pivot incentive is robust to the incidence question in a way a burden-based lobbying claim alone is not.

One implication of the model deserves emphasis alongside the paper's argument. At high pass-through the firm is indifferent to the *burden*—it suffers no concentrated loss from the Mandate and so has no concentrated interest in fighting the Mandate itself—even though its underlying return to conflict, $m'(H) > 0$, is untouched. Indifference to the instrument is easier to engineer than opposition to conflict and harder to organize against; the residual demand for conflict is then repriced at its democratic source, and the firm's own stake in conflict is the object of the production ceilings and anti-gaming provisions of Section 3.4. Where the paper argues political durability, the low-pass-through case is governed by the lobby-inversion result of §A.3, and the high-pass-through case, which the procurement evidence of §A.6 suggests is typical, by Corollary A.5; the two should be cited together.